



Final Comment Draft for the OpenMP 4.1 Specification

From the OpenMP Language Committee

July 23, 2015

We actively solicit comments. Please provide feedback on this document either to the Editor directly or in the OpenMP Forum at openmp.org

End of Public Comment Period: September 30, 2015

This public draft describes possible future directions or extensions to the OpenMP Specification.

The goal of this public draft is to build more widespread existing practice for an expanded OpenMP. It gives advice on extensions or future directions to those vendors who wish to provide them possibly for trial implementation, allows OpenMP to gather early feedback, support timing and scheduling differences between official OpenMP releases, and offers a preview to users of the future directions of OpenMP with the provision stated in the next paragraph.

This public draft is non-normative. Some of the components in this public draft may be considered for standardization in a future version of OpenMP, but they are not currently part of any OpenMP Specification. Some of the components in this public draft may never be standardized, others may be standardized in a substantially changed form, or it may be standardized as is in its entirety.



OpenMP Application Programming Interface

Version 4.1 rev4 July 2015

Copyright © 1997-2015 OpenMP Architecture Review Board.

Permission to copy without fee all or part of this material is granted, provided the OpenMP Architecture Review Board copyright notice and the title of this document appear. Notice is given that copying is by permission of OpenMP Architecture Review Board.

This page intentionally left blank in published version.

This is Revision 4 (TR3+) (23 July 2015) and includes the following internal tickets applied to the 4.0 LaTeX sources: 190, 267-273, 279-282, 284-286, 288-291, 293, 296-298, 300, 302, 304, 308-311, 314-316, 318, 323-326, 328, 330, 332-336, 338-341, 343-353, 355, 357-360, 362-369, 371, 372, 374-376, 379, 380, 382-386, 388, 390-392, 397, 402, 403

Contents

1	Introduction	1
1.1	Scope	1
1.2	Glossary	2
1.2.1	Threading Concepts	2
1.2.2	OpenMP Language Terminology	2
1.2.3	Loop Terminology	8
1.2.4	Synchronization Terminology	8
1.2.5	Tasking Terminology	9
1.2.6	Data Terminology	10
1.2.7	Implementation Terminology	13
1.3	Execution Model	14
1.4	Memory Model	16
1.4.1	Structure of the OpenMP Memory Model	16
1.4.2	Device Data Environments	18
1.4.3	The Flush Operation	18
1.4.4	OpenMP Memory Consistency	19
1.5	OpenMP Compliance	20
1.6	Normative References	21
1.7	Organization of this Document	23
2	Directives	25
2.1	Directive Format	26
2.1.1	Fixed Source Form Directives	28
2.1.2	Free Source Form Directives	29
2.1.3	Stand-Alone Directives	32
2.2	Conditional Compilation	33
2.2.1	Fixed Source Form Conditional Compilation Sentinels	33

2.2.2	Free Source Form Conditional Compilation Sentinel	34
2.3	Internal Control Variables	35
2.3.1	ICV Descriptions	36
2.3.2	ICV Initialization	37
2.3.3	Modifying and Retrieving ICV Values	39
2.3.4	How ICVs are Scoped	41
2.3.4.1	How the Per-Data Environment ICVs Work	42
2.3.5	ICV Override Relationships	42
2.4	Array Sections	43
2.5	parallel Construct	45
2.5.1	Determining the Number of Threads for a parallel Region	49
2.5.2	Controlling OpenMP Thread Affinity	51
2.6	Canonical Loop Form	52
2.7	Worksharing Constructs	55
2.7.1	Loop Construct	55
2.7.1.1	Determining the Schedule of a Worksharing Loop	62
2.7.2	sections Construct	63
2.7.3	single Construct	66
2.7.4	workshare Construct	68
2.8	SIMD Constructs	71
2.8.1	simd Construct	71
2.8.2	declare simd Construct	75
2.8.3	Loop SIMD Construct	79
2.9	Tasking Constructs	81
2.9.1	task Construct	81
2.9.2	taskloop Construct	85
2.9.3	taskloop simd Construct	89
2.9.4	taskyield Construct	90
2.9.5	Task Scheduling	92
2.10	Device Constructs	93
2.10.1	target data Construct	93
2.10.2	target Construct	95
2.10.3	target update Construct	99

2.10.4	declare target Directive	102
2.10.5	teams Construct	106
2.10.6	distribute Construct	109
2.10.7	distribute simd Construct	111
2.10.8	Distribute Parallel Loop Construct	113
2.10.9	Distribute Parallel Loop SIMD Construct	114
2.10.10	target enter data Construct	116
2.10.11	target exit data Construct	118
2.11	Combined Constructs	120
2.11.1	Parallel Loop Construct	121
2.11.2	parallel sections Construct	122
2.11.3	parallel workshare Construct	124
2.11.4	Parallel Loop SIMD Construct	125
2.11.5	target parallel Construct	126
2.11.6	Target Parallel Loop Construct	127
2.11.7	Target Parallel Loop SIMD Construct	129
2.11.8	target simd Construct	130
2.11.9	target teams Construct	132
2.11.10	teams distribute Construct	133
2.11.11	teams distribute simd Construct	134
2.11.12	target teams distribute Construct	135
2.11.13	target teams distribute simd Construct	137
2.11.14	Teams Distribute Parallel Loop Construct	138
2.11.15	Target Teams Distribute Parallel Loop Construct	139
2.11.16	Teams Distribute Parallel Loop SIMD Construct	141
2.11.17	Target Teams Distribute Parallel Loop SIMD Construct	142
2.12	if Clause	143
2.13	Master and Synchronization Constructs and Clauses	144
2.13.1	master Construct	145
2.13.2	critical Construct	146
2.13.3	barrier Construct	148
2.13.4	taskwait Construct	149
2.13.5	taskgroup Construct	150

2.13.6	atomic Construct	151
2.13.7	flush Construct	158
2.13.8	ordered Construct	163
2.13.9	depend Clause	165
2.14	Cancellation Constructs	168
2.14.1	cancel Construct	168
2.14.2	cancellation point Construct	172
2.15	Data Environment	174
2.15.1	Data-sharing Attribute Rules	175
2.15.1.1	Data-sharing Attribute Rules for Variables Referenced in a Construct	175
2.15.1.2	Data-sharing Attribute Rules for Variables Referenced in a Region but not in a Construct	179
2.15.2	threadprivate Directive	179
2.15.3	Data-Sharing Attribute Clauses	185
2.15.3.1	default Clause	186
2.15.3.2	shared Clause	187
2.15.3.3	private Clause	189
2.15.3.4	firstprivate Clause	193
2.15.3.5	lastprivate Clause	196
2.15.3.6	reduction Clause	198
2.15.3.7	linear Clause	204
2.15.4	Data Copying Clauses	208
2.15.4.1	copyin Clause	208
2.15.4.2	copyprivate Clause	210
2.15.5	Data-mapping Attribute Rules and Clauses	212
2.15.5.1	map Clause	213
2.15.5.2	defaultmap Clause	217
2.16	declare reduction Directive	217
2.17	Nesting of Regions	224
3	Runtime Library Routines	226
3.1	Runtime Library Definitions	227
3.2	Execution Environment Routines	228
3.2.1	omp_set_num_threads	228

3.2.2	<code>omp_get_num_threads</code>	229
3.2.3	<code>omp_get_max_threads</code>	230
3.2.4	<code>omp_get_thread_num</code>	232
3.2.5	<code>omp_get_num_procs</code>	233
3.2.6	<code>omp_in_parallel</code>	233
3.2.7	<code>omp_set_dynamic</code>	234
3.2.8	<code>omp_get_dynamic</code>	236
3.2.9	<code>omp_get_cancellation</code>	237
3.2.10	<code>omp_set_nested</code>	237
3.2.11	<code>omp_get_nested</code>	239
3.2.12	<code>omp_set_schedule</code>	240
3.2.13	<code>omp_get_schedule</code>	242
3.2.14	<code>omp_get_thread_limit</code>	243
3.2.15	<code>omp_set_max_active_levels</code>	243
3.2.16	<code>omp_get_max_active_levels</code>	245
3.2.17	<code>omp_get_level</code>	246
3.2.18	<code>omp_get_ancestor_thread_num</code>	247
3.2.19	<code>omp_get_team_size</code>	248
3.2.20	<code>omp_get_active_level</code>	249
3.2.21	<code>omp_in_final</code>	250
3.2.22	<code>omp_get_proc_bind</code>	251
3.2.23	<code>omp_get_num_places</code>	253
3.2.24	<code>omp_get_place_num_procs</code>	254
3.2.25	<code>omp_get_place_proc_ids</code>	255
3.2.26	<code>omp_get_place_num</code>	256
3.2.27	<code>omp_get_partition_num_places</code>	257
3.2.28	<code>omp_get_partition_place_nums</code>	258
3.2.29	<code>omp_set_default_device</code>	259
3.2.30	<code>omp_get_default_device</code>	260
3.2.31	<code>omp_get_num_devices</code>	260
3.2.32	<code>omp_get_num_teams</code>	261
3.2.33	<code>omp_get_team_num</code>	263
3.2.34	<code>omp_is_initial_device</code>	264

3.2.35	<code>omp_get_initial_device</code>	264
3.2.36	<code>omp_get_max_task_priority</code>	265
3.3	Lock Routines	267
3.3.1	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	269
3.3.2	<code>omp_init_lock_with_hint</code> and <code>omp_init_nest_lock_with_hint</code>	270
3.3.3	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	272
3.3.4	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	273
3.3.5	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	274
3.3.6	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	275
3.4	Timing Routines	276
3.4.1	<code>omp_get_wtime</code>	276
3.4.2	<code>omp_get_wtick</code>	278
3.5	Device Memory Routines	279
3.5.1	<code>omp_target_alloc</code>	279
3.5.2	<code>omp_target_free</code>	280
3.5.3	<code>omp_target_is_present</code>	281
3.5.4	<code>omp_target_memcpy</code>	282
3.5.5	<code>omp_target_memcpy_rect</code>	283
3.5.6	<code>omp_target_associate_ptr</code>	284
3.5.7	<code>omp_target_disassociate_ptr</code>	285
4	Environment Variables	287
4.1	<code>OMP_SCHEDULE</code>	289
4.2	<code>OMP_NUM_THREADS</code>	290
4.3	<code>OMP_DYNAMIC</code>	291
4.4	<code>OMP_PROC_BIND</code>	291
4.5	<code>OMP_PLACES</code>	292
4.6	<code>OMP_NESTED</code>	294
4.7	<code>OMP_STACKSIZE</code>	295
4.8	<code>OMP_WAIT_POLICY</code>	296
4.9	<code>OMP_MAX_ACTIVE_LEVELS</code>	297
4.10	<code>OMP_THREAD_LIMIT</code>	297
4.11	<code>OMP_CANCELLATION</code>	297
4.12	<code>OMP_DISPLAY_ENV</code>	298

4.13	OMP_DEFAULT_DEVICE	299
4.14	OMP_MAX_TASK_PRIORITY	300
A	Stubs for Runtime Library Routines	301
A.1	C/C++ Stub Routines	302
A.2	Fortran Stub Routines	313
B	Interface Declarations	323
B.1	Example of the omp.h Header File	324
B.2	Example of an Interface Declaration include File	328
B.3	Example of a Fortran Interface Declaration module	332
B.4	Example of a Generic Interface for a Library Routine	339
C	OpenMP Implementation-Defined Behaviors	340
D	Features History	345
D.1	Version 4.0 to 4.1 Differences	345
D.2	Version 3.1 to 4.0 Differences	347
D.3	Version 3.0 to 3.1 Differences	348
D.4	Version 2.5 to 3.0 Differences	349
	Index	352

CHAPTER 1

Introduction

The collection of compiler directives, library routines, and environment variables described in this document collectively define the specification of the OpenMP Application Program Interface (OpenMP API) for shared-memory parallelism in C, C++ and Fortran programs.

This specification provides a model for parallel programming that is portable across shared memory architectures from different vendors. Compilers from numerous vendors support the OpenMP API. More information about the OpenMP API can be found at the following web site

<http://www.openmp.org>

The directives, library routines, and environment variables defined in this document allow users to create and manage parallel programs while permitting portability. The directives extend the C, C++ and Fortran base languages with single program multiple data (SPMD) constructs, tasking constructs, device constructs, worksharing constructs, and synchronization constructs, and they provide support for sharing and privatizing data. The functionality to control the runtime environment is provided by library routines and environment variables. Compilers that support the OpenMP API often include a command line option to the compiler that activates and allows interpretation of all OpenMP directives.

1.1 Scope

The OpenMP API covers only user-directed parallelization, wherein the programmer explicitly specifies the actions to be taken by the compiler and runtime system in order to execute the program in parallel. OpenMP-compliant implementations are not required to check for data dependencies, data conflicts, race conditions, or deadlocks, any of which may occur in conforming programs. In addition, compliant implementations are not required to check for code sequences that cause a

program to be classified as non- conforming. Application developers are responsible for correctly using the OpenMP API to produce a conforming program. The OpenMP API does not cover compiler-generated automatic parallelization and directives to the compiler to assist such parallelization.

1.2 Glossary

1.2.1 Threading Concepts

thread An execution entity with a stack and associated static memory, called *threadprivate memory*.

OpenMP thread A *thread* that is managed by the OpenMP runtime system.

thread-safe routine A routine that performs the intended function even when executed concurrently (by more than one *thread*).

processor Implementation defined hardware unit on which one or more *OpenMP threads* can execute.

device An implementation defined logical execution engine.

COMMENT: A *device* could have one or more *processors*.

host device The *device* on which the *OpenMP program* begins execution

target device A device onto which code and data may be offloaded from the *host device*.

1.2.2 OpenMP Language Terminology

base language A programming language that serves as the foundation of the OpenMP specification.

COMMENT: See Section 1.6 on page 21 for a listing of current *base languages* for the OpenMP API.

base program A program written in a *base language*.

structured block For C/C++, an executable statement, possibly compound, with a single entry at the top and a single exit at the bottom, or an OpenMP *construct*.

For Fortran, a block of executable statements with a single entry at the top and a single exit at the bottom, or an OpenMP *construct*.

COMMENTS:

For all *base languages*,

- Access to the *structured block* must not be the result of a branch.
- The point of exit cannot be a branch out of the *structured block*.

For C/C++:

- The point of entry must not be a call to **setjmp()**.
- **longjmp()** and **throw()** must not violate the entry/exit criteria.
- Calls to **exit()** are allowed in a *structured block*.
- An expression statement, iteration statement, selection statement, or try block is considered to be a *structured block* if the corresponding compound statement obtained by enclosing it in { and } would be a *structured block*.

For Fortran:

- **STOP** statements are allowed in a *structured block*.

enclosing context In C/C++, the innermost scope enclosing an OpenMP *directive*.

In Fortran, the innermost scoping unit enclosing an OpenMP *directive*.

directive In C/C++, a **#pragma**, and in Fortran, a comment, that specifies *OpenMP program* behavior.

COMMENT: See Section 2.1 on page 26 for a description of OpenMP *directive* syntax.

white space A non-empty sequence of space and/or horizontal tab characters

OpenMP program A program that consists of a *base program*, annotated with OpenMP *directives* and runtime library routines.

conforming program An *OpenMP program* that follows all the rules and restrictions of the OpenMP specification.

declarative directive An OpenMP *directive* that may only be placed in a declarative context. A *declarative directive* results in one or more declarations only; it is not associated with the immediate execution of any user code.

1	executable directive	An OpenMP <i>directive</i> that is not declarative. That is, it may be placed in an
2		executable context.
3	stand-alone directive	An OpenMP <i>executable directive</i> that has no associated executable user code.
4	construct	An OpenMP <i>executable directive</i> (and for Fortran, the paired end directive , if any)
5		and the associated statement, loop or <i>structured block</i> , if any, not including the code
6		in any called routines. That is, in the lexical extent of an <i>executable directive</i> .
7	combined construct	A construct that is a shortcut for specifying one construct immediately nested inside
8		another construct. A combined construct is semantically identical to that of explicitly
9		specifying the first construct containing one instance of the second construct and no
10		other statements.
11	composite construct	A construct that is composed of two constructs but does not have identical semantics
12		to specifying one of the constructs immediately nested inside the other. A composite
13		construct either adds semantics not included in the constructs from which it is
14		composed or the nesting of the one construct inside the other is not conforming.
15	region	All code encountered during a specific instance of the execution of a given <i>construct</i>
16		or of an OpenMP library routine. A <i>region</i> includes any code in called routines as
17		well as any implicit code introduced by the OpenMP implementation. The generation
18		of a <i>task</i> at the point where a task directive is encountered is a part of the <i>region</i> of
19		the <i>encountering thread</i> , but the <i>explicit task region</i> associated with the task
20		<i>directive</i> is not. The point where a target or teams directive is encountered is a
21		part of the <i>region</i> of the <i>encountering thread</i> , but the <i>region</i> associated with the
22		target or teams directive is not.
23		COMMENTS:
24		A <i>region</i> may also be thought of as the dynamic or runtime extent of a
25		<i>construct</i> or of an OpenMP library routine.
26		During the execution of an <i>OpenMP program</i> , a <i>construct</i> may give rise to
27		many <i>regions</i> .
28	active parallel region	A parallel <i>region</i> that is executed by a <i>team</i> consisting of more than one <i>thread</i> .
29	inactive parallel region	A parallel <i>region</i> that is executed by a <i>team</i> of only one <i>thread</i> .
30	sequential part	All code encountered during the execution of an <i>initial task region</i> that is not part of
31		a parallel <i>region</i> corresponding to a parallel <i>construct</i> or a task <i>region</i>
32		corresponding to a task <i>construct</i> .
33		COMMENTS:
34		A <i>sequential part</i> is enclosed by an <i>implicit parallel region</i> .

1		Executable statements in called routines may be in both a <i>sequential part</i>
2		and any number of explicit parallel <i>regions</i> at different points in the
3		program execution.
4	master thread	The <i>thread</i> that encounters a parallel <i>construct</i> , creates a <i>team</i> , generates a set of
5		<i>implicit tasks</i> , then executes one of those <i>tasks</i> as <i>thread</i> number 0.
6	parent thread	The <i>thread</i> that encountered the parallel <i>construct</i> and generated a parallel
7		<i>region</i> is the <i>parent thread</i> of each of the <i>threads</i> in the <i>team</i> of that parallel
8		<i>region</i> . The <i>master thread</i> of a parallel <i>region</i> is the same <i>thread</i> as its <i>parent</i>
9		<i>thread</i> with respect to any resources associated with an <i>OpenMP thread</i> .
10	child thread	When a <i>thread</i> encounters a parallel <i>construct</i> , each of the <i>threads</i> in the
11		generated parallel <i>region</i> ’s <i>team</i> are <i>child threads</i> of the encountering <i>thread</i> .
12		The target or teams <i>region</i> ’s <i>initial thread</i> is not a <i>child thread</i> of the <i>thread</i> that
13		encountered the target or teams <i>construct</i> .
14	ancestor thread	For a given <i>thread</i> , its <i>parent thread</i> or one of its <i>parent thread</i> ’s <i>ancestor threads</i> .
15	descendent thread	For a given <i>thread</i> , one of its <i>child threads</i> or one of its <i>child threads</i> ’ <i>descendent</i>
16		<i>threads</i> .
17	team	A set of one or more <i>threads</i> participating in the execution of a parallel <i>region</i> .
18		COMMENTS:
19		For an <i>active parallel region</i> , the <i>team</i> comprises the <i>master thread</i> and at
20		least one additional <i>thread</i> .
21		For an <i>inactive parallel region</i> , the <i>team</i> comprises only the <i>master thread</i> .
22	league	The set of <i>thread teams</i> created by a target <i>construct</i> or a teams <i>construct</i> .
23	contention group	An <i>initial thread</i> and its <i>descendent threads</i> .
24	implicit parallel region	An <i>inactive parallel region</i> that generates an <i>initial task region</i> . <i>Implicit parallel</i>
25		<i>regions</i> surround the whole OpenMP program, all target <i>regions</i> , and all teams
26		<i>regions</i>
27	initial thread	A <i>thread</i> that executes an <i>implicit parallel region</i> .
28	nested construct	A <i>construct</i> (lexically) enclosed by another <i>construct</i> .
29	closely nested construct	A <i>construct</i> nested inside another <i>construct</i> with no other <i>construct</i> nested between
30		them.
31	nested region	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> encountered
32		during the execution of another <i>region</i> .
33		COMMENT: Some nestings are <i>conforming</i> and some are not. See
34		Section 2.17 on page 224 for the restrictions on nesting.

1	closely nested region	A <i>region</i> nested inside another <i>region</i> with no parallel <i>region</i> nested between
2		them.
3	strictly closely nested	A <i>region</i> nested inside another <i>region</i> with no parallel , teams or target
4	region	<i>region</i> nested between them.
5	all threads	All OpenMP <i>threads</i> participating in the <i>OpenMP</i> program.
6	current team	All <i>threads</i> in the <i>team</i> executing the innermost enclosing parallel <i>region</i> .
7	encountering thread	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
8	all tasks	All <i>tasks</i> participating in the <i>OpenMP</i> program.
9	current team tasks	All <i>tasks</i> encountered by the corresponding <i>team</i> . Note that the <i>implicit tasks</i>
10		constituting the parallel <i>region</i> and any <i>descendent tasks</i> encountered during the
11		execution of these <i>implicit tasks</i> are included in this set of tasks.
12	generating task	For a given <i>region</i> , the task whose execution by a <i>thread</i> generated the <i>region</i> .
13	binding thread set	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a
14		<i>region</i> .
15		The <i>binding thread</i> set for a given <i>region</i> can be <i>all threads</i> on a <i>device</i> , <i>all threads</i>
16		in a <i>contention group</i> , the <i>current team</i> , or the <i>encountering thread</i> .
17		COMMENT: The <i>binding thread</i> set for a particular <i>region</i> is described in
18		its corresponding subsection of this specification.
19	binding task set	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a
20		<i>region</i> .
21		The <i>binding task</i> set for a given <i>region</i> can be <i>all tasks</i> , the <i>current team tasks</i> , or the
22		<i>generating task</i> .
23		COMMENT: The <i>binding task</i> set for a particular <i>region</i> (if applicable) is
24		described in its corresponding subsection of this specification.

1	binding region	The enclosing <i>region</i> that determines the execution context and limits the scope of
2		the effects of the bound <i>region</i> is called the <i>binding region</i> .
3		<i>Binding region</i> is not defined for <i>regions</i> whose <i>binding thread</i> set is <i>all threads</i> or
4		the <i>encountering thread</i> , nor is it defined for <i>regions</i> whose <i>binding task set</i> is <i>all</i>
5		<i>tasks</i> .
6		COMMENTS:
7		The <i>binding region</i> for an ordered <i>region</i> is the innermost enclosing
8		<i>loop region</i> .
9		The <i>binding region</i> for a taskwait <i>region</i> is the innermost enclosing
10		<i>task region</i> .
11		The <i>binding region</i> for a cancel <i>region</i> is the innermost enclosing
12		<i>region</i> corresponding to the <i>construct-type-clause</i> of the cancel
13		construct.
14		The <i>binding region</i> for a cancellation point <i>region</i> is the
15		innermost enclosing <i>region</i> corresponding to the <i>construct-type-clause</i> of
16		the cancellation point construct.
17		For all other <i>regions</i> for which the <i>binding thread set</i> is the <i>current team</i>
18		or the <i>binding task set</i> is the <i>current team tasks</i> , the <i>binding region</i> is the
19		innermost enclosing parallel <i>region</i> .
20		For <i>regions</i> for which the <i>binding task set</i> is the <i>generating task</i> , the
21		<i>binding region</i> is the <i>region</i> of the <i>generating task</i> .
22		A parallel <i>region</i> need not be <i>active</i> nor explicit to be a <i>binding</i>
23		<i>region</i> .
24		A <i>task region</i> need not be explicit to be a <i>binding region</i> .
25		A <i>region</i> never binds to any <i>region</i> outside of the innermost enclosing
26		parallel <i>region</i> .
27	orphaned construct	A <i>construct</i> that gives rise to a <i>region</i> whose <i>binding thread set</i> is the <i>current team</i> ,
28		but is not nested within another <i>construct</i> giving rise to the <i>binding region</i> .
29	worksharing construct	A <i>construct</i> that defines units of work, each of which is executed exactly once by one
30		of the <i>threads</i> in the <i>team</i> executing the <i>construct</i> .
31		For C/C++, <i>worksharing constructs</i> are for , sections , and single .
32		For Fortran, <i>worksharing constructs</i> are do , sections , single and
33		workshare .

place

1		Unordered set of <i>processors</i> that is treated by the execution environment as a location
2		unit when dealing with OpenMP thread affinity.
3	place list	The ordered list that describes all OpenMP <i>places</i> available to the execution
4		environment.
5	place partition	An ordered list that corresponds to a contiguous interval in the OpenMP <i>place list</i> . It
6		describes the <i>places</i> currently available to the execution environment for a given
7		parallel region.
8	place number	A number that uniquely identifies a <i>place</i> in the <i>place list</i> , with zero identifying the
9		first <i>place</i> in the <i>place list</i> , and each consecutive whole number identifying the next
10		<i>place</i> in the <i>place list</i> .
11	SIMD instruction	A single machine instruction that can operate on multiple data elements.
12	SIMD lane	A software or hardware mechanism capable of processing one data element from a
13		<i>SIMD instruction</i> .
14	SIMD chunk	A set of iterations executed concurrently, each by a <i>SIMD lane</i> , by a single <i>thread</i> by
15		means of <i>SIMD instructions</i> .

16 1.2.3 Loop Terminology

17	loop directive	An OpenMP <i>executable</i> directive whose associated user code must be a loop nest
18		that is a <i>structured block</i> .
19	associated loop(s)	The loop(s) controlled by a <i>loop directive</i> .
20		COMMENT: If the <i>loop directive</i> contains a collapse or an
21		ordered (n) clause then there may be more than one <i>associated loop</i> .
22	sequential loop	A loop that is not associated with any OpenMP <i>loop directive</i> .
23	SIMD loop	A loop that includes at least one <i>SIMD chunk</i> .
24	doacross loop nest	A loop nest that has cross-iteration dependence. An iteration is dependent on one or
25		more lexicographically earlier iterations.
26		COMMENT: The loop directive with an ordered clause with the
27		parameter identifies the loop(s) associated for the <i>doacross loop nest</i> .

28 1.2.4 Synchronization Terminology

1	barrier	A point in the execution of a program encountered by a <i>team</i> of <i>threads</i> , beyond
2		which no <i>thread</i> in the team may execute until all <i>threads</i> in the <i>team</i> have reached
3		the barrier and all <i>explicit tasks</i> generated by the <i>team</i> have executed to completion.
4		If <i>cancellation</i> has been requested, threads may proceed to the end of the canceled
5		<i>region</i> even if some threads in the team have not reached the <i>barrier</i> .
6	cancellation	An action that cancels (that is, aborts) an OpenMP <i>region</i> and causes executing
7		<i>implicit</i> or <i>explicit</i> tasks to proceed to the end of the canceled <i>region</i> .
8	cancellation point	A point at which implicit and explicit tasks check if cancellation has been requested.
9		If cancellation has been observed, they perform the <i>cancellation</i> .
10		COMMENT: For a list of cancellation points, see Section 2.14.1 on
11		page 168

12 1.2.5 Tasking Terminology

13	task	A specific instance of executable code and its <i>data environment</i> , generated when a
14		<i>thread</i> encounters a task construct or a parallel construct .
15	task region	A <i>region</i> consisting of all code encountered during the execution of a <i>task</i> .
16		COMMENT: A parallel region consists of one or more implicit <i>task</i>
17		<i>regions</i> .
18	explicit task	A <i>task</i> generated when a task construct is encountered during execution.
19	implicit task	A <i>task</i> generated by an <i>implicit parallel region</i> or generated when a parallel
20		<i>construct</i> is encountered during execution.
21	initial task	An <i>implicit task</i> associated with an <i>implicit parallel region</i> .
22	current task	For a given <i>thread</i> , the <i>task</i> corresponding to the <i>task region</i> in which it is executing.
23	child task	A <i>task</i> is a <i>child task</i> of its generating <i>task region</i> . A <i>child task region</i> is not part of
24		its generating <i>task region</i> .
25	sibling tasks	<i>Tasks</i> that are <i>child tasks</i> of the same <i>task region</i> .
26	descendent task	A <i>task</i> that is the <i>child task</i> of a <i>task region</i> or of one of its <i>descendent task regions</i> .
27	task completion	<i>Task completion</i> occurs when the end of the <i>structured block</i> associated with the
28		<i>construct</i> that generated the <i>task</i> is reached.
29		COMMENT: Completion of the <i>initial task</i> occurs at program exit.

1	task scheduling point	A point during the execution of the current <i>task region</i> at which it can be suspended
2		to be resumed later; or the point of <i>task completion</i> , after which the executing thread
3		may switch to a different <i>task region</i> .
4		COMMENT: For a list of task scheduling points, see Section 2.9.5 on
5		page 92.
6	task switching	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
7	tied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same
8		<i>thread</i> that suspended it. That is, the <i>task</i> is tied to that <i>thread</i> .
9	untied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in the
10		team. That is, the <i>task</i> is not tied to any <i>thread</i> .
11	undelayed task	A <i>task</i> for which execution is not deferred with respect to its generating <i>task region</i> .
12		That is, its generating <i>task region</i> is suspended until execution of the <i>undelayed task</i>
13		is completed.
14	included task	A <i>task</i> for which execution is sequentially included in the generating <i>task region</i> .
15		That is, an <i>included task</i> is <i>undelayed</i> and executed immediately by the <i>encountering</i>
16		<i>thread</i> .
17	merged task	A <i>task</i> whose <i>data environment</i> , inclusive of ICVs, is the same as that of its
18		generating <i>task region</i> .
19	final task	A <i>task</i> that forces all of its <i>child tasks</i> to become <i>final</i> and <i>included tasks</i> .
20	task dependence	An ordering relation between two <i>sibling tasks</i> : the <i>dependent task</i> and a previously
21		generated <i>predecessor task</i> . The <i>task dependence</i> is fulfilled when the <i>predecessor</i>
22		<i>task</i> has completed.
23	dependent task	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor</i>
24		<i>tasks</i> have completed.
25	predecessor task	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
26	task synchronization construct	A taskwait , taskgroup , or a barrier construct.
27	target task	A <i>merged task</i> that is executed immediately.
28	taskgroup set	A set of tasks that are logically grouped by a taskgroup region.

29 1.2.6 Data Terminology

1	variable	A named data storage block, whose value can be defined and redefined during the
2		execution of a program.
<div>▼</div>		
3	Note – An array or structure element is a variable that is part of another variable.	
<div>▲</div>		
4	scalar variable	For C/C++: A scalar variable.
5		For Fortran: A scalar variable with intrinsic type, excluding character type.
6	array section	A designated subset of the elements of an array.
7	array item	An array, an array section or an array element.
8	structure	A structure is a variable that contains one or more variables. Structures are
9		implemented using structs in C or C++, classes in C++, and derived types in Fortran.
10	private variable	With respect to a given set of <i>task regions</i> or <i>SIMD lanes</i> that bind to the same
11		parallel region , a <i>variable</i> whose name provides access to a different block of
12		storage for each <i>task region</i> or <i>SIMD lane</i> .
13		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
14		made private independently of other components.
15	shared variable	With respect to a given set of <i>task regions</i> that bind to the same parallel region , a
16		<i>variable</i> whose name provides access to the same block of storage for each <i>task</i>
17		<i>region</i> .
18		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
19		<i>shared</i> independently of the other components, except for static data members of
20		C++ classes.
21	threadprivate variable	A <i>variable</i> that is replicated, one instance per <i>thread</i> , by the OpenMP
22		implementation. Its name then provides access to a different block of storage for each
23		<i>thread</i> .
24		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
25		<i>threadprivate</i> independently of the other components, except for static data
26		members of C++ classes.
27	threadprivate memory	The set of <i>threadprivate variables</i> associated with each <i>thread</i> .
28	data environment	The <i>variables</i> associated with the execution of a given <i>region</i> .
29	device data environment	The initial <i>data environment</i> associated with a device.
30	device address	An <i>implementation defined</i> reference to an address in a <i>device data environment</i> .

1	device pointer	A <i>variable</i> that contains a <i>device address</i> .
2	mapped variable	An original <i>variable</i> in a <i>data environment</i> with a corresponding <i>variable</i> in a device
3		<i>data environment</i> .
4		COMMENT: The original and corresponding <i>variables</i> may share storage.
5	mappable type	A type that is valid for a <i>mapped variable</i> . If a type is composed from other types
6		(such as the type of an array or structure element) and any of the other types are not
7		mappable then the type is not mappable.
8		COMMENT: Pointer types are <i>mappable</i> but the memory block to which
9		the pointer refers is not <i>mapped</i> .
10		For C: The type must be a complete type.
11		For C++: The type must be a complete type.
12		In addition, for class types:
13		• All member functions accessed in any target region must appear in a
14		declare target directive.
15		• All data members must be non-static.
16		• A <i>mappable type</i> cannot contain virtual members.
17		For Fortran: The type must be definable.
18		In addition, for derived types:
19		• All type-bound procedures accessed in any target region must appear in a declare
20		target directive.
21	defined	For <i>variables</i> , the property of having a valid value.
22		For C: For the contents of <i>variables</i> , the property of having a valid value.
23		For C++: For the contents of <i>variables</i> of POD (plain old data) type, the property of
24		having a valid value.
25		For <i>variables</i> of non-POD class type, the property of having been constructed but not
26		subsequently destructed.
27		For Fortran: For the contents of <i>variables</i> , the property of having a valid value. For
28		the allocation or association status of <i>variables</i> , the property of having a valid status.
29		COMMENT: Programs that rely upon <i>variables</i> that are not <i>defined</i> are
30		<i>non-conforming programs</i> .
31	class type	For C++: <i>Variables</i> declared with one of the class , struct , or union keywords

1	sequentially consistent atomic construct	An atomic construct for which the seq_cst clause is specified.
2	non-sequentially consistent atomic construct	An atomic construct for which the seq_cst clause is not specified

3 1.2.7 Implementation Terminology

4	supporting n levels of parallelism	Implies allowing an <i>active parallel region</i> to be enclosed by $n-1$ <i>active parallel regions</i> .
5		
6	supporting the OpenMP API	Supporting at least one level of parallelism.
7	supporting nested parallelism	Supporting more than one level of parallelism.
8	internal control variable	A conceptual variable that specifies runtime behavior of a set of <i>threads</i> or <i>tasks</i> in an <i>OpenMP program</i> .
9		
10		COMMENT: The acronym ICV is used interchangeably with the term
11		<i>internal control variable</i> in the remainder of this specification.
12	compliant implementation	An implementation of the OpenMP specification that compiles and executes any <i>conforming program</i> as defined by the specification.
13		
14		COMMENT: A <i>compliant implementation</i> may exhibit <i>unspecified behavior</i> when compiling or executing a <i>non-conforming program</i> .
15		
16	unspecified behavior	A behavior or result that is not specified by the OpenMP specification or not known prior to the compilation or execution of an <i>OpenMP program</i> .
17		
18		Such <i>unspecified behavior</i> may result from:
19		• Issues documented by the OpenMP specification as having <i>unspecified behavior</i> .
20		• A <i>non-conforming program</i> .
21		• A <i>conforming program</i> exhibiting an <i>implementation defined behavior</i> .
22	implementation defined	Behavior that must be documented by the implementation, and is allowed to vary among different <i>compliant implementations</i> . An implementation is allowed to define this behavior as <i>unspecified</i> .
23		
24		
25		COMMENT: All features that have <i>implementation defined behavior</i> are
26		documented in Appendix C.

deprecated Implies a construct, clause or other feature is normative in the current specification but is considered obsolescent and will be removed in the future.

1.3 Execution Model

The OpenMP API uses the fork-join model of parallel execution. Multiple threads of execution perform tasks defined implicitly or explicitly by OpenMP directives. The OpenMP API is intended to support programs that will execute correctly both as parallel programs (multiple threads of execution and a full OpenMP support library) and as sequential programs (directives ignored and a simple OpenMP stubs library). However, it is possible and permitted to develop a program that executes correctly as a parallel program but not as a sequential program, or that produces different results when executed as a parallel program compared to when it is executed as a sequential program. Furthermore, using different numbers of threads may result in different numeric results because of changes in the association of numeric operations. For example, a serial addition reduction may have a different pattern of addition associations than a parallel reduction. These different associations may change the results of floating-point addition.

An OpenMP program begins as a single thread of execution, called an initial thread. An initial thread executes sequentially, as if enclosed in an implicit task region, called an initial task region, that is defined by the implicit parallel region surrounding the whole program.

The thread that executes the implicit parallel region that surrounds the whole program executes on the *host device*. An implementation may support other *target devices*. If supported, one or more devices are available to the host device for offloading code and data. Each device has its own threads that are distinct from threads that execute on another device. Threads cannot migrate from one device to another device. The execution model is host-centric such that the host device offloads **target** regions to target devices.

The initial thread that executes the implicit parallel region that surrounds the **target** region may execute on a *target device*. An initial thread executes sequentially, as if enclosed in an implicit task region, called an initial task region, that is defined by an implicit inactive **parallel** region that surrounds the entire **target** region.

When a **target** construct is encountered, the **target** region is executed by the implicit device task. The task that encounters the **target** construct waits at the end of the construct until execution of the region completes. If the target device does not exist or the implementation does not support the target device, all **target** regions associated with that device are executed by the host device.

The implementation must ensure that the **target** region executes as if it were executed in the data environment of the target device unless an **if** clause is present and the **if** clause expression evaluates to *false*.

The **teams** construct creates a *league of thread teams* where the master thread of each team executes the region. Each of these master threads is an initial thread, and executes sequentially, as if enclosed in an implicit task region that is defined by an implicit parallel region that surrounds the entire **teams** region.

If a construct creates a data environment, the data environment is created at the time the construct is encountered. Whether a construct creates a data environment is defined in the description of the construct.

When any thread encounters a **parallel** construct, the thread creates a team of itself and zero or more additional threads and becomes the master of the new team. A set of implicit tasks, one per thread, is generated. The code for each task is defined by the code inside the **parallel** construct. Each task is assigned to a different thread in the team and becomes tied; that is, it is always executed by the thread to which it is initially assigned. The task region of the task being executed by the encountering thread is suspended, and each member of the new team executes its implicit task. There is an implicit barrier at the end of the **parallel** construct. Only the master thread resumes execution beyond the end of the **parallel** construct, resuming the task region that was suspended upon encountering the **parallel** construct. Any number of **parallel** constructs can be specified in a single program.

parallel regions may be arbitrarily nested inside each other. If nested parallelism is disabled, or is not supported by the OpenMP implementation, then the new team that is created by a thread encountering a **parallel** construct inside a **parallel** region will consist only of the encountering thread. However, if nested parallelism is supported and enabled, then the new team can consist of more than one thread. A **parallel** construct may include a **proc_bind** clause to specify the places to use for the threads in the team within the **parallel** region.

When any team encounters a worksharing construct, the work inside the construct is divided among the members of the team, and executed cooperatively instead of being executed by every thread. There is a default barrier at the end of each worksharing construct unless the **nowait** clause is present. Redundant execution of code by every thread in the team resumes after the end of the worksharing construct.

When any thread encounters a **task** construct, a new explicit task is generated. Execution of explicitly generated tasks is assigned to one of the threads in the current team, subject to the thread's availability to execute work. Thus, execution of the new task could be immediate, or deferred until later according to task scheduling constraints and thread availability. Threads are allowed to suspend the current task region at a task scheduling point in order to execute a different task. If the suspended task region is for a tied task, the initially assigned thread later resumes execution of the suspended task region. If the suspended task region is for an untied task, then any thread may resume its execution. Completion of all explicit tasks bound to a given parallel region is guaranteed before the master thread leaves the implicit barrier at the end of the region. Completion of a subset of all explicit tasks bound to a given parallel region may be specified through the use of task synchronization constructs. Completion of all explicit tasks bound to the implicit parallel region is guaranteed by the time the program exits.

When any thread encounters a **simd** construct, the iterations of the loop associated with the construct may be executed concurrently using the SIMD lanes that are available to the thread.

The **cancel** construct can alter the previously described flow of execution in an OpenMP region. The effect of the **cancel** construct depends on its *construct-type-clause*. If a task encounters a **cancel** construct with a **taskgroup** *construct-type-clause*, then the task activates cancellation and continues execution at the end of its **task** region, which implies completion of that task. Any other task in that **taskgroup** that has begun executing completes execution unless it encounters a **cancellation point** construct, in which case it continues execution at the end of its **task** region, which implies its completion. Other tasks in that **taskgroup** region that have not begun execution are aborted, which implies their completion.

For all other *construct-type-clause* values, if a thread encounters a **cancel** construct, it activates cancellation of the innermost enclosing region of the type specified and the thread continues execution at the end of that region. Threads check if cancellation has been activated for their region at cancellation points and, if so, also resume execution at the end of the canceled region.

If cancellation has been activated regardless of *construct-type-clause*, threads that are waiting inside a barrier other than an implicit barrier at the end of the canceled region exit the barrier and resume execution at the end of the canceled region. This action can occur before the other threads reach that barrier.

Synchronization constructs and library routines are available in the OpenMP API to coordinate tasks and data access in **parallel** regions. In addition, library routines and environment variables are available to control or to query the runtime environment of OpenMP programs.

The OpenMP specification makes no guarantee that input or output to the same file is synchronous when executed in parallel. In this case, the programmer is responsible for synchronizing input and output statements (or routines) using the provided synchronization constructs or library routines. For the case where each thread accesses a different file, no synchronization by the programmer is necessary.

1.4 Memory Model

1.4.1 Structure of the OpenMP Memory Model

The OpenMP API provides a relaxed-consistency, shared-memory model. All OpenMP threads have access to a place to store and to retrieve variables, called the *memory*. In addition, each thread is allowed to have its own *temporary view* of the memory. The temporary view of memory for each thread is not a required part of the OpenMP memory model, but can represent any kind of intervening structure, such as machine registers, cache, or other local storage, between the thread

and the memory. The temporary view of memory allows the thread to cache variables and thereby to avoid going to memory for every reference to a variable. Each thread also has access to another type of memory that must not be accessed by other threads, called *threadprivate memory*.

A directive that accepts data-sharing attribute clauses determines two kinds of access to variables used in the directive's associated structured block: shared and private. Each variable referenced in the structured block has an original variable, which is the variable by the same name that exists in the program immediately outside the construct. Each reference to a shared variable in the structured block becomes a reference to the original variable. For each private variable referenced in the structured block, a new version of the original variable (of the same type and size) is created in memory for each task or SIMD lane that contains code associated with the directive. Creation of the new version does not alter the value of the original variable. However, the impact of attempts to access the original variable during the region associated with the directive is unspecified; see Section 2.15.3.3 on page 189 for additional details. References to a private variable in the structured block refer to the private version of the original variable for the current task or SIMD lane. The relationship between the value of the original variable and the initial or final value of the private version depends on the exact clause that specifies it. Details of this issue, as well as other issues with privatization, are provided in Section 2.15 on page 174.

The minimum size at which a memory update may also read and write back adjacent variables that are part of another variable (as array or structure elements) is implementation defined but is no larger than required by the base language.

A single access to a variable may be implemented with multiple load or store instructions, and hence is not guaranteed to be atomic with respect to other accesses to the same variable. Accesses to variables smaller than the implementation defined minimum size or to C or C++ bit-fields may be implemented by reading, modifying, and rewriting a larger unit of memory, and may thus interfere with updates of variables or fields in the same unit of memory.

If multiple threads write without synchronization to the same memory unit, including cases due to atomicity considerations as described above, then a data race occurs. Similarly, if at least one thread reads from a memory unit and at least one thread writes without synchronization to that same memory unit, including cases due to atomicity considerations as described above, then a data race occurs. If a data race occurs then the result of the program is unspecified.

A private variable in a task region that eventually generates an inner nested **parallel** region is permitted to be made shared by implicit tasks in the inner **parallel** region. A private variable in a task region can be shared by an explicit **task** region generated during its execution. However, it is the programmer's responsibility to ensure through synchronization that the lifetime of the variable does not end before completion of the explicit **task** region sharing it. Any other access by one task to the private variables of another task results in unspecified behavior.

1 1.4.2 Device Data Environments

2 When an OpenMP program begins, an implicit **target data** region for each device surrounds
3 the whole program. Each device has a device data environment that is defined by its implicit
4 **target data** region. Any **declare target** directives and the directives that accept
5 data-mapping attribute clauses determine how an original variable in a data environment is mapped
6 to a corresponding variable in a device data environment.

7 When an original variable is mapped to a device data environment and the associated
8 corresponding variable is not present in the device data environment, a new corresponding variable
9 (of the same type and size as the original variable) is created in the device data environment. The
10 initial value of the new corresponding variable is determined from the clauses and the data
11 environment of the encountering thread.

12 The corresponding variable in the device data environment may share storage with the original
13 variable. Writes to the corresponding variable may alter the value of the original variable. The
14 impact of this on memory consistency is discussed in Section 1.4.4 on page 19. When a task
15 executes in the context of a device data environment, references to the original variable refer to the
16 corresponding variable in the device data environment.

17 The relationship between the value of the original variable and the initial or final value of the
18 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with
19 mapping a variable, are provided in Section 2.15.5.1 on page 213.

20 The original variable in a data environment and the corresponding variable(s) in one or more device
21 data environments may share storage. Without intervening synchronization data races can occur.

22 1.4.3 The Flush Operation

23 The memory model has relaxed-consistency because a thread's temporary view of memory is not
24 required to be consistent with memory at all times. A value written to a variable can remain in the
25 thread's temporary view until it is forced to memory at a later time. Likewise, a read from a variable
26 may retrieve the value from the thread's temporary view, unless it is forced to read from memory.
27 The OpenMP flush operation enforces consistency between the temporary view and memory.

28 The flush operation is applied to a set of variables called the *flush-set*. The flush operation restricts
29 reordering of memory operations that an implementation might otherwise do. Implementations
30 must not reorder the code for a memory operation for a given variable, or the code for a flush
31 operation for the variable, with respect to a flush operation that refers to the same variable.

32 If a thread has performed a write to its temporary view of a shared variable since its last flush of
33 that variable, then when it executes another flush of the variable, the flush does not complete until
34 the value of the variable has been written to the variable in memory. If a thread performs multiple

writes to the same variable between two flushes of that variable, the flush ensures that the value of the last write is written to the variable in memory. A flush of a variable executed by a thread also causes its temporary view of the variable to be discarded, so that if its next memory operation for that variable is a read, then the thread will read from memory when it may again capture the value in the temporary view. When a thread executes a flush, no later memory operation by that thread for a variable involved in that flush is allowed to start until the flush completes. The completion of a flush of a set of variables executed by a thread is defined as the point at which all writes to those variables performed by the thread before the flush are visible in memory to all other threads and that thread's temporary view of all variables involved is discarded.

The flush operation provides a guarantee of consistency between a thread's temporary view and memory. Therefore, the flush operation can be used to guarantee that a value written to a variable by one thread may be read by a second thread. To accomplish this, the programmer must ensure that the second thread has not written to the variable since its last flush of the variable, and that the following sequence of events happens in the specified order:

1. The value is written to the variable by the first thread.
2. The variable is flushed by the first thread.
3. The variable is flushed by the second thread.
4. The value is read from the variable by the second thread.

Note – OpenMP synchronization operations, described in Section 2.13 on page 144 and in Section 3.3 on page 267, are recommended for enforcing this order. Synchronization through variables is possible but is not recommended because the proper timing of flushes is difficult.

1.4.4 OpenMP Memory Consistency

The restrictions in Section 1.4.3 on page 18 on reordering with respect to flush operations guarantee the following:

- If the intersection of the flush-sets of two flushes performed by two different threads is non-empty, then the two flushes must be completed as if in some sequential order, seen by all threads.
- If two operations performed by the same thread either access, modify, or flush the same variable, then they must be completed as if in that thread's program order, as seen by all threads.
- If the intersection of the flush-sets of two flushes is empty, the threads can observe these flushes in any order.

The flush operation can be specified using the **flush** directive, and is also implied at various locations in an OpenMP program: see Section 2.13.7 on page 158 for details.

Note – Since flush operations by themselves cannot prevent data races, explicit flush operations are only useful in combination with non-sequentially consistent atomic directives.

OpenMP programs that:

- do not use non-sequentially consistent atomic directives,
- do not rely on the accuracy of a *false* result from **omp_test_lock** and **omp_test_nest_lock**, and
- correctly avoid data races as required in Section 1.4.1 on page 16

behave as though operations on shared variables were simply interleaved in an order consistent with the order in which they are performed by each thread. The relaxed consistency model is invisible for such programs, and any explicit flush operations in such programs are redundant.

Implementations are allowed to relax the ordering imposed by implicit flush operations when the result is only visible to programs using non-sequentially consistent atomic directives.

1.5 OpenMP Compliance

An implementation of the OpenMP API is compliant if and only if it compiles and executes all conforming programs according to the syntax and semantics laid out in Chapters 1, 2, 3 and 4. Appendices A, B, C, D, E and F and sections designated as Notes (see Section 1.7 on page 23) are for information purposes only and are not part of the specification.

The OpenMP API defines constructs that operate in the context of the base language that is supported by an implementation. If the base language does not support a language construct that appears in this document, a compliant OpenMP implementation is not required to support it, with the exception that for Fortran, the implementation must allow case insensitivity for directive and API routines names, and must allow identifiers of more than six characters

All library, intrinsic and built-in routines provided by the base language must be thread-safe in a compliant implementation. In addition, the implementation of the base language must also be thread-safe. For example, **ALLOCATE** and **DEALLOCATE** statements must be thread-safe in Fortran. Unsynchronized concurrent use of such routines by different threads must produce correct results (although not necessarily the same as serial execution results, as in the case of random number generation routines).

Starting with Fortran 90, variables with explicit initialization have the **SAVE** attribute implicitly. This is not the case in Fortran 77. However, a compliant OpenMP Fortran implementation must give such a variable the **SAVE** attribute, regardless of the underlying base language version.

Appendix C lists certain aspects of the OpenMP API that are implementation defined. A compliant implementation is required to define and document its behavior for each of the items in Appendix C.

1.6 Normative References

- ISO/IEC 9899:1990, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:1990 as C90.
- ISO/IEC 9899:1999, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:1999 as C99.
- ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:1998 as C++.
- ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003. The following features are not supported:
 - IEEE Arithmetic issues covered in Fortran 2003 Section 14
 - Parameterized derived types
 - The **PASS** attribute
 - Procedures bound to a type as operators
 - Overriding a type-bound procedure
 - Polymorphic entities

- 1 – **SELECT TYPE** construct
- 2 – Deferred bindings and abstract types
- 3 – Controlling IEEE underflow
- 4 – Another IEEE class value

5 Where this OpenMP API specification refers to C, C++ or Fortran, reference is made to the base
6 language supported by the implementation.

1.7 Organization of this Document

The remainder of this document is structured as follows:

- Chapter 2 “Directives”
- Chapter 3 “Runtime Library Routines”
- Chapter 4 “Environment Variables”
- Appendix A “Stubs for Runtime Library Routines”
- Appendix B “Interface Declarations”
- Appendix C “OpenMP Implementation-Defined Behaviors”
- Appendix D “Features History”

Some sections of this document only apply to programs written in a certain base language. Text that applies only to programs whose base language is C or C++ is shown as follows:

▼ C / C++ ▼

C/C++ specific text...

▲ C / C++ ▲

Text that applies only to programs whose base language is C only is shown as follows:

▼ C ▼

C specific text...

▲ C ▲

Text that applies only to programs whose base language is C90 only is shown as follows:

▼ C90 ▼

C90 specific text...

▲ C90 ▲

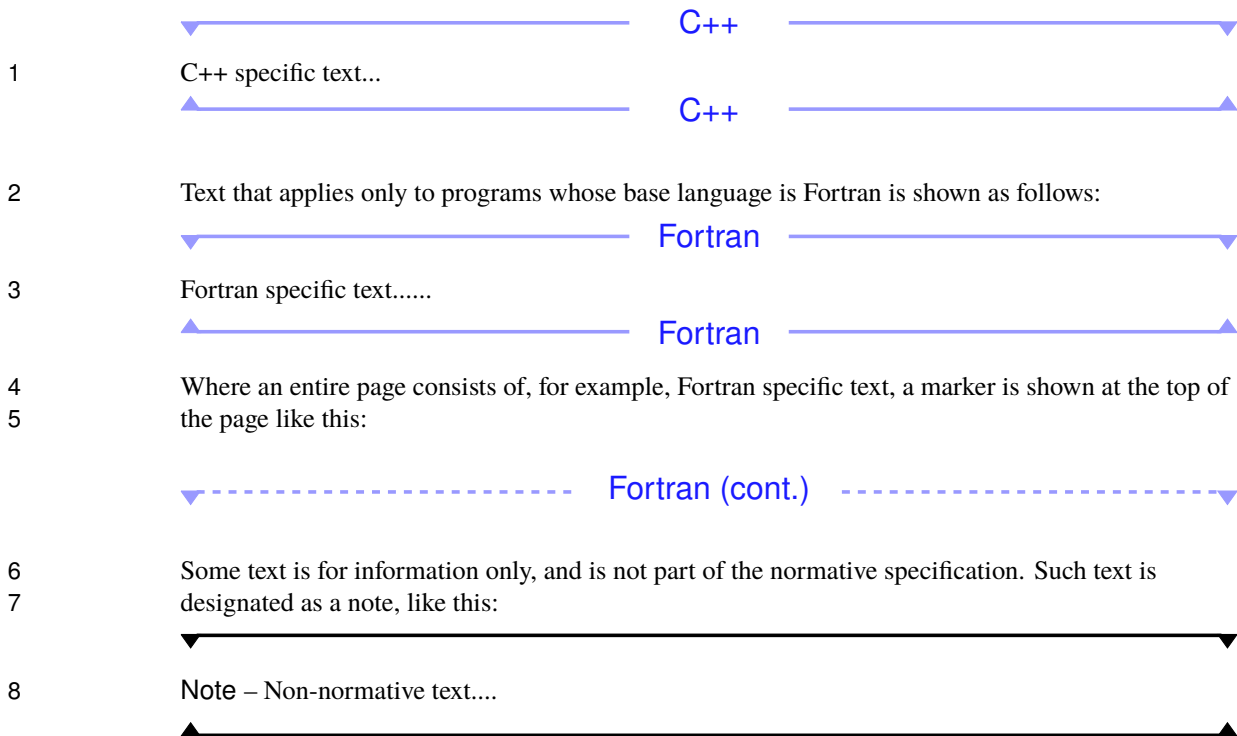
Text that applies only to programs whose base language is C99 only is shown as follows:

▼ C99 ▼

C99 specific text...

▲ C99 ▲

Text that applies only to programs whose base language is C++ only is shown as follows:



2 Directives

3 This chapter describes the syntax and behavior of OpenMP directives, and is divided into the
4 following sections:

- 5 • The language-specific directive format (Section 2.1 on page 26)
- 6 • Mechanisms to control conditional compilation (Section 2.2 on page 33)
- 7 • How to specify and to use array sections for all base languages (Section 2.4 on page 43)
- 8 • Control of OpenMP API ICVs (Section 2.3 on page 35)
- 9 • Details of each OpenMP directive (Section 2.5 on page 45 to Section 2.17 on page 224)

▼ C / C++ ▼

10 In C/C++, OpenMP directives are specified by using the **#pragma** mechanism provided by the C
11 and C++ standards.

▲ C / C++ ▲

▼ Fortran ▼

12 In Fortran, OpenMP directives are specified by using special comments that are identified by
13 unique sentinels. Also, a special comment form is available for conditional compilation.

▲ Fortran ▲

14 Compilers can therefore ignore OpenMP directives and conditionally compiled code if support of
15 the OpenMP API is not provided or enabled. A compliant implementation must provide an option
16 or interface that ensures that underlying support of all OpenMP directives and OpenMP conditional
17 compilation mechanisms is enabled. In the remainder of this document, the phrase *OpenMP*
18 *compilation* is used to mean a compilation with these OpenMP features enabled.

Restrictions

The following restriction applies to all OpenMP directives:

- OpenMP directives may not appear in **PURE** or **ELEMENTAL** procedures.

2.1 Directive Format

OpenMP directives for C/C++ are specified with the **pragma** preprocessing directive. The syntax of an OpenMP directive is as follows:

```
#pragma omp directive-name [clause[ [, ] clause] ... ] new-line
```

Each directive starts with **#pragma omp**. The remainder of the directive follows the conventions of the C and C++ standards for compiler directives. In particular, white space can be used before and after the **#**, and sometimes white space must be used to separate the words in a directive. Preprocessing tokens following the **#pragma omp** are subject to macro replacement.

Some OpenMP directives may be composed of consecutive **#pragma** preprocessing directives if specified in their syntax.

Directives are case-sensitive.

An OpenMP executable directive applies to at most one succeeding statement, which must be a structured block.

Fortran

OpenMP directives for Fortran are specified as follows:

sentinel directive-name [*clause*[*,* *clause*]...]

All OpenMP compiler directives must begin with a directive *sentinel*. The format of a *sentinel* differs between fixed and free-form source files, as described in Section 2.1.1 on page 28 and Section 2.1.2 on page 29.

Directives are case insensitive. Directives cannot be embedded within continued statements, and statements cannot be embedded within directives.

In order to simplify the presentation, free form is used for the syntax of OpenMP directives for Fortran in the remainder of this document, except as noted.

Fortran

Only one *directive-name* can be specified per directive (note that this includes combined directives, see Section 2.11 on page 120). The order in which clauses appear on directives is not significant. Clauses on directives may be repeated as needed, subject to the restrictions listed in the description of each clause.

Some data-sharing attribute clauses (Section 2.15.3 on page 185), data copying clauses (Section 2.15.4 on page 208), the **threadprivate** directive (Section 2.15.2 on page 179), the **flush** directive (Section 2.13.7 on page 158), and the **link** clause of the **declare target** directive (Section 2.10.4 on page 102) accept a *list*. The **to** clause of the **declare target** directive (Section 2.10.4 on page 102) accepts an *extended-list*. A *list* consists of a comma-separated collection of one or more *list items*. A *extended-list* consists of a comma-separated collection of one or more *extended list items*.

C / C++

A *list item* is a variable or array section. An *extended list item* is a *list item* or a function name.

C / C++

Fortran

A *list item* is a variable, array section or common block name (enclosed in slashes). An *extended list item* is a *list item* or a procedure name.

Fortran

For all base languages, a *list item* or an *extended list item* is subject to the restrictions specified in Section 2.4 on page 43 and in each of the sections describing clauses and directives for which the *list* or *extended-list* appears.

2.1.1 Fixed Source Form Directives

The following sentinels are recognized in fixed form source files:

!\$omp | c\$omp | *\$omp

Sentinels must start in column 1 and appear as a single word with no intervening characters.

Fortran fixed form line length, white space, continuation, and column rules apply to the directive line. Initial directive lines must have a space or zero in column 6, and continuation directive lines must have a character other than a space or a zero in column 6.

Comments may appear on the same line as a directive. The exclamation point initiates a comment when it appears after column 6. The comment extends to the end of the source line and is ignored. If the first non-blank character after the directive sentinel of an initial or continuation directive line is an exclamation point, the line is ignored.

Note – in the following example, the three formats for specifying the directive are equivalent (the first line represents the position of the first 9 columns):

c23456789

!\$omp parallel do shared(a,b,c)

c\$omp parallel do

c\$omp+shared(a,b,c)

c\$omp paralleldoshared(a,b,c)

2.1.2 Free Source Form Directives

The following sentinel is recognized in free form source files:

!\$omp

The sentinel can appear in any column as long as it is preceded only by white space (spaces and tab characters). It must appear as a single word with no intervening character. Fortran free form line length, white space, and continuation rules apply to the directive line. Initial directive lines must have a space after the sentinel. Continued directive lines must have an ampersand (&) as the last non-blank character on the line, prior to any comment placed inside the directive. Continuation directive lines can have an ampersand after the directive sentinel with optional white space before and after the ampersand.

Comments may appear on the same line as a directive. The exclamation point (!) initiates a comment. The comment extends to the end of the source line and is ignored. If the first non-blank character after the directive sentinel is an exclamation point, the line is ignored.

One or more blanks or horizontal tabs must be used to separate adjacent keywords in directives in free source form, except in the following cases, where white space is optional between the given set of keywords:

```

declare reduction
declare simd
declare target
distribute parallel do
distribute parallel do simd
distribute simd
do simd
end atomic
end critical
end distribute
end distribute parallel do
end distribute parallel do simd
    
```

```
1      end distribute simd
2      end do
3      end do simd
4      end master
5      end ordered
6      end parallel
7      end parallel do
8      end parallel do simd
9      end parallel sections
10     end parallel workshare
11     end sections
12     end simd
13     end single
14     end target
15     end target data
16     end target parallel
17     end target parallel do
18     end target parallel do simd
19     end target simd
20     end target teams
21     end target teams distribute
22     end target teams distribute parallel do
23     end target teams distribute parallel do simd
24     end target teams distribute simd
25     end task
26     end taskgroup
27     end taskloop
```



```

1      end teams
2      end teams distribute
3      end teams distribute parallel do
4      end teams distribute parallel do simd
5      end teams distribute simd
6      end workshare
7      parallel do
8      parallel do simd
9      parallel sections
10     parallel workshare
11     target data
12     target enter data
13     target exit data
14     target parallel
15     target parallel do
16     target parallel do simd
17     target simd
18     target teams
19     target teams distribute
20     target teams distribute parallel do
21     target teams distribute parallel do simd
22     target teams distribute simd
23     target update
24     teams distribute
25     teams distribute parallel do
26     teams distribute parallel do simd
27     teams distribute simd

```

Note – in the following example the three formats for specifying the directive are equivalent (the first line represents the position of the first 9 columns):

```

1      !23456789
2      !$omp parallel do
3          !$omp shared(a,b,c)
4
5      !$omp parallel
6      !$omp&do shared(a,b,c)
7
8      !$omp paralleldo shared(a,b,c)
9
10

```

Fortran

2.1.3 Stand-Alone Directives

Summary

Stand-alone directives are executable directives that have no associated user code.

Description

Stand-alone directives do not have any associated executable user code. Instead, they represent executable statements that typically do not have succinct equivalent statements in the base languages. There are some restrictions on the placement of a stand-alone directive within a program. A stand-alone directive may be placed only at a point where a base language executable statement is allowed.

Restrictions

C / C++

For C/C++, a stand-alone directive may not be used in place of the statement following an **if**, **while**, **do**, **switch**, or **label**.

C / C++

Fortran

For Fortran, a stand-alone directive may not be used as the action statement in an **if** statement or as the executable statement following a label if the label is referenced in the program.

Fortran

1 2.2 Conditional Compilation

2 In implementations that support a preprocessor, the `_OPENMP` macro name is defined to have the
3 decimal value `yyyymm` where `yyyy` and `mm` are the year and month designations of the version of
4 the OpenMP API that the implementation supports.

5 If this macro is the subject of a `#define` or a `#undef` preprocessing directive, the behavior is
6 unspecified.



7 The OpenMP API requires Fortran lines to be compiled conditionally, as described in the following
8 sections.

9 2.2.1 Fixed Source Form Conditional Compilation 10 Sentinels

11 The following conditional compilation sentinels are recognized in fixed form source files:

<code>!\$ *\$ c\$</code>

12 To enable conditional compilation, a line with a conditional compilation sentinel must satisfy the
13 following criteria:

- 14
- The sentinel must start in column 1 and appear as a single word with no intervening white space.
 - After the sentinel is replaced with two spaces, initial lines must have a space or zero in column 6 and only white space and numbers in columns 1 through 5.
 - After the sentinel is replaced with two spaces, continuation lines must have a character other than a space or zero in column 6 and only white space in columns 1 through 5.

17
18
19 If these criteria are met, the sentinel is replaced by two spaces. If these criteria are not met, the line
20 is left unchanged.

Note – in the following example, the two forms for specifying conditional compilation in fixed source form are equivalent (the first line represents the position of the first 9 columns):

```
c23456789
!$ 10 iam = omp_get_thread_num() +
!$   &          index

#ifdef _OPENMP
    10 iam = omp_get_thread_num() +
        &          index
#endif
```

2.2.2 Free Source Form Conditional Compilation Sentinel

The following conditional compilation sentinel is recognized in free form source files:

```
!$
```

To enable conditional compilation, a line with a conditional compilation sentinel must satisfy the following criteria:

- The sentinel can appear in any column but must be preceded only by white space.
- The sentinel must appear as a single word with no intervening white space.
- Initial lines must have a space after the sentinel.
- Continued lines must have an ampersand as the last non-blank character on the line, prior to any comment appearing on the conditionally compiled line. Continued lines can have an ampersand after the sentinel, with optional white space before and after the ampersand.

If these criteria are met, the sentinel is replaced by two spaces. If these criteria are not met, the line is left unchanged.

Note – in the following example, the two forms for specifying conditional compilation in free source form are equivalent (the first line represents the position of the first 9 columns):

```
c23456789
!$ iam = omp_get_thread_num() +      &
!$&    index

#ifdef _OPENMP
    iam = omp_get_thread_num() +      &
    index
#endif
```

Fortran

2.3 Internal Control Variables

An OpenMP implementation must act as if there are internal control variables (ICVs) that control the behavior of an OpenMP program. These ICVs store information such as the number of threads to use for future **parallel** regions, the schedule to use for worksharing loops and whether nested parallelism is enabled or not. The ICVs are given values at various times (described below) during the execution of the program. They are initialized by the implementation itself and may be given values through OpenMP environment variables and through calls to OpenMP API routines. The program can retrieve the values of these ICVs only through OpenMP API routines.

For purposes of exposition, this document refers to the ICVs by certain names, but an implementation is not required to use these names or to offer any way to access the variables other than through the ways shown in Section 2.3.2 on page 37.

1 2.3.1 ICV Descriptions

2 The following ICVs store values that affect the operation of **parallel** regions.

- 3 • *dyn-var* - controls whether dynamic adjustment of the number of threads is enabled for
4 encountered **parallel** regions. There is one copy of this ICV per data environment.
- 5 • *nest-var* - controls whether nested parallelism is enabled for encountered **parallel** regions.
6 There is one copy of this ICV per data environment.
- 7 • *nthreads-var* - controls the number of threads requested for encountered **parallel** regions.
8 There is one copy of this ICV per data environment.
- 9 • *thread-limit-var* - controls the maximum number of threads participating in the contention
10 group. There is one copy of this ICV per data environment.
- 11 • *max-active-levels-var* - controls the maximum number of nested active **parallel** regions.
12 There is one copy of this ICV per device.
- 13 • *place-partition-var* – controls the place partition available to the execution environment for
14 encountered **parallel** regions. There is one copy of this ICV per implicit task.
- 15 • *active-levels-var* - the number of nested, active parallel regions enclosing the current task such
16 that all of the **parallel** regions are enclosed by the outermost initial task region on the current
17 device. There is one copy of this ICV per data environment.
- 18 • *levels-var* - the number of nested parallel regions enclosing the current task such that all of the
19 **parallel** regions are enclosed by the outermost initial task region on the current device.
20 There is one copy of this ICV per data environment.
- 21 • *bind-var* - controls the binding of OpenMP threads to places. When binding is requested, the
22 variable indicates that the execution environment is advised not to move threads between places.
23 The variable can also provide default thread affinity policies. There is one copy of this ICV per
24 data environment.

25 The following ICVs store values that affect the operation of loop regions.

- 26 • *run-sched-var* - controls the schedule that the **runtime** schedule clause uses for loop regions.
27 There is one copy of this ICV per data environment.
- 28 • *def-sched-var* - controls the implementation defined default scheduling of loop regions. There is
29 one copy of this ICV per device.

30 The following ICVs store values that affect the program execution.

- 31 • *stacksize-var* - controls the stack size for threads that the OpenMP implementation creates. There
32 is one copy of this ICV per device.
- 33 • *wait-policy-var* - controls the desired behavior of waiting threads. There is one copy of this ICV
34 per device.

- *cancel-var* - controls the desired behavior of the **cancel** construct and cancellation points. There is one copy of this ICV for the whole program.
- *default-device-var* - controls the default target device. There is one copy of this ICV per data environment.
- *max-task-priority-var* - controls the maximum priority value that can be specified in the **priority** clause of the **task** construct. There is one copy of this ICV for the whole program.

2.3.2 ICV Initialization

Table 2.1 shows the ICVs, associated environment variables, and initial values.

ICV	Environment Variable	Initial value
<i>dyn-var</i>	OMP_DYNAMIC	See comments below
<i>nest-var</i>	OMP_NESTED	<i>false</i>
<i>nthreads-var</i>	OMP_NUM_THREADS	Implementation defined
<i>run-sched-var</i>	OMP_SCHEDULE	Implementation defined
<i>def-sched-var</i>	(none)	Implementation defined
<i>bind-var</i>	OMP_PROC_BIND	Implementation defined
<i>stacksize-var</i>	OMP_STACKSIZE	Implementation defined
<i>wait-policy-var</i>	OMP_WAIT_POLICY	Implementation defined
<i>thread-limit-var</i>	OMP_THREAD_LIMIT	Implementation defined
<i>max-active-levels-var</i>	OMP_MAX_ACTIVE_LEVELS	See comments below
<i>active-levels-var</i>	(none)	<i>zero</i>
<i>levels-var</i>	(none)	<i>zero</i>
<i>place-partition-var</i>	OMP_PLACES	Implementation defined
<i>cancel-var</i>	OMP_CANCELLATION	<i>false</i>
<i>default-device-var</i>	OMP_DEFAULT_DEVICE	Implementation defined

table continued on next page

table continued from previous page

ICV	Environment Variable	Initial value
<i>max-task-priority-var</i>	OMP_MAX_TASK_PRIORITY	zero

Table 2.1.: ICV Initial Values

Comments

- Each device has its own ICVs.
- The value of the *nthreads-var* ICV is a list.
- The value of the *bind-var* ICV is a list.
- The initial value of *dyn-var* is implementation defined if the implementation supports dynamic adjustment of the number of threads; otherwise, the initial value is *false*.
- The initial value of *max-active-levels-var* is the number of levels of parallelism that the implementation supports. See the definition of *supporting n levels of parallelism* in Section 1.2.7 on page 13 for further details.

The host and target device ICVs are initialized before any OpenMP API construct or OpenMP API routine executes. After the initial values are assigned, the values of any OpenMP environment variables that were set by the user are read and the associated ICVs for the host device are modified accordingly. The method for initializing a target device's ICVs is implementation defined.

Cross References

- **OMP_SCHEDULE** environment variable, see Section 4.1 on page 289.
- **OMP_NUM_THREADS** environment variable, see Section 4.2 on page 290.
- **OMP_DYNAMIC** environment variable, see Section 4.3 on page 291.
- **OMP_PROC_BIND** environment variable, see Section 4.4 on page 291.
- **OMP_PLACES** environment variable, see Section 4.5 on page 292.
- **OMP_NESTED** environment variable, see Section 4.6 on page 294.
- **OMP_STACKSIZE** environment variable, see Section 4.7 on page 295.
- **OMP_WAIT_POLICY** environment variable, see Section 4.8 on page 296.
- **OMP_MAX_ACTIVE_LEVELS** environment variable, see Section 4.9 on page 297.
- **OMP_THREAD_LIMIT** environment variable, see Section 4.10 on page 297.
- **OMP_CANCELLATION** environment variable, see Section 4.11 on page 297.

- **OMP_DEFAULT_DEVICE** environment variable, see Section 4.13 on page 299.
- **OMP_MAX_TASK_PRIORITY** environment variable, see Section 4.14 on page 300.

2.3.3 Modifying and Retrieving ICV Values

Table 2.2 shows the method for modifying and retrieving the values of ICVs through OpenMP API routines.

ICV	Ways to modify value	Ways to retrieve value
<i>dyn-var</i>	<code>omp_set_dynamic()</code>	<code>omp_get_dynamic()</code>
<i>nest-var</i>	<code>omp_set_nested()</code>	<code>omp_get_nested()</code>
<i>nthreads-var</i>	<code>omp_set_num_threads()</code>	<code>omp_get_max_threads()</code>
<i>run-sched-var</i>	<code>omp_set_schedule()</code>	<code>omp_get_schedule()</code>
<i>def-sched-var</i>	(none)	(none)
<i>bind-var</i>	(none)	<code>omp_get_proc_bind()</code>
<i>stacksize-var</i>	(none)	(none)
<i>wait-policy-var</i>	(none)	(none)
<i>thread-limit-var</i>	<code>thread_limit</code> clause	<code>omp_get_thread_limit()</code>
<i>max-active-levels-var</i>	<code>omp_set_max_active_levels()</code>	<code>omp_get_max_active_levels()</code>
<i>active-levels-var</i>	(none)	<code>omp_get_active_level()</code>
<i>levels-var</i>	(none)	<code>omp_get_level()</code>
<i>place-partition-var</i>	(none)	<code>omp_get_partition_num_places()</code>
<i>cancel-var</i>	(none)	<code>omp_get_cancellation()</code>
<i>default-device-var</i>	<code>omp_set_default_device()</code>	<code>omp_get_default_device()</code>
<i>max-task-priority-var</i>	(none)	<code>omp_get_max_task_priority()</code>

Table 2.2.: Ways to Modify and to Retrieve ICV Values

Comments

- The value of the *nthreads-var* ICV is a list. The runtime call `omp_set_num_threads()` sets the value of the first element of this list, and `omp_get_max_threads()` retrieves the value of the first element of this list.

- The value of the *bind-var* ICV is a list. The runtime call `omp_get_proc_bind()` retrieves the value of the first element of this list.
- Detailed values in the *place-partition-var* ICV are retrieved using the runtime calls `omp_get_partition_num_places()`, `omp_get_partition_place_nums()`, `omp_get_place_num_procs()`, and `omp_get_place_proc_ids()`.

Cross References

- `thread_limit` clause of the `teams` construct, see Section 2.10.5 on page 106.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 228.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 230.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 234.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 236.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 237.
- `omp_set_nested` routine, see Section 3.2.10 on page 237.
- `omp_get_nested` routine, see Section 3.2.11 on page 239.
- `omp_set_schedule` routine, see Section 3.2.12 on page 240.
- `omp_get_schedule` routine, see Section 3.2.13 on page 242.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 243.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 243.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 245.
- `omp_get_level` routine, see Section 3.2.17 on page 246.
- `omp_get_active_level` routine, see Section 3.2.20 on page 249.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 251.
- `omp_get_place_num_procs()` routine, see Section 3.2.24 on page 254.
- `omp_get_place_proc_ids()` routine, see Section 3.2.25 on page 255.
- `omp_get_partition_num_places()` routine, see Section 3.2.27 on page 257.
- `omp_get_partition_place_nums()` routine, see Section 3.2.28 on page 258.
- `omp_set_default_device` routine, see Section 3.2.29 on page 259.
- `omp_get_default_device` routine, see Section 3.2.30 on page 260.
- `omp_get_max_task_priority` routine, see Section 3.2.36 on page 265.

2.3.4 How ICVs are Scoped

Table 2.3 shows the ICVs and their scope.

ICV	Scope
<i>dyn-var</i>	data environment
<i>nest-var</i>	data environment
<i>nthreads-var</i>	data environment
<i>run-sched-var</i>	data environment
<i>def-sched-var</i>	device
<i>bind-var</i>	data environment
<i>stacksize-var</i>	device
<i>wait-policy-var</i>	device
<i>thread-limit-var</i>	data environment
<i>max-active-levels-var</i>	device
<i>active-levels-var</i>	data environment
<i>levels-var</i>	data environment
<i>place-partition-var</i>	implicit task
<i>cancel-var</i>	global
<i>default-device-var</i>	data environment
<i>max-task-priority-var</i>	global

Table 2.3.: Scopes of ICVs

Comments

- There is one copy per device of each ICV with device scope
 - Each data environment has its own copies of ICVs with data environment scope
 - Each implicit task has its own copy of ICVs with implicit task scope
- Calls to OpenMP API routines retrieve or modify data environment scoped ICVs in the data environment of their binding tasks.

1 **2.3.4.1 How the Per-Data Environment ICVs Work**

2 When a **task** construct or **parallel** construct is encountered, the generated task(s) inherit the
3 values of the data environment scoped ICVs from the generating task's ICV values.

4 When a **task** construct is encountered, the generated task inherits the value of *nthreads-var* from
5 the generating task's *nthreads-var* value. When a **parallel** construct is encountered, and the
6 generating task's *nthreads-var* list contains a single element, the generated task(s) inherit that list as
7 the value of *nthreads-var*. When a **parallel** construct is encountered, and the generating task's
8 *nthreads-var* list contains multiple elements, the generated task(s) inherit the value of *nthreads-var*
9 as the list obtained by deletion of the first element from the generating task's *nthreads-var* value.
10 The *bind-var* ICV is handled in the same way as the *nthreads-var* ICV.

11 When a **target** construct is encountered, the construct's data environment uses the values of the
12 data environment scoped ICVs from the device data environment ICV values of the device that will
13 execute the region. If a **teams** construct with a **thread_limit** clause is encountered, the
14 *thread-limit-var* ICV of the construct's data environment is instead set to a value that is less than or
15 equal to the value specified in the clause.

16 When encountering a loop worksharing region with **schedule(runtime)**, all implicit task
17 regions that constitute the binding parallel region must have the same value for *run-sched-var* in
18 their data environments. Otherwise, the behavior is unspecified.

19 **2.3.5 ICV Override Relationships**

20 Table 2.4 shows the override relationships among construct clauses and ICVs.

ICV	construct clause, if used
<i>dyn-var</i>	(none)
<i>nest-var</i>	(none)
<i>nthreads-var</i>	num_threads
<i>run-sched-var</i>	schedule
<i>def-sched-var</i>	schedule
<i>bind-var</i>	proc_bind
<i>stacksize-var</i>	(none)

table continued on next page

table continued from previous page

ICV	construct clause, if used
<i>wait-policy-var</i>	(none)
<i>thread-limit-var</i>	(none)
<i>max-active-levels-var</i>	(none)
<i>active-levels-var</i>	(none)
<i>levels-var</i>	(none)
<i>place-partition-var</i>	(none)
<i>cancel-var</i>	(none)
<i>default-device-var</i>	(none)
<i>max-task-priority-var</i>	(none)

Table 2.4.: ICV Override Relationships

Comments

- The **num_threads** clause overrides the value of the first element of the *nthreads-var* ICV.
- If *bind-var* is not set to *false* then the **proc_bind** clause overrides the value of the first element of the *bind-var* ICV; otherwise, the **proc_bind** clause has no effect.

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- **proc_bind** clause, Section 2.5 on page 45.
- **num_threads** clause, see Section 2.5.1 on page 49.
- Loop construct, see Section 2.7.1 on page 55.
- **schedule** clause, see Section 2.7.1.1 on page 62.

2.4 Array Sections

An array section designates a subset of the elements in an array. An array section can appear only in clauses where it is explicitly allowed.

To specify an array section in an OpenMP construct, array subscript expressions are extended with the following syntax:

[*lower-bound* : *length*] or

[*lower-bound* :] or

[: *length*] or

[:]

The array section must be a subset of the original array.

Array sections are allowed on multidimensional arrays. Base language array subscript expressions can be used to specify length-one dimensions of multidimensional array sections.

The *lower-bound* and *length* are integral type expressions. When evaluated they represent a set of integer values as follows:

{ *lower-bound*, *lower-bound* + 1, *lower-bound* + 2, ..., *lower-bound* + *length* - 1 }

The *lower-bound* and *length* must evaluate to non-negative integers.

When the size of the array dimension is not known, the *length* must be specified explicitly.

When the *length* is absent, it defaults to the size of the array dimension minus the *lower-bound*.

When the *lower-bound* is absent it defaults to 0.

Note – The following are examples of array sections:

a[0:6]

a[:6]

a[1:10]

a[1:]

b[10][:][:0]

c[1:10][42][0:6]

The first two examples are equivalent. If **a** is declared to be an eleven element array, the third and fourth examples are equivalent. The fifth example is a zero-length array section. The last example is not contiguous.

Fortran

Fortran has built-in support for array sections but the following restrictions apply for OpenMP constructs:

- A stride expression may not be specified.
- The upper bound for the last dimension of an assumed-size dummy array must be specified.

Fortran

Restrictions

Restrictions to array sections are as follows:

- An array section can appear only in clauses where it is explicitly allowed.

C / C++

- An array section can only be specified for a base language identifier.
- The type of the variable appearing in an array section must be array, pointer, reference to array, or reference to pointer.

C / C++

C++

- An array section cannot be used in a C++ user-defined `[]`-operator.

C++

2.5 parallel Construct

Summary

This fundamental construct starts parallel execution. See Section 1.3 on page 14 for a general description of the OpenMP execution model.

Syntax

C / C++

The syntax of the **parallel** construct is as follows:

```
#pragma omp parallel [clause[ [, ] clause] ... ] new-line
    structured-block
```

where *clause* is one of the following:

```
if ([parallel :] scalar-expression)
num_threads (integer-expression)
default (shared | none)
private (list)
firstprivate (list)
shared (list)
copyin (list)
reduction (reduction-identifier : list)
proc_bind (master | close | spread)
```

C / C++

Fortran

The syntax of the **parallel** construct is as follows:

```
!$omp parallel [clause[ [, ] clause] ... ]
    structured-block
!$omp end parallel
```


where *clause* is one of the following:

```
if ([parallel :] scalar-logical-expression)  
num_threads (scalar-integer-expression)  
default (private | firstprivate | shared | none)  
private (list)  
firstprivate (list)  
shared (list)  
copyin (list)  
reduction (reduction-identifier : list)  
proc_bind (master | close | spread)
```

The **end parallel** directive denotes the end of the **parallel** construct.

Fortran

Binding

The binding thread set for a **parallel** region is the encountering thread. The encountering thread becomes the master thread of the new team.

Description

When a thread encounters a **parallel** construct, a team of threads is created to execute the **parallel** region (see Section 2.5.1 on page 49 for more information about how the number of threads in the team is determined, including the evaluation of the **if** and **num_threads** clauses). The thread that encountered the **parallel** construct becomes the master thread of the new team, with a thread number of zero for the duration of the new **parallel** region. All threads in the new team, including the master thread, execute the region. Once the team is created, the number of threads in the team remains constant for the duration of that **parallel** region.

The optional **proc_bind** clause, described in Section 2.5.2 on page 51, specifies the mapping of OpenMP threads to places within the current place partition, that is, within the places listed in the *place-partition-var* ICV for the implicit task of the encountering thread.

Within a **parallel** region, thread numbers uniquely identify each thread. Thread numbers are consecutive whole numbers ranging from zero for the master thread up to one less than the number of threads in the team. A thread may obtain its own thread number by a call to the **omp_get_thread_num** library routine.

A set of implicit tasks, equal in number to the number of threads in the team, is generated by the encountering thread. The structured block of the **parallel** construct determines the code that will be

executed in each implicit task. Each task is assigned to a different thread in the team and becomes tied. The task region of the task being executed by the encountering thread is suspended and each thread in the team executes its implicit task. Each thread can execute a path of statements that is different from that of the other threads

The implementation may cause any thread to suspend execution of its implicit task at a task scheduling point, and switch to execute any explicit task generated by any of the threads in the team, before eventually resuming execution of the implicit task (for more details see Section 2.9 on page 81).

There is an implied barrier at the end of a **parallel** region. After the end of a **parallel** region, only the master thread of the team resumes execution of the enclosing task region.

If a thread in a team executing a **parallel** region encounters another **parallel** directive, it creates a new team, according to the rules in Section 2.5.1 on page 49, and it becomes the master of that new team.

If execution of a thread terminates while inside a **parallel** region, execution of all threads in all teams terminates. The order of termination of threads is unspecified. All work done by a team prior to any barrier that the team has passed in the program is guaranteed to be complete. The amount of work done by each thread after the last barrier that it passed and before it terminates is unspecified.

Restrictions

Restrictions to the **parallel** construct are as follows:

- A program that branches into or out of a **parallel** region is non-conforming.
- A program must not depend on any ordering of the evaluations of the clauses of the **parallel** directive, or on any side effects of the evaluations of the clauses.
- At most one **if** clause can appear on the directive.
- At most one **proc_bind** clause can appear on the directive.
- At most one **num_threads** clause can appear on the directive. The **num_threads** expression must evaluate to a positive integer value.

————— C / C++ —————

A **throw** executed inside a **parallel** region must cause execution to resume within the same **parallel** region, and the same thread that threw the exception must catch it.

————— C / C++ —————

————— Fortran —————

Unsynchronized use of Fortran I/O statements by multiple threads on the same unit has unspecified behavior.

————— Fortran —————

Cross References

- **if** Clause, see Section 2.12 on page 143.
- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section 2.15.3 on page 185.
- **copyin** clause, see Section 2.15.4 on page 208.
- **omp_get_thread_num** routine, see Section 3.2.4 on page 232.

2.5.1 Determining the Number of Threads for a **parallel** Region

When execution encounters a **parallel** directive, the value of the **if** clause or **num_threads** clause (if any) on the directive, the current parallel context, and the values of the *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs are used to determine the number of threads to use in the region.

Note that using a variable in an **if** or **num_threads** clause expression of a **parallel** construct causes an implicit reference to the variable in all enclosing constructs. The **if** clause expression and the **num_threads** clause expression are evaluated in the context outside of the **parallel** construct, and no ordering of those evaluations is specified. It is also unspecified whether, in what order, or how many times any side effects of the evaluation of the **num_threads** or **if** clause expressions occur.

When a thread encounters a **parallel** construct, the number of threads is determined according to Algorithm 2.1.

Algorithm 2.1

```
let ThreadsBusy be the number of OpenMP threads currently executing in this
contention group;

let ActiveParRegions be the number of enclosing active parallel regions;

if an if clause exists
then let IfClauseValue be the value of the if clause expression;
else let IfClauseValue = true;

if a num_threads clause exists
then let ThreadsRequested be the value of the num_threads clause expression;
```

```

1      else let ThreadsRequested = value of the first element of nthreads-var;
2      let ThreadsAvailable = (thread-limit-var - ThreadsBusy + 1);
3      if (IfClauseValue = false)
4      then number of threads = 1;
5      else if (ActiveParRegions >= 1) and (nest-var = false)
6      then number of threads = 1;
7      else if (ActiveParRegions = max-active-levels-var)
8      then number of threads = 1;
9      else if (dyn-var = true) and (ThreadsRequested <= ThreadsAvailable)
10     then number of threads = [ 1 : ThreadsRequested ];
11     else if (dyn-var = true) and (ThreadsRequested > ThreadsAvailable)
12     then number of threads = [ 1 : ThreadsAvailable ];
13     else if (dyn-var = false) and (ThreadsRequested <= ThreadsAvailable)
14     then number of threads = ThreadsRequested;
15     else if (dyn-var = false) and (ThreadsRequested > ThreadsAvailable)
16     then behavior is implementation defined;

```

Note – Since the initial value of the *dyn-var* ICV is implementation defined, programs that depend on a specific number of threads for correct execution should explicitly disable dynamic adjustment of the number of threads

Cross References

- *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs, see Section 2.3 on page 35.

2.5.2 Controlling OpenMP Thread Affinity

When a thread encounters a parallel directive without a **proc_bind** clause, the *bind-var* ICV is used to determine the policy for assigning OpenMP threads to places within the current place partition, that is, the places listed in the *place-partition-var* ICV for the implicit task of the encountering thread. If the parallel directive has a **proc_bind** clause then the binding policy specified by the **proc_bind** clause overrides the policy specified by the first element of the *bind-var* ICV. Once a thread in the team is assigned to a place, the OpenMP implementation should not move it to another place.

The **master** thread affinity policy instructs the execution environment to assign every thread in the team to the same place as the master thread. The place partition is not changed by this policy, and each implicit task inherits the *place-partition-var* ICV of the parent implicit task.

The **close** thread affinity policy instructs the execution environment to assign the threads in the team to places close to the place of the parent thread. The place partition is not changed by this policy, and each implicit task inherits the *place-partition-var* ICV of the parent implicit task. If T is the number of threads in the team, and P is the number of places in the parent's place partition, then the assignment of threads in the team to places is as follows:

- $T \leq P$. The master thread executes on the place of the parent thread. The thread with the next smallest thread number executes on the next place in the place partition, and so on, with wrap around with respect to the place partition of the master thread.
- $T > P$. Each place P will contain S_p threads with consecutive thread numbers, where $\lfloor (T/P) \rfloor \leq S_p \leq \lceil (T/P) \rceil$. The first S_0 threads (including the master thread) are assigned to the place of the parent thread. The next S_1 threads are assigned to the next place in the place partition, and so on, with wrap around with respect to the place partition of the master thread. When P does not divide T evenly, the exact number of threads in a particular place is implementation defined.

The purpose of the **spread** thread affinity policy is to create a sparse distribution for a team of T threads among the P places of the parent's place partition. A sparse distribution is achieved by first subdividing the parent partition into T subpartitions if $T \leq P$, or P subpartitions if $T > P$. Then one thread ($T \leq P$) or a set of threads ($T > P$) is assigned to each subpartition. The *place-partition-var* ICV of each implicit task is set to its subpartition. The subpartitioning is not only a mechanism for achieving a sparse distribution, it also defines a subset of places for a thread to use when creating a nested parallel region. The assignment of threads to places is as follows:

- $T \leq P$. The parent thread's place partition is split into T subpartitions, where each subpartition contains $\lfloor (P/T) \rfloor$ or $\lceil (P/T) \rceil$ consecutive places. A single thread is assigned to each subpartition. The master thread executes on the place of the parent thread and is assigned to the subpartition that includes that place. The thread with the next smallest thread number is assigned to the first place in the next subpartition, and so on, with wrap around with respect to the original place partition of the master thread.

- $T > P$. The parent thread's place partition is split into P subpartitions, each consisting of a single place. Each subpartition is assigned S_p threads with consecutive thread numbers, where $\lfloor (T/P) \rfloor \leq S_p \leq \lceil (T/P) \rceil$. The first S_0 threads (including the master thread) are assigned to the subpartition containing the place of the parent thread. The next S_1 threads are assigned to the next subpartition, and so on, with wrap around with respect to the original place partition of the master thread. When P does not divide T evenly, the exact number of threads in a particular subpartition is implementation defined.

The determination of whether the affinity request can be fulfilled is implementation defined. If the affinity request cannot be fulfilled, then the affinity of threads in the team is implementation defined.

Note - Wrap around is needed if the end of a place partition is reached before all thread assignments are done. For example, wrap around may be needed in the case of **close** and $T \leq P$, if the master thread is assigned to a place other than the first place in the place partition. In this case, thread 1 is assigned to the place after the place of the master place, thread 2 is assigned to the place after that, and so on. The end of the place partition may be reached before all threads are assigned. In this case, assignment of threads is resumed with the first place in the place partition.

2.6 Canonical Loop Form

C / C++

A loop has *canonical loop form* if it conforms to the following:

for (*init-expr*; *test-expr*; *incr-expr*) *structured-block*

init-expr One of the following:
 var = *lb*
 integer-type var = *lb*
 random-access-iterator-type var = *lb*
 pointer-type var = *lb*

test-expr One of the following:
 var relational-op b
 b relational-op var

continued on next page

continued from previous page

incr-expr One of the following:

```

++var
var++
-- var
var --
var += incr
var -= incr
var = var + incr
var = incr + var
var = var - incr

```

var One of the following:

A variable of a signed or unsigned integer type.
 For C++, a variable of a random access iterator type.
 For C, a variable of a pointer type.

If this variable would otherwise be shared, it is implicitly made private in the loop construct. This variable must not be modified during the execution of the *for-loop* other than in *incr-expr*. Unless the variable is specified **lastprivate** on the loop construct, its value after the loop is unspecified.

relational-op One of the following:

```

<
<=
>
>=

```

lb and *b* Loop invariant expressions of a type compatible with the type of *var*.

incr A loop invariant integer expression.

The canonical form allows the iteration count of all associated loops to be computed before executing the outermost loop. The computation is performed for each loop in an integer type. This type is derived from the type of *var* as follows:

- If *var* is of an integer type, then the type is the type of *var*.
- For C++, if *var* is of a random access iterator type, then the type is the type that would be used by *std::distance* applied to variables of the type of *var*.
- For C, if *var* is of a pointer type, then the type is **ptrdiff_t**.

The behavior is unspecified if any intermediate result required to compute the iteration count cannot be represented in the type determined above.

There is no implied synchronization during the evaluation of the *lb*, *b*, or *incr* expressions. It is unspecified whether, in what order, or how many times any side effects within the *lb*, *b*, or *incr* expressions occur.

Note – Random access iterators are required to support random access to elements in constant time. Other iterators are precluded by the restrictions since they can take linear time or offer limited functionality. It is therefore advisable to use tasks to parallelize those cases.

Restrictions

The following restrictions also apply:

- If *test-expr* is of the form *var relational-op b* and *relational-op* is *<* or *<=* then *incr-expr* must cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var relational-op b* and *relational-op* is *>* or *>=* then *incr-expr* must cause *var* to decrease on each iteration of the loop.
- If *test-expr* is of the form *b relational-op var* and *relational-op* is *<* or *<=* then *incr-expr* must cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *b relational-op var* and *relational-op* is *>* or *>=* then *incr-expr* must cause *var* to increase on each iteration of the loop.
- For C++, in the **simd** construct the only random access iterator types that are allowed for *var* are pointer types.

C / C++

1 2.7 Worksharing Constructs

2 A worksharing construct distributes the execution of the associated region among the members of
3 the team that encounters it. Threads execute portions of the region in the context of the implicit
4 tasks each one is executing. If the team consists of only one thread then the worksharing region is
5 not executed in parallel.

6 A worksharing region has no barrier on entry; however, an implied barrier exists at the end of the
7 worksharing region, unless a **nowait** clause is specified. If a **nowait** clause is present, an
8 implementation may omit the barrier at the end of the worksharing region. In this case, threads that
9 finish early may proceed straight to the instructions following the worksharing region without
10 waiting for the other members of the team to finish the worksharing region, and without performing
11 a flush operation.

12 The OpenMP API defines the following worksharing constructs, and these are described in the
13 sections that follow:

- 14 • **loop** construct
- 15 • **sections** construct
- 16 • **single** construct
- 17 • **workshare** construct

18 Restrictions

19 The following restrictions apply to worksharing constructs:

- 20 • Each worksharing region must be encountered by all threads in a team or by none at all, unless
21 cancellation has been requested for the innermost enclosing parallel region.
- 22 • The sequence of worksharing regions and **barrier** regions encountered must be the same for
23 every thread in a team

24 2.7.1 Loop Construct

25 Summary

26 The loop construct specifies that the iterations of one or more associated loops will be executed in
27 parallel by threads in the team in the context of their implicit tasks. The iterations are distributed
28 across threads that already exist in the team executing the **parallel** region to which the loop
29 region binds.

Syntax

C / C++

The syntax of the loop construct is as follows:

```
#pragma omp for [clause[ [, ] clause] ... ] new-line
    for-loops
```

where *clause* is one of the following:

```
private (list)
firstprivate (list)
lastprivate (list)
linear (list)
reduction (reduction-identifier : list)
schedule ([modifier : ]kind[ , chunk_size])
collapse (n)
ordered[ (n) ]
nowait
```

The **for** directive places restrictions on the structure of all associated *for-loops*. Specifically, all associated *for-loops* must have *canonical loop form* (see Section 2.6 on page 52).

C / C++

Fortran

The syntax of the loop construct is as follows:

```
!$omp do [clause[ [, ] clause] ... ]
    do-loops
[!$omp end do [nowait]]
```

where *clause* is one of the following:

```

1      private (list)
2      firstprivate (list)
3      lastprivate (list)
4      linear (list)
5      reduction (reduction-identifier : list)
6      schedule ([modifier : ]kind[ , chunk_size])
7      collapse (n)
8      ordered[ (n) ]

```

If an **end do** directive is not specified, an **end do** directive is assumed at the end of the *do-loops*.

Any associated *do-loop* must be a *do-construct* or an *inner-shared-do-construct* as defined by the Fortran standard. If an **end do** directive follows a *do-construct* in which several loop statements share a **DO** termination statement, then the directive can only be specified for the outermost of these **DO** statements.

If any of the loop iteration variables would otherwise be shared, they are implicitly made private on the loop construct.



Fortran

Binding

The binding thread set for a loop region is the current team. A loop region binds to the innermost enclosing **parallel** region. Only the threads of the team executing the binding **parallel** region participate in the execution of the loop iterations and the implied barrier of the loop region if the barrier is not eliminated by a **nowait** clause.

Description

The loop construct is associated with a loop nest consisting of one or more loops that follow the directive.

There is an implicit barrier at the end of a loop construct unless a **nowait** clause is specified.

The **collapse** clause or the **ordered** clause with the parameter may be used to specify how many loops are associated with the loop construct. An **ordered** clause without a parameter is equivalent to the **ordered** clause with a parameter value of one specified. The number of loops associated is determined by the parameters of the **collapse** clause and the **ordered** clause, which must be constant positive integer expressions.

If neither the **collapse** nor the **ordered** clause is present, the only loop that is associated with the loop construct is the one that immediately follows the loop directive. If the value of the

parameter in the **collapse** or **ordered** clause is larger than the number of nested loops following the construct, the behavior is unspecified.

If both the **collapse** clause and the **ordered** clause with a parameter are specified, the **collapse** clause applies to the loops that immediately follow the directive. The **ordered** clause with the parameter then applies to the resulting loop nest (that is, the collapsed loop and the remaining loops). If the value of the parameter in the **ordered** clause is larger than the number of the loops in the resulting loop nest, the behavior is unspecified.

If more than one loop is associated with the loop construct and a **collapse** clause is specified, then the iterations of all associated loops are collapsed into one larger iteration space that is then divided according to the schedule clause.

If an **ordered** clause with the parameter is specified for the loop construct, then the associated loops form a *doacross loop nest*.

The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space.

The iteration count for each associated loop is computed before entry to the outermost loop. If execution of any associated loop changes any of the values used to compute any of the iteration counts, then the behavior is unspecified.

The integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is implementation defined.

A worksharing loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations, and the logical numbering denotes the sequence in which the iterations would be executed if the associated loop(s) were executed by a single thread. The **schedule** clause specifies how iterations of the associated loops are divided into contiguous non-empty subsets, called chunks, and how these chunks are distributed among threads of the team. Each thread executes its assigned chunk(s) in the context of its implicit task. The *chunk_size* expression is evaluated using the original list items of any variables that are made private in the loop construct. It is unspecified whether, in what order, or how many times, any side effects of the evaluation of this expression occur. The use of a variable in a **schedule** clause expression of a loop construct causes an implicit reference to the variable in all enclosing constructs.

Different loop regions with the same schedule and iteration count, even if they occur in the same parallel region, can distribute iterations among threads differently. The only exception is for the **static** schedule as specified in Table 2.5. Programs that depend on which thread executes a particular iteration under any other circumstances are non-conforming.

See Section 2.7.1.1 on page 62 for details of how the schedule for a worksharing loop is determined.

The schedule *kind* can be one of those specified in Table 2.5.

The schedule *modifier* can be one of those specified in Table 2.6.

static

When **schedule(static, chunk_size)** is specified, iterations are divided into chunks of size *chunk_size*, and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.

When no *chunk_size* is specified, the iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each thread. Note that the size of the chunks is unspecified in this case.

A compliant implementation of the **static** schedule must ensure that the same assignment of logical iteration numbers to threads will be used in two loop regions if the following conditions are satisfied: 1) both loop regions have the same number of loop iterations, 2) both loop regions have the same value of *chunk_size* specified, or both loop regions have no *chunk_size* specified, 3) both loop regions bind to the same parallel region, and 4) neither loop is associated with a SIMD construct. A data dependence between the same logical iterations in two such loops is guaranteed to be satisfied allowing safe use of the **nowait** clause.

dynamic

When **schedule(dynamic, chunk_size)** is specified, the iterations are distributed to threads in the team in chunks as the threads request them. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be distributed.

Each chunk contains *chunk_size* iterations, except for the last chunk to be distributed, which may have fewer iterations.

When no *chunk_size* is specified, it defaults to 1

guided

When **schedule(guided, chunk_size)** is specified, the iterations are assigned to threads in the team in chunks as the executing threads request them. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be assigned.

For a *chunk_size* of 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads in the team, decreasing to 1. For a *chunk_size* with value *k* (greater than 1), the size of each chunk is determined in the same way, with the restriction that the chunks do not contain fewer than *k* iterations (except for the last chunk to be assigned, which may have fewer than *k* iterations).

table continued on next page

table continued from previous page

	When no <i>chunk_size</i> is specified, it defaults to 1.
auto	When schedule(auto) is specified, the decision regarding scheduling is delegated to the compiler and/or runtime system. The programmer gives the implementation the freedom to choose any possible mapping of iterations to threads in the team.
runtime	When schedule(runtime) is specified, the decision regarding scheduling is deferred until run time, and the schedule and chunk size are taken from the <i>run-sched-var</i> ICV. If the ICV is set to auto , the schedule is implementation defined.

Table 2.5.: **schedule** Clause *kind* Values

Note – For a team of p threads and a loop of n iterations, let $\lceil n/p \rceil$ be the integer q that satisfies $n = p * q - r$, with $0 \leq r < p$. One compliant implementation of the **static** schedule (with no specified *chunk_size*) would behave as though *chunk_size* had been specified with value q . Another compliant implementation would assign q iterations to the first $p - r$ threads, and $q - 1$ iterations to the remaining r threads. This illustrates why a conforming program must not rely on the details of a particular implementation.

A compliant implementation of the **guided** schedule with a *chunk_size* value of k would assign $q = \lceil n/p \rceil$ iterations to the first available thread and set n to the larger of $n - q$ and $p * k$. It would then repeat this process until q is greater than or equal to the number of remaining iterations, at which time the remaining iterations form the final chunk. Another compliant implementation could use the same method, except with $q = \lceil n/(2p) \rceil$, and set n to the larger of $n - q$ and $2 * p * k$.

simd	When the simd modifier is specified and the loop is associated with a SIMD construct, the <i>chunk_size</i> for all chunks except the first and last chunks is $new_chunk_size = \lceil chunk_size / simd_width \rceil * simd_width$ where <i>simd_width</i> is an implementation-defined value. The first chunk will have at least <i>new_chunk_size</i> iterations except if it is also the last chunk. The last chunk may have fewer iterations than <i>new_chunk_size</i> . If the simd modifier is specified and the loop is not associated with a SIMD construct, the modifier is ignored.
-------------	---

table continued on next page

2 Table 2.6.: **schedule** Clause *modifier* Values

3 **Restrictions**

4 Restrictions to the loop construct are as follows:

- 5 • All loops associated with the loop construct must be perfectly nested; that is, there must be no
6 intervening code nor any OpenMP directive between any two loops.
- 7 • The values of the loop control expressions of the loops associated with the loop construct must
8 be the same for all the threads in the team.
- 9 • Only one **schedule** clause can appear on a loop directive.
- 10 • Only one **collapse** clause can appear on a loop directive.
- 11 • *chunk_size* must be a loop invariant integer expression with a positive value.
- 12 • The value of the *chunk_size* expression must be the same for all threads in the team.
- 13 • The value of the *run-sched-var* ICV must be the same for all threads in the team.
- 14 • When **schedule(runtime)** or **schedule(auto)** is specified, *chunk_size* must not be
15 specified.
- 16 • Only one **ordered** clause can appear on a loop directive.
- 17 • The **ordered** clause must be present on the loop construct if any **ordered** region ever binds
18 to a loop region arising from the loop construct.
- 19 • The loop iteration variable may not appear in a **threadprivate** directive.

▼ C / C++ ▼

- 20 • The associated *for-loops* must be structured blocks.
- 21 • Only an iteration of the innermost associated loop may be curtailed by a **continue** statement.
- 22 • No statement can branch to any associated **for** statement.
- 23 • Only one **nowait** clause can appear on a **for** directive.
- 24 • A throw executed inside a loop region must cause execution to resume within the same iteration
25 of the loop region, and the same thread that threw the exception must catch it.

▲ C / C++ ▲

- The associated *do-loops* must be structured blocks.
- Only an iteration of the innermost associated loop may be curtailed by a **CYCLE** statement.
- No statement in the associated loops other than the **DO** statements can cause a branch out of the loops.
- The *do-loop* iteration variable must be of type integer.
- The *do-loop* cannot be a **DO WHILE** or a **DO** loop without loop control.

Cross References

- **private**, **firstprivate**, **lastprivate**, **linear**, and **reduction** clauses, see Section 2.15.3 on page 185.
- **OMP_SCHEDULE** environment variable, see Section 4.1 on page 289.
- **ordered** construct, see Section 2.13.8 on page 163.
- **depend** clause, see Section 2.13.9 on page 165.

2.7.1.1 Determining the Schedule of a Worksharing Loop

When execution encounters a loop directive, the **schedule** clause (if any) on the directive, and the *run-sched-var* and *def-sched-var* ICVs are used to determine how loop iterations are assigned to threads. See Section 2.3 on page 35 for details of how the values of the ICVs are determined. If the loop directive does not have a **schedule** clause then the current value of the *def-sched-var* ICV determines the schedule. If the loop directive has a **schedule** clause that specifies the **runtime** schedule kind then the current value of the *run-sched-var* ICV determines the schedule. Otherwise, the value of the **schedule** clause determines the schedule. Figure 2.1 describes how the schedule for a worksharing loop is determined.

Cross References

- ICVs, see Section 2.3 on page 35

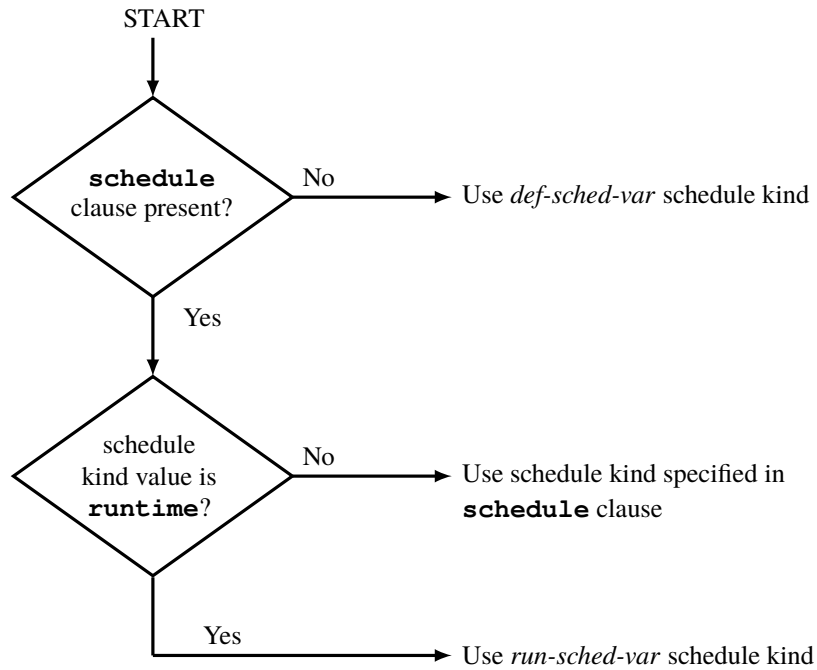


Figure 2.1.: Determining the **schedule** for a Worksharing Loop

1 2.7.2 sections Construct

2 Summary

3 The **sections** construct is a non-iterative worksharing construct that contains a set of structured
 4 blocks that are to be distributed among and executed by the threads in a team. Each structured
 5 block is executed once by one of the threads in the team in the context of its implicit task.

Syntax

C / C++

The syntax of the **sections** construct is as follows:

```
#pragma omp sections [clause[ [, ] clause] ... ] new-line
{
  [#pragma omp section new-line]
    structured-block
  [#pragma omp section new-line]
    structured-block
  ...
}
```

where *clause* is one of the following:

```
private(list)
firstprivate(list)
lastprivate(list)
reduction(reduction-identifier : list)
nowait
```

C / C++

The syntax of the **sections** construct is as follows:

```
!$omp sections [clause[ [, ] clause] ... ]
  [!$omp section]
    structured-block
  [!$omp section
    structured-block]
  ...
!$omp end sections [nowait]
```

where *clause* is one of the following:

```
private(list)
firstprivate(list)
lastprivate(list)
reduction(reduction-identifier : list)
```

Binding

The binding thread set for a **sections** region is the current team. A **sections** region binds to the innermost enclosing **parallel** region. Only the threads of the team executing the binding **parallel** region participate in the execution of the structured blocks and the implied barrier of the **sections** region if the barrier is not eliminated by a **nowait** clause.

Description

Each structured block in the **sections** construct is preceded by a **section** directive except possibly the first block, for which a preceding **section** directive is optional.

The method of scheduling the structured blocks among the threads in the team is implementation defined.

There is an implicit barrier at the end of a **sections** construct unless a **nowait** clause is specified.

Restrictions

Restrictions to the **sections** construct are as follows:

- Orphaned **section** directives are prohibited. That is, the **section** directives must appear within the **sections** construct and must not be encountered elsewhere in the **sections** region.
- The code enclosed in a **sections** construct must be a structured block.
- Only a single **nowait** clause can appear on a **sections** directive.
- A throw executed inside a **sections** region must cause execution to resume within the same section of the **sections** region, and the same thread that threw the exception must catch it.

C++

C++

Cross References

- **private**, **firstprivate**, **lastprivate**, and **reduction** clauses, see Section 2.15.3 on page 185.

2.7.3 single Construct

Summary

The **single** construct specifies that the associated structured block is executed by only one of the threads in the team (not necessarily the master thread), in the context of its implicit task. The other threads in the team, which do not execute the block, wait at an implicit barrier at the end of the **single** construct unless a **nowait** clause is specified.

Syntax

C / C++

The syntax of the single construct is as follows:

```
#pragma omp single [clause[ [, ] clause] ... ] new-line  
    structured-block
```

where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
copyprivate (list)  
nowait
```

C / C++

Fortran

The syntax of the **single** construct is as follows:

```
!$omp single [clause [ , ] clause ] ... ]  
    structured-block  
!$omp end single [end_clause [ , ] end_clause ] ... ]
```

where *clause* is one of the following:

```
private (list)  
firstprivate (list)
```

and *end_clause* is one of the following:

```
copyprivate (list)  
nowait
```

Fortran

Binding

The binding thread set for a **single** region is the current team. A **single** region binds to the innermost enclosing **parallel** region. Only the threads of the team executing the binding **parallel** region participate in the execution of the structured block and the implied barrier of the **single** region if the barrier is not eliminated by a **nowait** clause.

Description

The method of choosing a thread to execute the structured block is implementation defined. There is an implicit barrier at the end of the **single** construct unless a **nowait** clause is specified.

Restrictions

Restrictions to the **single** construct are as follows:

- The **copyprivate** clause must not be used with the **nowait** clause.
- At most one **nowait** clause can appear on a **single** construct.

C++

- A throw executed inside a **single** region must cause execution to resume within the same **single** region, and the same thread that threw the exception must catch it.

C++

Cross References

- **private** and **firstprivate** clauses, see Section 2.15.3 on page 185.
- **copyprivate** clause, see Section 2.15.4.2 on page 210.

Fortran

2.7.4 workshare Construct

Summary

The **workshare** construct divides the execution of the enclosed structured block into separate units of work, and causes the threads of the team to share the work such that each unit is executed only once by one thread, in the context of its implicit task.

Syntax

The syntax of the **workshare** construct is as follows:

```
!$omp workshare
    structured-block
!$omp end workshare [nowait]
```

The enclosed structured block must consist of only the following:

- array assignments
- scalar assignments
- **FORALL** statements

1 • **FORALL** constructs

2 • **WHERE** statements

3 • **WHERE** constructs

4 • **atomic** constructs

5 • **critical** constructs

6 • **parallel** constructs

7 Statements contained in any enclosed **critical** construct are also subject to these restrictions.

8 Statements in any enclosed **parallel** construct are not restricted.

9 Binding

10 The binding thread set for a **workshare** region is the current team. A **workshare** region binds
 11 to the innermost enclosing **parallel** region. Only the threads of the team executing the binding
 12 **parallel** region participate in the execution of the units of work and the implied barrier of the
 13 **workshare** region if the barrier is not eliminated by a **nowait** clause.

14 Description

15 There is an implicit barrier at the end of a **workshare** construct unless a **nowait** clause is
 16 specified.

17 An implementation of the **workshare** construct must insert any synchronization that is required
 18 to maintain standard Fortran semantics. For example, the effects of one statement within the
 19 structured block must appear to occur before the execution of succeeding statements, and the
 20 evaluation of the right hand side of an assignment must appear to complete prior to the effects of
 21 assigning to the left hand side.

22 The statements in the **workshare** construct are divided into units of work as follows:

- 23 • For array expressions within each statement, including transformational array intrinsic functions
 24 that compute scalar values from arrays:
 - 25 – Evaluation of each element of the array expression, including any references to **ELEMENTAL**
 26 functions, is a unit of work.
 - 27 – Evaluation of transformational array intrinsic functions may be freely subdivided into any
 28 number of units of work.
- 29 • For an array assignment statement, the assignment of each element is a unit of work.
- 30 • For a scalar assignment statement, the assignment operation is a unit of work.

- For a **WHERE** statement or construct, the evaluation of the mask expression and the masked assignments are each a unit of work.
- For a **FORALL** statement or construct, the evaluation of the mask expression, expressions occurring in the specification of the iteration space, and the masked assignments are each a unit of work
- For an **atomic** construct, the atomic operation on the storage location designated as x is a unit of work.
- For a **critical** construct, the construct is a single unit of work.
- For a **parallel** construct, the construct is a unit of work with respect to the **workshare** construct. The statements contained in the **parallel** construct are executed by a new thread team.
- If none of the rules above apply to a portion of a statement in the structured block, then that portion is a unit of work.

The transformational array intrinsic functions are **MATMUL**, **DOT_PRODUCT**, **SUM**, **PRODUCT**, **MAXVAL**, **MINVAL**, **COUNT**, **ANY**, **ALL**, **SPREAD**, **PACK**, **UNPACK**, **RESHAPE**, **TRANSPOSE**, **EOSHIFT**, **CSHIFT**, **MINLOC**, and **MAXLOC**.

It is unspecified how the units of work are assigned to the threads executing a **workshare** region.

If an array expression in the block references the value, association status, or allocation status of private variables, the value of the expression is undefined, unless the same value would be computed by every thread.

If an array assignment, a scalar assignment, a masked array assignment, or a **FORALL** assignment assigns to a private variable in the block, the result is unspecified.

The **workshare** directive causes the sharing of work to occur only in the **workshare** construct, and not in the remainder of the **workshare** region.

Restrictions

The following restrictions apply to the **workshare** construct:

- All array assignments, scalar assignments, and masked array assignments must be intrinsic assignments.
- The construct must not contain any user defined function calls unless the function is **ELEMENTAL**.

Fortran

1 2.8 SIMD Constructs

2 2.8.1 `simd` Construct

3 Summary

4 The **`simd`** construct can be applied to a loop to indicate that the loop can be transformed into a
5 SIMD loop (that is, multiple iterations of the loop can be executed concurrently using SIMD
6 instructions).

7 Syntax

8 The syntax of the **`simd`** construct is as follows:

▼ C / C++ ▼

```
#pragma omp simd [clause[ [, ] clause] ... ] new-line  
for-loops
```

9 where *clause* is one of the following:

```
10 safelen(length)  
11 simdlen(length)  
12 linear(list [ : linear-step])  
13 aligned(list [ : alignment])  
14 private(list)  
15 lastprivate(list)  
16 reduction(reduction-identifier : list)  
17 collapse(n)
```

18 The **`simd`** directive places restrictions on the structure of the associated *for-loops*. Specifically, all
19 associated *for-loops* must have *canonical loop form* (Section 2.6 on page 52).

▲ C / C++ ▲

```

1  !$omp simd [clause[ [, ] clause ... ]
    do-loops
2  [!$omp end simd]

```

where *clause* is one of the following:

```

3      safelen (length)
4      simdlen (length)
5      linear (list[ : linear-step])
6      aligned (list[ : alignment])
7      private (list)
8      lastprivate (list)
9      reduction (reduction-identifier : list)
10     collapse (n)

```

If an **end simd** directive is not specified, an **end simd** directive is assumed at the end of the *do-loops*.

All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end simd** directive follows a *do-construct* in which several loop statements share a **DO** termination statement, then the directive can only be specified for the outermost of these **DO** statements.

Binding

A **simd** region binds to the current task region. The binding thread set of the **simd** region is the current team.

1 **Description**

2 The **simd** construct enables the execution of multiple iterations of the associated loops
3 concurrently by means of SIMD instructions.

4 The **collapse** clause may be used to specify how many loops are associated with the construct.
5 The parameter of the **collapse** clause must be a constant positive integer expression. If no
6 **collapse** clause is present, the only loop that is associated with the loop construct is the one that
7 immediately follows the directive.

8 If more than one loop is associated with the **simd** construct, then the iterations of all associated
9 loops are collapsed into one larger iteration space that is then executed with SIMD instructions.
10 The sequential execution of the iterations in all associated loops determines the order of the
11 iterations in the collapsed iteration space.

12 The iteration count for each associated loop is computed before entry to the outermost loop. If
13 execution of any associated loop changes any of the values used to compute any of the iteration
14 counts, then the behavior is unspecified.

15 The integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is
16 implementation defined.

17 A SIMD loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations,
18 and the logical numbering denotes the sequence in which the iterations would be executed if the
19 associated loop(s) were executed with no SIMD instructions. If the **safelen** clause is used then
20 no two iterations executed concurrently with SIMD instructions can have a greater distance in the
21 logical iteration space than its value. The parameter of the **safelen** clause must be a constant
22 positive integer expression. If used, the **simdlen** clause specifies the preferred number of
23 iterations to be executed concurrently. The parameter of the **simdlen** clause must be a constant
24 positive integer. The number of iterations that are executed concurrently at any given time is
25 implementation defined. Each concurrent iteration will be executed by a different SIMD lane. Each
26 set of concurrent iterations is a SIMD chunk. Lexical forward dependencies in the iterations of the
27 original loop must be preserved within each SIMD chunk.

▼ C / C++ ▼

28 The **aligned** clause declares that the object to which each list item points is aligned to the
29 number of bytes expressed in the optional parameter of the **aligned** clause

▲ C / C++ ▲

▼ Fortran ▼

30 The **aligned** clause declares that the target of each list item is aligned to the number of bytes
31 expressed in the optional parameter of the **aligned** clause.

▲ Fortran ▲

32 The optional parameter of the **aligned** clause, *alignment*, must be a constant positive integer
33 expression. If no optional parameter is specified, implementation-defined default alignments for
34 SIMD instructions on the target platforms are assumed.

Restrictions

- All loops associated with the construct must be perfectly nested; that is, there must be no intervening code nor any OpenMP directive between any two loops.
- The associated loops must be structured blocks.
- A program that branches into or out of a **simd** region is non-conforming.
- Only one **collapse** clause can appear on a **simd** directive.
- A *list-item* cannot appear in more than one **aligned** clause.
- Only one **safelen** clause can appear on a **simd** directive.
- Only one **simdlen** clause can appear on a **simd** directive.
- If both **simdlen** and **safelen** clauses are specified, the value of the **simdlen** parameter must be less than or equal to the value of the **safelen** parameter.
- An ordered construct with the **simd** clause is the only OpenMP construct that can appear in the **simd** region.

C / C++

- The **simd** region cannot contain calls to the **longjmp** or **setjmp** functions.

C / C++

C

- The type of list items appearing in the **aligned** clause must be array or pointer.

C

C++

- The type of list items appearing in the **aligned** clause must be array, pointer, reference to array, or reference to pointer.

- No exception can be raised in the **simd** region.

C++

Fortran

- The *do-loop* iteration variable must be of type **integer**.
- The *do-loop* cannot be a **DO WHILE** or a **DO** loop without loop control.
- The type of list items appearing in the **aligned** clause must be **C_PTR** or Cray pointer, or the list item must have the **POINTER** or **ALLOCATABLE** attribute.

Fortran

Cross References

- **private**, **lastprivate**, **linear** and **reduction** clauses, see Section 2.15.3 on page 185.

2.8.2 declare simd Construct

Summary

The **declare simd** construct can be applied to a function (C, C++ and Fortran) or a subroutine (Fortran) to enable the creation of one or more versions that can process multiple arguments using SIMD instructions from a single invocation from a SIMD loop. The **declare simd** directive is a declarative directive. There may be multiple **declare simd** directives for a function (C, C++, Fortran) or subroutine (Fortran).

Syntax

The syntax of the **declare simd** construct is as follows:

C / C++

```
#pragma omp declare simd [clause[ [, ] clause] ... ] new-line
[#pragma omp declare simd [clause[ [, ] clause] ... ] new-line]
[ ... ]
    function definition or declaration
```

where *clause* is one of the following:

```
simdlen(length)
linear(linear-list [ : constant-linear-step])
aligned(argument-list [ : alignment])
uniform(argument-list)
inbranch
notinbranch
```

C / C++

Fortran

```
!$omp declare simd [ (proc-name) ] [clause [ , ] clause] ... ]
```

where *clause* is one of the following:

```
simdlen (length)  
linear (linear-list[ : constant-linear-step])  
aligned (argument-list[ : alignment])  
uniform (argument-list)  
inbranch  
notinbranch
```

Fortran

Description

C / C++

The use of a **declare simd** construct on a function enables the creation of SIMD versions of the associated function that can be used to process multiple arguments from a single invocation from a SIMD loop concurrently.

The expressions appearing in the clauses of this directive are evaluated in the scope of the arguments of the function declaration or definition

C / C++

Fortran

The use of a **declare simd** construct enables the creation of SIMD versions of the specified subroutine or function that can be used to process multiple arguments from a single invocation from a SIMD loop concurrently.

Fortran

If a **declare simd** directive contains multiple SIMD declarations, then one or more SIMD versions will be created for each declaration.

If a SIMD version is created, the number of concurrent arguments for the function is determined by the **simdlen** clause. If the **simdlen** clause is used its value corresponds to the number of concurrent arguments of the function. The parameter of the **simdlen** clause must be a constant positive integer expression. Otherwise, the number of concurrent arguments for the function is implementation defined.

C++

The special *this* pointer can be used as if was one of the arguments to the function in any of the **linear**, **aligned**, or **uniform** clauses.

C++

The **uniform** clause declares one or more arguments to have an invariant value for all concurrent invocations of the function in the execution of a single SIMD loop.

C / C++

The **aligned** clause declares that the object to which each list item points is aligned to the number of bytes expressed in the optional parameter of the **aligned** clause.

C / C++

Fortran

The **aligned** clause declares that the target of each list item is aligned to the number of bytes expressed in the optional parameter of the **aligned** clause.

Fortran

The optional parameter of the **aligned** clause, *alignment*, must be a constant positive integer expression. If no optional parameter is specified, implementation-defined default alignments for SIMD instructions on the target platforms are assumed.

The **inbranch** clause specifies that the function will always be called from inside a conditional statement of a SIMD loop. The **notinbranch** clause specifies that the function will never be called from inside a conditional statement of a SIMD loop. If neither clause is specified, then the function may or may not be called from inside a conditional statement of a SIMD loop.

Restrictions

- Each argument can appear in at most one **uniform** or **linear** clause.
- At most one **simdlen** clause can appear in a **declare simd** directive.
- Either **inbranch** or **notinbranch** may be specified, but not both.
- When a *constant-linear-step* expression is specified in a **linear** clause it must be a constant positive integer expression.
- The function or subroutine body must be a structured block.
- The execution of the function or subroutine, when called from a SIMD loop, cannot result in the execution of an OpenMP construct except for an **ordered** construct with the **simd** clause.
- The execution of the function or subroutine cannot have any side effects that would alter its execution for concurrent iterations of a SIMD chunk.
- A program that branches into or out of the function is non-conforming.

C / C++

- If the function has any declarations, then the **declare simd** construct for any declaration that has one must be equivalent to the one specified for the definition. Otherwise, the result is unspecified.
- The function cannot contain calls to the *longjmp* or *setjmp* functions.

C / C++

C

- The type of list items appearing in the **aligned** clause must be array or pointer.

C

C++

- The function cannot contain any calls to **throw**.
- The type of list items appearing in the **aligned** clause must be array, pointer, reference to array, or reference to pointer.

C++

- *proc-name* must not be a generic name, procedure pointer or entry name.
- If *proc-name* is omitted, the **declare simd** directive must appear in the specification part of a subroutine subprogram or a function subprogram for which creation of the SIMD versions is enabled.
- Any **declare simd** directive must appear in the specification part of a subroutine subprogram, function subprogram or interface body to which it applies.
- If a **declare simd** directive is specified in an interface block for a procedure, it must match a **declare simd** directive in the definition of the procedure.
- If a procedure is declared via a procedure declaration statement, the procedure *proc-name* should appear in the same specification.
- If a **declare simd** directive is specified for a procedure name with explicit interface and a **declare simd** directive is also specified for the definition of the procedure then the two **declare simd** directives must match. Otherwise the result is unspecified.
- Procedure pointers may not be used to access versions created by the **declare simd** directive.
- The type of list items appearing in the **aligned** clause must be **C_PTR** or Cray pointer, or the list item must have the **POINTER** or **ALLOCATABLE** attribute.

Cross References

- **reduction** clause, see Section 2.15.3.6 on page 198.
- **linear** clause, see Section 2.15.3.7 on page 204.

2.8.3 Loop SIMD Construct

Summary

The loop SIMD construct specifies that the iterations of one or more associated loops will be distributed across threads that already exist in the team and that the iterations executed by each thread can also be executed concurrently using SIMD instructions. The loop SIMD construct is a composite construct.

Syntax

C / C++

```
#pragma omp for simd [clause[ [, ] clause] ... ] new-line  
for-loops
```

where *clause* can be any of the clauses accepted by the **for** or **simd** directives with identical meanings and restrictions.

C / C++

Fortran

```
!$omp do simd [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end do simd [nowait]]
```

where *clause* can be any of the clauses accepted by the **simd** or **do** directives, with identical meanings and restrictions.

If an **end do simd** directive is not specified, an **end do simd** directive is assumed at the end of the *do-loops*.

Fortran

Description

The loop SIMD construct will first distribute the iterations of the associated loop(s) across the implicit tasks of the parallel region in a manner consistent with any clauses that apply to the loop construct. The resulting chunks of iterations will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

Restrictions

All restrictions to the loop construct and the **simd** construct apply to the loop SIMD construct. In addition, the following restriction applies:

- No **ordered** clause can be specified.
- A list item may appear in a **linear** or **firstprivate** clause but not both.

Cross References

- loop construct, see Section 2.7.1 on page 55.
- **simd** construct, see Section 2.8.1 on page 71.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.9 Tasking Constructs

2.9.1 task Construct

Summary

The **task** construct defines an explicit task.

Syntax

C / C++

The syntax of the **task** construct is as follows:

```
#pragma omp task [clause[ [, ] clause] ... ] new-line  
structured-block
```

where *clause* is one of the following:

```
if ([ task : ] scalar-expression)  
final (scalar-expression)  
untied  
default (shared | none)  
mergeable  
private (list)  
firstprivate (list)  
shared (list)  
depend (dependence-type : list)  
priority (priority-value)
```

C / C++
Fortran

The syntax of the **task** construct is as follows:

```
!$omp task [clause [ , ] clause] ... ]  
    structured-block  
!$omp end task
```

where *clause* is one of the following:

```
if ([ task : ] scalar-logical-expression)  
final (scalar-logical-expression)  
untied  
default (private | firstprivate | shared | none)  
mergeable  
private (list)  
firstprivate (list)  
shared (list)  
depend (dependence-type : list)  
priority (priority-value)
```

Fortran

Binding

The binding thread set of the **task** region is the current team. A **task** region binds to the innermost enclosing **parallel** region.

Description

When a thread encounters a **task** construct, a task is generated from the code for the associated structured block. The data environment of the task is created according to the data-sharing attribute clauses on the **task** construct, per-data environment ICVs, and any defaults that apply.

The encountering thread may immediately execute the task, or defer its execution. In the latter case, any thread in the team may be assigned the task. Completion of the task can be guaranteed using task synchronization constructs. A **task** construct may be nested inside an outer task, but the **task** region of the inner task is not a part of the **task** region of the outer task.

When an **if** clause is present on a **task** construct, and the **if** clause expression evaluates to *false*, an undeferred task is generated, and the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated task is completed. Note that the use of a variable in an **if** clause expression of a **task** construct causes an implicit reference to the variable in all enclosing constructs.

When a **final** clause is present on a **task** construct and the **final** clause expression evaluates to *true*, the generated task will be a final task. All **task** constructs encountered during execution of a final task will generate final and included tasks. Note that the use of a variable in a **final** clause expression of a **task** construct causes an implicit reference to the variable in all enclosing constructs.

The **if** clause expression and the **final** clause expression are evaluated in the context outside of the **task** construct, and no ordering of those evaluations is specified.

A thread that encounters a task scheduling point within the **task** region may temporarily suspend the **task** region. By default, a task is tied and its suspended task region can only be resumed by the thread that started its execution. If the **untied** clause is present on a **task** construct, any thread in the team can resume the **task** region after a suspension. The **untied** clause is ignored if a **final** clause is present on the same **task** construct and the **final** clause expression evaluates to *true*, or if a task is an included task.

The **task** construct includes a task scheduling point in the task region of its generating task, immediately following the generation of the explicit task. Each explicit **task** region includes a task scheduling point at its point of completion.

When a **mergeable** clause is present on a **task** construct, and the generated task is an undeferred task or an included task, the implementation may generate a merged task instead.

The **priority** clause is a hint for the priority of the generated task. The *priority-value* is a non-negative numerical scalar expression that provides a hint for task execution order. Among all tasks ready to be executed, higher priority tasks (those with a higher numerical value in the

priority clause expression) are recommended to execute before lower priority ones. The default *priority-value* when no **priority** clause is specified is zero (the lowest priority). If a value is specified in the **priority** clause that is higher than the *max-task-priority-var* ICV then the implementation will use the value of that ICV. A program that relies on task execution order being determined by this *priority-value* may have unspecified behavior.

Note – When storage is shared by an explicit **task** region, it is the programmer’s responsibility to ensure, by adding proper synchronization, that the storage does not reach the end of its lifetime before the explicit **task** region completes its execution.

Restrictions

Restrictions to the **task** construct are as follows:

- A program that branches into or out of a **task** region is non-conforming.
- A program must not depend on any ordering of the evaluations of the clauses of the **task** directive, or on any side effects of the evaluations of the clauses.
- At most one **if** clause can appear on the directive.
- At most one **final** clause can appear on the directive.
- At most one **priority** clause can appear on the directive.

C / C++

- A throw executed inside a **task** region must cause execution to resume within the same **task** region, and the same thread that threw the exception must catch it.

C / C++

Fortran

- Unsynchronized use of Fortran I/O statements by multiple tasks on the same unit has unspecified behavior

Fortran

Cross References

- Task scheduling constraints, see Section 2.9.5 on page 92.
- **depend** clause, see Section 2.13.9 on page 165.
- **if** Clause, see Section 2.12 on page 143.

1 2.9.2 taskloop Construct

2 Summary

3 The **taskloop** construct specifies that the iterations of one or more associated loops will be
4 executed in parallel using OpenMP tasks. The iterations are distributed across tasks created by the
5 construct and scheduled to be executed.

6 Syntax

C / C++

7 The syntax of the **taskloop** construct is as follows:

```
#pragma omp taskloop [clause[[,] clause] ...] new-line
    for-loops
```

8 where *clause* is one of the following:

```
9     if ([ taskloop :] scalar-expr)
10     shared (list)
11     private (list)
12     firstprivate (list)
13     lastprivate (list)
14     default (shared | none)
15     grainsize (grain-size)
16     num_tasks (num-tasks)
17     collapse (n)
18     final (scalar-expr)
19     priority (priority-value)
20     untied
21     mergeable
22     nogroup
```

23 The **taskloop** directive places restrictions on the structure of all associated *for-loops*.
24 Specifically, all associated *for-loops* must have canonical loop form (see Section 2.6 on page 52).

C / C++

The syntax of the **taskloop** construct is as follows:

```
!$omp taskloop [clause[[,] clause] ...]
      do-loops
[!$omp end taskloop]
```

where *clause* is one of the following:

```
if ([ taskloop :] scalar-logical-expr)
shared (list)
private (list)
firstprivate (list)
lastprivate (list)
default (private | firstprivate | shared | none)
grainsize (grain-size)
num_tasks (num-tasks)
collapse (n)
final (scalar-logical-expr)
priority (priority-value)
untied
mergeable
nogroup
```

If **end taskloop** directive is not specified, an **end taskloop** directive is assumed at the end of the *do-loops*.

All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end do** directive follows a *do-construct* in which several loop statements share a **DO** termination statement, then the directive can only be specified for the outermost of these **DO** statements.

If any of the loop iteration variables would otherwise be shared, they are implicitly made private for the loop-iteration tasks created by the **taskloop** construct. Unless the loop iteration variables are specified in a **lastprivate** clause on the **taskloop** construct, their values after the loop are unspecified.

Binding

The binding thread set of the **taskloop** construct is the current team. A **taskloop** region binds to the innermost enclosing **parallel** region.

Description

When a thread encounters a **taskloop** construct, the construct partitions the associated loops into tasks for parallel execution of the loops' iterations. The data environment of the created tasks is created according to the data-sharing attribute clauses on the **taskloop** construct, per-data environment ICVs, and any defaults that apply. The order of the creation of the loop tasks is unspecified.

If a **grainsize** clause is present on the **taskloop** construct, the number of logical loop iterations assigned to each created task is larger than or equal to the value of the *grain-size* expression, but less than two times the value of the *grain-size* expression and less than or equal to the number of logical loop iterations. The parameter of the **grainsize** clause must be a positive integer expression. If **num_tasks** is specified, the **taskloop** construct creates as many tasks as the minimum of the *num-tasks* expression and the number of logical loop iterations. The parameter of the **num_tasks** clause must evaluate to a positive integer. If neither a **grainsize** nor **num_tasks** clause is present, the number of loop tasks created and the number of logical loop iterations assigned to these tasks is implementation defined.

The **collapse** clause may be used to specify how many loops are associated with the **taskloop** construct. The parameter of the **collapse** clause must be a constant positive integer expression. If no **collapse** clause is present, the only loop that is associated with the **taskloop** construct is the one that immediately follows the **taskloop** directive.

If more than one loop is associated with the **taskloop** construct, then the iterations of all associated loops are collapsed into one larger iteration space that is then divided according to the **grainsize** and **num_tasks** clauses. The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space.

The iteration count for each associated loop is computed before entry to the outermost loop. If execution of any associated loop changes any of the values used to compute any of the iteration counts, then the behavior is unspecified. The integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is implementation defined.

When an **if** clause is present on a **taskloop** construct, and if the **if** clause expression evaluates to *false*, undeferred tasks are generated. The use of a variable in an **if** clause expression of a **taskloop** construct causes an implicit reference to the variable in all enclosing constructs.

When a **final** clause is present on a **taskloop** construct and the **final** clause expression evaluates to *true*, the generated tasks will be final tasks. The use of a variable in a **final** clause expression of a **taskloop** construct causes an implicit reference to the variable in all enclosing constructs.

When a **priority** clause is present on a **taskloop** construct, the generated tasks will have the *priority-value* as if it was specified for each individual task. If the **priority** clause has not been specified, tasks generated by the **taskloop** construct have the default task priority (zero).

If the **untied** clause is specified, all tasks created by the **taskloop** construct are untied tasks.

When a **mergeable** clause is present on a **taskloop** construct, and the generated tasks are undeferred or included tasks, the implementation may generate merged tasks instead.

By default, the **taskloop** construct executes as if it was enclosed in a **taskgroup** construct with no statements or directives outside of the **taskloop** construct. Thus, the **taskloop** construct creates an implicit **taskgroup** region. If the **nogroup** clause is present, no implicit **taskgroup** region is created.

Note – When storage is shared by a **taskloop** region, the programmer must ensure, by adding proper synchronization, that the storage does not reach the end of its lifetime before the **taskloop** region and its descendant tasks complete their execution.

Restrictions

The restrictions of the **taskloop** construct are as follows:

- A program that branches into or out of a **taskloop** region is non-conforming.
- All loops associated with the **taskloop** construct must be perfectly nested; that is, there must be no intervening code nor any OpenMP directive between any two loops.
- At most one **grainsize** clause can appear on a **taskloop** directive.
- At most one **num_tasks** clause can appear on a **taskloop** directive.
- The **grainsize** clause and **num_tasks** clause are mutually exclusive and may not appear on the same **taskloop** directive.
- At most one **collapse** clause can appear on a **taskloop** directive.
- At most one **if** clause can appear on the directive.
- At most one **final** clause can appear on the directive.
- At most one **priority** clause can appear on the directive.

Cross References

- **task** construct, Section 2.9.1 on page 81.
- **taskgroup** construct, Section 2.13.5 on page 150.
- Data-sharing attribute clauses, Section 2.15.3 on page 185.
- **if** Clause, see Section 2.12 on page 143.

2.9.3 taskloop simd Construct

Summary

The **taskloop simd** construct specifies a loop that can be executed concurrently using SIMD instructions and that those iterations will also be executed in parallel using OpenMP tasks. The **taskloop simd** construct is a composite construct.

Syntax

C / C++

The syntax of the **taskloop simd** construct is as follows:

```
#pragma omp taskloop simd [clause[[,] clause] ...] new-line  
for-loops
```

where *clause* can be any of the clauses accepted by the **taskloop** or **simd** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **taskloop simd** construct is as follows:

```
!$omp taskloop simd [clause[[,] clause] ...]  
do-loops  
[!$omp end taskloop simd]
```

where *clause* can be any of the clauses accepted by the **taskloop** or **simd** directives with identical meanings and restrictions.

If **end taskloop simd** directive is not specified, an **end taskloop simd** directive is assumed at the end of the *do-loops*.

Fortran

Binding

The binding thread set of the **taskloop simd** construct is the current team. A **taskloop simd** region binds to the innermost enclosing parallel region.

Description

The **taskloop simd** construct will first distribute the iterations of the associated loop(s) across tasks in a manner consistent with any clauses that apply to the **taskloop** construct. The resulting tasks will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

Restrictions

- The restrictions for the **taskloop** and **simd** constructs apply.
- No **reduction** clause can be specified.

Cross References

- **taskloop** construct, see Section 2.9.2 on page 85.
- **simd** construct, see Section 2.8.1 on page 71.
- Data-sharing attribute clauses, see Section 2.15.3 on page 185.

2.9.4 taskyield Construct

Summary

The **taskyield** construct specifies that the current task can be suspended in favor of execution of a different task. The **taskyield** construct is a stand-alone directive.

Syntax

▼ C / C++ ▼

The syntax of the **taskyield** construct is as follows:

```
#pragma omp taskyield new-line
```

▲ C / C++ ▲

The syntax of the **taskyield** construct is as follows:

```
!$omp taskyield
```

Binding

A **taskyield** region binds to the current task region. The binding thread set of the **taskyield** region is the current team.

Description

The **taskyield** region includes an explicit task scheduling point in the current task region.

Cross References

- Task scheduling, see Section [2.9.5](#) on page [92](#).

1 2.9.5 Task Scheduling

2 Whenever a thread reaches a task scheduling point, the implementation may cause it to perform a
3 task switch, beginning or resuming execution of a different task bound to the current team. Task
4 scheduling points are implied at the following locations:

- 5 • the point immediately following the generation of an explicit task
- 6 • after the point of completion of a **task** region
- 7 • in a **taskyield** region
- 8 • in a **taskwait** region
- 9 • at the end of a **taskgroup** region
- 10 • in an implicit and explicit **barrier** region
- 11 • the point immediately following the generation of a **target** region
- 12 • at the beginning and end of a **target data** region
- 13 • in a **target update** region
- 14 • in either the **omp_target_memcpy** or the **omp_target_memcpy_rect** routine

15 When a thread encounters a task scheduling point it may do one of the following, subject to the
16 *Task Scheduling Constraints* (below):

- 17 • begin execution of a tied task bound to the current team
- 18 • resume any suspended task region, bound to the current team, to which it is tied
- 19 • begin execution of an untied task bound to the current team
- 20 • resume any suspended untied task region bound to the current team.

21 If more than one of the above choices is available, it is unspecified as to which will be chosen.

22 *Task Scheduling Constraints* are as follows:

- 23 1. An included task is executed immediately after generation of the task.
- 24 2. Scheduling of new tied tasks is constrained by the set of task regions that are currently tied to the
25 thread, and that are not suspended in a **barrier** region. If this set is empty, any new tied task
26 may be scheduled. Otherwise, a new tied task may be scheduled only if it is a descendent task of
27 every task in the set.
- 28 3. A dependent task shall not be scheduled until its task dependences are fulfilled.
- 29 4. When an explicit task is generated by a construct containing an **if** clause for which the
30 expression evaluated to *false*, and the previous constraints are already met, the task is executed
31 immediately after generation of the task.

A program relying on any other assumption about task scheduling is non-conforming.

Note – Task scheduling points dynamically divide task regions into parts. Each part is executed uninterrupted from start to end. Different parts of the same task region are executed in the order in which they are encountered. In the absence of task synchronization constructs, the order in which a thread executes parts of different schedulable tasks is unspecified.

A correct program must behave correctly and consistently with all conceivable scheduling sequences that are compatible with the rules above.

For example, if **threadprivate** storage is accessed (explicitly in the source code or implicitly in calls to library routines) in one part of a task region, its value cannot be assumed to be preserved into the next part of the same task region if another schedulable task exists that modifies it.

As another example, if a lock acquire and release happen in different parts of a task region, no attempt should be made to acquire the same lock in any part of another task that the executing thread may schedule. Otherwise, a deadlock is possible. A similar situation can occur when a critical region spans multiple parts of a task and another schedulable task contains a critical region with the same name.

The use of threadprivate variables and the use of locks or critical sections in an explicit task with an **if** clause must take into account that when the **if** clause evaluates to *false*, the task is executed immediately, without regard to *Task Scheduling Constraint 2*.

2.10 Device Constructs

2.10.1 target data Construct

Summary

Map variables to a device data environment for the extent of the region.

Syntax

C / C++

The syntax of the **target data** construct is as follows:

```
#pragma omp target data [clause[ [, ] clause] ... ] new-line
structured-block
```

where *clause* is one of the following:

```
if([ target data :] scalar-expression)
device(integer-expression)
map([[map-type-modifier[,]] map-type: ] list)
```

C / C++

Fortran

The syntax of the **target data** construct is as follows:

```
!$omp target data [clause[ [, ] clause] ... ]
structured-block
!$omp end target data
```

where *clause* is one of the following:

```
if([ target data :] scalar-logical-expression)
device(scalar-integer-expression)
map([[map-type-modifier[,]] map-type: ] list)
```

The **end target data** directive denotes the end of the **target data** construct.

Fortran

Binding

The binding task region for a **target data** construct is the encountering task. The target region binds to the enclosing parallel or task region.

Description

When a **target data** construct is encountered, the encountering task executes the region. If there is no **device** clause, the default device is determined by the *default-device-var* ICV. Variables are mapped for the extent of the region, according to any data-mapping clauses, from the data environment of the encountering task to the device data environment. When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

List items that appear in a **use_device_ptr** clause are converted into device pointers to the corresponding list item in the device data environment.

Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target data** directive, or on any side effects of the evaluations of the clauses.
- At most one **device** clause can appear on the directive. The **device** expression must evaluate to a non-negative integer value strictly less than the value of `omp_get_num_devices()`.
- At most one **if** clause can appear on the directive.
- A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- At least one **map** clause must appear on the directive.
- A list item in a **use_device_ptr** clause must have a corresponding list item in the device data environment.
- References in the construct to a list item that appears in a **use_device_ptr** clause must be to the address of the list item.

Cross References

- **map** clause, see Section [2.15.5.1](#) on page [213](#).
- *default-device-var*, see Section [2.3](#) on page [35](#).
- **if** Clause, see Section [2.12](#) on page [143](#).

2.10.2 target Construct

Summary

Map variables to a device data environment and execute the construct on that device.

Syntax

C / C++

The syntax of the **target** construct is as follows:

```
#pragma omp target [clause[ [, ] clause] ... ] new-line  
structured-block
```

where *clause* is one of the following:

```
if([ target :] scalar-expression)  
device(integer-expression)  
private(list)  
firstprivate(list)  
map([[map-type-modifier[,]] map-type: ] list)  
is_device_ptr(list)  
defaultmap(tofrom:scalar)  
nowait  
depend(dependence-type: list)
```

C / C++

Fortran

The syntax of the **target** construct is as follows:

```
!$omp target [clause[ [, ] clause] ... ]  
structured-block  
!$omp end target
```

where *clause* is one of the following:

```
if([ target :] scalar-logical-expression)
device (scalar-integer-expression)
private (list)
firstprivate (list)
map ([[map-type-modifier[,]] map-type: ] list)
is_device_ptr (list)
defaultmap (tofrom: scalar)
nowait
depend (dependence-type : list)
```

The **end target** directive denotes the end of the **target** construct

Fortran

Binding

The binding task for a **target** construct is the encountering task. The target region binds to the enclosing parallel or task region.

Description

The **target** construct provides a superset of the functionality and restrictions provided by the **target data** directive, excepting the **use_device_ptr** clause.

The functionality added to the **target** directive is the inclusion of an executable region to be executed by a device. That is, the **target** directive is an executable directive.

The **target** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated *target task* is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target** construct.

When an **if** clause is present and the **if** clause expression evaluates to *false*, the **target** region is executed by the host device in the host data environment.

The **is_device_ptr** clause is used to indicate that a list item is a device pointer already in the device data environment and that it should be used directly. Support for pointers created outside of OpenMP, specifically outside of the **omp_target_alloc** routine and the **use_device_ptr** clause, is implementation defined.

▼ C / C++ ▼

If an array section is a list item in a **map** clause, and the array section is derived from a variable for which the type is pointer, reference to array or reference to pointer then the data-sharing attribute for that variable in the construct is private. Prior to the execution of the construct, the private variable is initialized with the address of the storage location of the corresponding array section in the device data environment.

▲ C / C++ ▲

Restrictions

- If a **target**, **target update**, **target data**, **target enter data**, or **target exit data** construct appears within a **target** region the behavior is unspecified.
- The result of an **omp_set_default_device**, **omp_get_default_device**, or **omp_get_num_devices** routine called within a **target** region is unspecified.
- The effect of an access to a **threadprivate** variable in a target region is unspecified.
- If a list item in a **map** clause is a structure element, any other element of that structure that is referenced in the **target** construct must also appear as a list item in a **map** clause.
- A variable referenced in a **target** region but not the **target** construct that is not declared in the **target** region must appear in a **declare target** directive.
- The restrictions for the **task** construct apply.
- At most one **defaultmap** clause can appear on the directive.
- A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- A list item that appears in an **is_device_ptr** must be a valid device pointer in the device data environment.

▼ C / C++ ▼

- If a list item in a **map** clause is an array section, and the array section is derived from a variable for which the type is pointer, reference to array or reference to pointer then the behavior is unspecified if that variable's corresponding private variable is modified in the region.

▲ C / C++ ▲

C++

- A list item that appears in an **is_device_ptr** clause must have a type of pointer, array, reference to pointer or reference to array.
- A throw executed inside a **target** region must cause execution to resume within the same **target** region, and the same thread that threw the exception must catch it.

C++

Fortran

- A list item that appears in an **is_device_ptr** clause must be a dummy argument.
- If a list item in a **map** clause is an array section, and the array section is derived from a variable with a **POINTER** or **ALLOCATABLE** attribute then the behavior is unspecified if the corresponding list item's variable is modified in the region.

Fortran

Cross References

- *default-device-var*, see Section 2.3 on page 35.
- **target data** construct, see Section 2.10.1 on page 93.
- **task** construct, see Section 2.9.1 on page 81.
- **task** scheduling constraints, see Section 2.9.5 on page 92
- **if** Clause, see Section 2.12 on page 143.
- Data-mapping Attribute Rules and Clauses, see Section 2.15.5 on page 212.
- **private** and **firstprivate** clauses, see Section 2.15.3 on page 185.

2.10.3 target update Construct

Summary

The **target update** directive makes the corresponding list items in the device data environment consistent with their original list items, according to the specified motion clauses. The **target update** construct is a stand-alone directive.

Syntax

C / C++

The syntax of the **target update** construct is as follows:

```
#pragma omp target update clause[ [, ] clause] ... ] new-line
```

where *clause* is either *motion-clause* or one of the following:

```
if([ target update :] scalar-expression)
```

```
device(integer-expression)
```

```
nowait
```

```
depend (dependence-type : list)
```

and *motion-clause* is one of the following:

```
to(list)
```

```
from(list)
```

C / C++

Fortran

The syntax of the **target update** construct is as follows:

```
!$omp target update clause[ [, ] clause] ... ]
```

where *clause* is either *motion-clause* or one of the following:

```
if([target update :] scalar-logical-expression)
```

```
device(scalar-integer-expression)
```

```
nowait
```

```
depend (dependence-type : list)
```

and *motion-clause* is one of the following:

```
to(list)
```

```
from(list)
```

Fortran

Binding

The binding task for a **target update** construct is the encountering task. The **target update** directive is a stand-alone directive.

Description

For each list item in a **to** or **from** clause there is a corresponding list item and an original list item. If the corresponding list item is not present in the device data environment then no assignment occurs to or from the original list item. Otherwise, each corresponding list item in the device data environment has an original list item in the current task's data environment.

For each list item in a **from** clause the value of the corresponding list item is assigned to the original list item.

For each list item in a **to** clause the value of the original list item is assigned to the corresponding list item.

The list items that appear in the **to** or **from** clauses may include array sections.

The **target update** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target update** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target update** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated *target task* is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target update** construct.

The device is specified in the **device** clause. If there is no **device** clause, the device is determined by the *default-device-var* ICV. When an **if** clause is present and the **if** clause expression evaluates to *false* then no assignments occur.

Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target update** directive, or on any side effects of the evaluations of the clauses.
- At least one *motion-clause* must be specified.
- If a list item is an array section it must specify contiguous storage.

- A list item can only appear in a **to** or **from** clause, but not both.
- A list item in a **to** or **from** clause must have a mappable type
- At most one **device** clause can appear on the directive. The **device** expression must evaluate to a non-negative integer value.
- At most one **if** clause can appear on the directive.
- The restrictions for the **task** construct apply.

Cross References

- *default-device-var*, see Section 2.3 on page 35.
- **target data**, see Section 2.10.1 on page 93.
- Array sections, Section 2.4 on page 43
- **task** construct, see Section 2.9.1 on page 81.
- **task** scheduling constraints, see Section 2.9.5 on page 92
- **if** Clause, see Section 2.12 on page 143.

2.10.4 declare target Directive

Summary

The **declare target** directive specifies that variables, functions (C, C++ and Fortran), and subroutines (Fortran) are mapped to a device. The **declare target** directive is a declarative directive.

Syntax



The syntax of the **declare target** directive takes either of the following forms:

```
#pragma omp declare target new-line
declaration-definition-seq
#pragma omp end declare target new-line
```

or


```
#pragma omp declare target clause[ [, ] clause] ... ] new-line
```

where *clause* is one of the following:

[to] (extended-list)

link (*list*)

C / C++

Fortran

The syntax of the **declare target** directive is as follows:

```
!$omp declare target [clause[ [, ] clause] ... ]
```

where *clause* is one of the following:

[to] (extended-list)

link (*list*)

Fortran

Description

The **declare target** directive ensures that procedures and global variables can be executed or accessed on a device. Variables are mapped for all device executions, or for specific device executions through a **link** clause.

If a list item of a **to** clause is a function (C, C++, Fortran) or subroutine (Fortran) then a device-specific version of the routine is created that can be called from a **target** region.

If a list item of a **to** clause is a variable then the original variable is mapped to a corresponding variable in the device data environment of all devices as if it had appeared in a **map** clause with the *map-type to* on the implicit **target data** construct for each device. The list item is never removed from those device data environments as if its reference count is initialized to positive infinity.

The list items of a **link** clause are not mapped by the **declare target** directive. Instead, their mapping is deferred until they are mapped by **target data** or **target** constructs. They are mapped only for such regions.

If an *extended-list* is present with no clause then the **to** clause is assumed for it.

▼ C / C++ ▼

The form of the **declare target** directive that has no clauses and requires a matching **end declare target** directive defines an implicit *extended-list* to an implicit **to** clause. The implicit *extended-list* consists of the variable names of any variable declarations at file or namespace scope that appear between the two directives and of the function names of any function declarations at file, namespace or class scope that appear between the two directives.

▲ C / C++ ▲
▼ Fortran ▼

If a **declare target** does not have any clauses then an implicit *extended-list* to an implicit **to** clause of one item is formed from the name of the enclosing subroutine subprogram, function subprogram or interface body to which it applies.

▲ Fortran ▲

Restrictions

- A threadprivate variable cannot appear in a **declare target** directive.
- A variable declared in a **declare target** directive must have a mappable type.
- The same list item must not appear in both a **to** clause and a **link** clause.

C / C++

- All declarations and definitions for a function must have a **declare target** directive if one is specified for any of them. Otherwise, the result is unspecified.
- The *declaration-definition-seq* defined by a **declare target** directive and an **end declare target** directive must not contain any **declare target** directives.

C / C++

Fortran

- If a list item is a procedure name, it must not be a generic name, procedure pointer or entry name.
- Any **declare target** directive with clauses must appear in a specification part of a subroutine subprogram, function subprogram, program or module.
- Any **declare target** directive without clauses must appear in a specification part of a subroutine subprogram, function subprogram or interface body to which it applies.
- If a **declare target** directive is specified in an interface block for a procedure, it must match a **declare target** directive in the definition of the procedure.
- If an external procedure is a type-bound procedure of a derived type and a **declare target** directive is specified in the definition of the external procedure, such a directive must appear in the interface block that is accessible to the derived type definition.
- If any procedure is declared via a procedure declaration statement that is not in the type-bound procedure part of a derived-type definition, any **declare target** with the procedure name must appear in the same specification part.
- A variable that is part of another variable (as an array or structure element) cannot appear in a **declare target** directive.
- The **declare target** directive must appear in the declaration section of a scoping unit in which the common block or variable is declared. Although variables in common blocks can be accessed by use association or host association, common block names cannot. This means that a common block name specified in a **declare target** directive must be declared to be a common block in the same scoping unit in which the **declare target** directive appears.
- If a **declare target** directive specifying a common block name appears in one program unit, then such a directive must also appear in every other program unit that contains a **COMMON** statement specifying the same name. It must appear after the last such **COMMON** statement in the program unit.

- If a list item is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **declare target** directive in the C program.
- A blank common block cannot appear in a **declare target** directive.
- A variable can only appear in a **declare target** directive in the scope in which it is declared. It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- A variable that appears in a **declare target** directive must be declared in the Fortran scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

2.10.5 teams Construct

Summary

The **teams** construct creates a league of thread teams and the master thread of each team executes the region.

Syntax

C / C++

The syntax of the **teams** construct is as follows:

```
#pragma omp teams [clause[ [, ] clause] ... ] new-line
structured-block
```

where *clause* is one of the following:

```
num_teams (integer-expression)
thread_limit (integer-expression)
default (shared | none)
private (list)
firstprivate (list)
shared (list)
reduction (reduction-identifier : list)
```

C / C++

The syntax of the **teams** construct is as follows:

```
!$omp teams [clause[ [, ] clause] ... ]
structured-block
!$omp end teams
```

where *clause* is one of the following:

```
num_teams (scalar-integer-expression)
thread_limit (scalar-integer-expression)
default (shared | firstprivate | private | none)
private (list)
firstprivate (list)
shared (list)
reduction (reduction-identifier : list)
```

The **end teams** directive denotes the end of the **teams** construct.

Binding

The binding thread set for a **teams** region is the encountering thread.

Description

When a thread encounters a **teams** construct, a league of thread teams is created and the master thread of each thread team executes the **teams** region.

The number of teams created is implementation defined, but is less than or equal to the value specified in the **num_teams** clause.

The maximum number of threads participating in the contention group that each team initiates is implementation defined, but is less than or equal to the value specified in the **thread_limit** clause.

Once the teams are created, the number of teams remains constant for the duration of the **teams** region.

Within a **teams** region, team numbers uniquely identify each team. Team numbers are consecutive whole numbers ranging from zero to one less than the number of teams. A thread may obtain its own team number by a call to the **omp_get_team_num** library routine.

After the teams have completed execution of the **teams** region, the encountering thread resumes execution of the enclosing **target** region.

There is no implicit barrier at the end of a **teams** construct.

Restrictions

Restrictions to the **teams** construct are as follows:

- A program that branches into or out of a **teams** region is non-conforming.
- A program must not depend on any ordering of the evaluations of the clauses of the **teams** directive, or on any side effects of the evaluation of the clauses.
- At most one **thread_limit** clause can appear on the directive. The **thread_limit** expression must evaluate to a positive integer value.
- At most one **num_teams** clause can appear on the directive. The **num_teams** expression must evaluate to a positive integer value.
- If specified, a **teams** construct must be contained within a **target** construct. That **target** construct must contain no statements, declarations or directives outside of the **teams** construct.
- **distribute**, **parallel**, **parallel sections**, **parallel workshare**, and the parallel loop and parallel loop SIMD constructs are the only OpenMP constructs that can be closely nested in the **teams** region.

Cross References

- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section 2.15.3 on page 185.
- **omp_get_num_teams** routine, see Section 3.2.32 on page 261.
- **omp_get_team_num** routine, see Section 3.2.33 on page 263.

2.10.6 distribute Construct

Summary

The **distribute** construct specifies that the iterations of one or more loops will be executed by the thread teams in the context of their implicit tasks. The iterations are distributed across the master threads of all teams that execute the **teams** region to which the **distribute** region binds.

Syntax

C / C++

The syntax of the **distribute** construct is as follows:

```
#pragma omp distribute [clause[ [, ] clause] ... ] new-line  
for-loops
```

Where *clause* is one of the following:

```
private(list)  
firstprivate(list)  
collapse(n)  
dist_schedule(kind[, chunk_size])
```

All associated *for-loops* must have the canonical form described in Section 2.6 on page 52.

C / C++

Fortran

The syntax of the **distribute** construct is as follows:

```
!$omp distribute [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end distribute]
```

Where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
collapse (n)  
dist_schedule (kind[, chunk_size])
```

If an **end distribute** directive is not specified, an **end distribute** directive is assumed at the end of the *do-loops*.

All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end do** directive follows a *do-construct* in which several loop statements share a **DO** termination statement, then the directive can only be specified for the outermost of these **DO** statements.

Fortran

Binding

The binding thread set for a **distribute** region is the set of master threads created by a **teams** construct. A **distribute** region binds to the innermost enclosing **teams** region. Only the threads executing the binding **teams** region participate in the execution of the loop iterations.

Description

The **distribute** construct is associated with a loop nest consisting of one or more loops that follow the directive.

There is no implicit barrier at the end of a **distribute** construct.

The **collapse** clause may be used to specify how many loops are associated with the **distribute** construct. The parameter of the **collapse** clause must be a constant positive integer expression. If no **collapse** clause is present, the only loop that is associated with the **distribute** construct is the one that immediately follows the **distribute** construct.

If more than one loop is associated with the **distribute** construct, then the iteration of all associated loops are collapsed into one larger iteration space. The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space.

If **dist_schedule** is specified, *kind* must be **static**. If specified, iterations are divided into chunks of size *chunk_size*, chunks are assigned to the teams of the league in a round-robin fashion in the order of the team number. When no *chunk_size* is specified, the iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each team of the league. Note that the size of the chunks is unspecified in this case.

When no **dist_schedule** clause is specified, the schedule is implementation defined.

Restrictions

Restrictions to the **distribute** construct are as follows:

- The **distribute** construct inherits the restrictions of the loop construct.
- A **distribute** construct must be closely nested in a **teams** region.

Cross References

- loop construct, see Section 2.7.1 on page 55.
- **teams** construct, see Section 2.10.5 on page 106

2.10.7 distribute simd Construct

Summary

The **distribute simd** construct specifies a loop that will be distributed across the master threads of the **teams** region and executed concurrently using SIMD instructions. The **distribute simd** construct is a composite construct.

Syntax

The syntax of the **distribute simd** construct is as follows:

C / C++

```
#pragma omp distribute simd [clause[ [, ] clause] ... ]  
    for-loops
```

where *clause* can be any of the clauses accepted by the **distribute** or **simd** directives with identical meanings and restrictions.

C / C++

```

!$omp distribute simd [clause[ [, ] clause] ... ]
    do-loops
[!$omp end distribute simd]

```

where *clause* can be any of the clauses accepted by the **distribute** or **simd** directives with identical meanings and restrictions.

If an **end distribute simd** directive is not specified, an **end distribute simd** directive is assumed at the end of the *do-loops*

Description

The **distribute simd** construct will first distribute the iterations of the associated loop(s) according to the semantics of the **distribute** construct and any clauses that apply to the distribute construct. The resulting chunks of iterations will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

Restrictions

- The restrictions for the **distribute** and **simd** constructs apply.
- A list item may appear in a **linear** or **firstprivate** clause but not both.
- A list item may appear in a **linear** or **lastprivate** clause but not both.

Cross References

- **simd** construct, see Section [2.8.1](#) on page [71](#).
- **distribute** construct, see Section [2.10.6](#) on page [109](#).
- Data attribute clauses, see Section [2.15.3](#) on page [185](#).

1 2.10.8 Distribute Parallel Loop Construct

2 Summary

3 The distribute parallel loop construct specifies a loop that can be executed in parallel by multiple
4 threads that are members of multiple teams. The distribute parallel loop construct is a composite
5 construct.

6 Syntax

7 The syntax of the distribute parallel loop construct is as follows:

▼ C / C++ ▼

```
#pragma omp distribute parallel for [clause[ [, ] clause] ... ]  
    for-loops
```

8 where *clause* can be any of the clauses accepted by the **distribute** or parallel loop directives
9 with identical meanings and restrictions.

▲ C / C++ ▲
▼ Fortran ▼

```
!$omp distribute parallel do [clause[ [, ] clause] ... ]  
    do-loops  
[!$omp end distribute parallel do]
```

10 where *clause* can be any of the clauses accepted by the **distribute** or parallel loop directives
11 with identical meanings and restrictions.

12 If an **end distribute parallel do** directive is not specified, an
13 **end distribute parallel do** directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

14 Description

15 The distribute parallel loop construct will first distribute the iterations of the associated loop(s) into
16 chunks according to the semantics of the **distribute** construct and any clauses that apply to the
17 **distribute** construct. Each of these chunks will form a loop. Each resulting loop will then be
18 distributed across the threads within the teams region to which the **distribute** construct binds
19 in a manner consistent with any clauses that apply to the parallel loop construct. The effect of any
20 clause that applies to both constructs is as if it were applied to both constructs separately except the
21 **collapse** clause, which is applied once.

Restrictions

- The restrictions for the **distribute** and parallel loop constructs apply.
- A list item may appear in a **linear** or **firstprivate** clause but not both.
- A list item may appear in a **linear** or **lastprivate** clause but not both.

Cross References

- **distribute** construct, see Section 2.10.6 on page 109.
- Parallel loop construct, see Section 2.11.1 on page 121.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.10.9 Distribute Parallel Loop SIMD Construct

Summary

The distribute parallel loop SIMD construct specifies a loop that can be executed concurrently using SIMD instructions in parallel by multiple threads that are members of multiple teams. The distribute parallel loop SIMD construct is a composite construct.

Syntax

C / C++

The syntax of the distribute parallel loop SIMD construct is as follows:

```
#pragma omp distribute parallel for simd [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **distribute** or parallel loop SIMD directives with identical meanings and restrictions

C / C++

The syntax of the distribute parallel loop SIMD construct is as follows:

```
!$omp distribute parallel do simd [clause[ [, ] clause] ... ]
      do-loops
[/!$omp end distribute parallel do simd]
```

where *clause* can be any of the clauses accepted by the **distribute** or parallel loop SIMD directives with identical meanings and restrictions.

If an **end distribute parallel do simd** directive is not specified, an **end distribute parallel do simd** directive is assumed at the end of the *do-loops*.

Description

The distribute parallel loop SIMD construct will first distribute the iterations of the associated loop(s) according to the semantics of the **distribute** construct and any clauses that apply to the **distribute** construct. The resulting loops will then be distributed across the threads contained within the **teams** region to which the **distribute** construct binds in a manner consistent with any clauses that apply to the parallel loop construct. The resulting chunks of iterations will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

Restrictions

- The restrictions for the **distribute** and parallel loop SIMD constructs apply.
- A list item may appear in a **linear** or **firstprivate** clause but not both.
- A list item may appear in a **linear** or **lastprivate** clause but not both.

Cross References

- **distribute** construct, see Section 2.10.6 on page 109.
- Parallel loop SIMD construct, see Section 2.11.4 on page 125.
- Data attribute clauses, see Section 2.15.3 on page 185.

1 2.10.10 target enter data Construct

2 Summary

3 The **target enter data** directive specifies that variables are mapped to a device data
4 environment. The **target enter data** directive is a stand-alone directive.

5 Syntax

▼ C / C++ ▼

6 The syntax of the **target enter data** construct is as follows:

```
#pragma omp target enter data [clause[[,] clause]...] new-line
```

7 where *clause* is one of the following:

```
8     if([ target enter data :] scalar-expression)  
9     device(integer-expression)  
10    map([ [map-type-modifier[,] map-type : ] list)  
11    depend(dependence-type : list)  
12    nowait
```

▲ C / C++ ▲
▼ Fortran ▼

13 The syntax of the **target enter data** is as follows:

```
!$omp target enter data [clause[[,] clause]...]
```

14 where *clause* is one of the following:

```
15     if([ target enter data :] scalar-logical-expression)  
16     device(scalar-integer-expression)  
17    map([ [map-type-modifier[,] map-type : ] list)  
18    depend(dependence-type : list)  
19    nowait
```

▲ Fortran ▲

Binding

The binding task for a **target enter data** construct is the encountering task.

Description

When a **target enter data** construct is encountered, the list items are mapped to the device data environment according to the **map** clause semantics.

The **target enter data** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target enter data** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target enter data** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated target task is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target enter data** construct.

If there is no **device** clause, the default device is determined by the *default-device-var* ICV.

When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target enter data** directive, or on any side effects of the evaluations of the clauses.
- At least one **map** clause must appear on the directive.
- At most one **device** clause can appear on the directive. The **device** expression must evaluate to a non-negative integer value.
- At most one **if** clause can appear on the directive.
- A *map-type* must be specified in all **map** clauses and must be either **to** or **alloc**.
- The restrictions for the **task** construct apply.

Cross References

- **target data**, see Section 2.10.1 on page 93.
- **target exit data**, see Section 2.10.11 on page 118.
- **map** clause, see Section 2.15.5.1 on page 213.
- *default-device-var*, see Section 2.3.1 on page 36.
- **task**, see Section 2.9.1 on page 81.
- **task scheduling constraints**, see Section 2.9.5 on page 92.
- **if** Clause, see Section 2.12 on page 143.

2.10.11 target exit data Construct

Summary

The **target exit data** directive specifies that list items are unmapped from a device data environment. The **target exit data** directive is a stand-alone directive.

Syntax

C / C++

The syntax of the **target exit data** construct is as follows:

```
#pragma omp target exit data [clause[[,] clause]...] new-line
```

where *clause* is one of the following:

```
if([ target exit data :] scalar-expression)
```

```
device(integer-expression)
```

```
map([ [map-type-modifier[,]] map-type : ] list)
```

```
depend(dependence-type : list)
```

```
nowait
```

C / C++

Fortran

The syntax of the **target exit data** is as follows:


```
!$omp target exit data [clause[[,] clause]...]
```

where clause is one of the following:

```
if([ target exit data :] scalar-logical-expression)
device (scalar-integer-expression)
map([ [map-type-modifier[,]] map-type : ] list)
depend (dependence-type : list)
nowait
```

Fortran

Binding

The binding task for a **target exit data** construct is the encountering task.

Description

When a **target exit data** construct is encountered, the list items in the **map** clauses are unmapped from the device data environment according to the **map** clause semantics.

The **target exit data** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target exit data** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target exit data** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated *target task* is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target exit data** construct.

If there is no **device** clause, the default device is determined by the *default-device-var* ICV.

When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target exit data** directive, or on any side effects of the evaluations of the clauses.
- At least one **map** clause must appear on the directive.
- At most one **device** clause can appear on the directive. The **device** expression must evaluate to a non-negative integer value.
- At most one **if** clause can appear on the directive.
- A *map-type* must be specified in all **map** clauses and must be either **from**, **release**, or **delete**.
- The restrictions for the **task** construct apply.

Cross References

- **target data**, see Section 2.10.1 on page 93.
- **target enter data**, see Section 2.10.10 on page 116.
- **map** clause, see Section 2.15.5.1 on page 213.
- *default-device-var*, see Section 2.3.1 on page 36.
- **task**, see Section 2.9.1 on page 81.
- **task scheduling constraints**, see Section 2.9.5 on page 92.
- **if** Clause, see Section 2.12 on page 143.

2.11 Combined Constructs

Combined constructs are shortcuts for specifying one construct immediately nested inside another construct. The semantics of the combined constructs are identical to that of explicitly specifying the first construct containing one instance of the second construct and no other statements.

Some combined constructs have clauses that are permitted on both constructs that were combined. Where specified, the effect is as if applying the clauses to one or both constructs. If not specified and applying the clause to one construct would result in different program behavior than applying the clause to the other construct then the program's behavior is unspecified.

2.11.1 Parallel Loop Construct

Summary

The parallel loop construct is a shortcut for specifying a **parallel** construct containing one or more associated loops and no other statements.

Syntax

C / C++

The syntax of the parallel loop construct is as follows:

```
#pragma omp parallel for [clause[ [, ] clause] ... ] new-line  
for-loop
```

where *clause* can be any of the clauses accepted by the **parallel** or **for** directives, except the **nowait** clause, with identical meanings and restrictions.

C / C++

Fortran

The syntax of the parallel loop construct is as follows:

```
!$omp parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end parallel do]
```

where *clause* can be any of the clauses accepted by the **parallel** or **do** directives, with identical meanings and restrictions.

If an **end parallel do** directive is not specified, an **end parallel do** directive is assumed at the end of the *do-loops*. **nowait** may not be specified on an **end parallel do** directive.

Fortran

Description

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a loop directive.

Restrictions

- The restrictions for the **parallel** construct and the loop construct apply

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- loop construct, see Section 2.7.1 on page 55.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.2 parallel sections Construct

Summary

The **parallel sections** construct is a shortcut for specifying a **parallel** construct containing one **sections** construct and no other statements.

Syntax

C / C++

The syntax of the **parallel sections** construct is as follows:

```
#pragma omp parallel sections [clause[ [, ] clause] ... ] new-line
{
    [#pragma omp section new-line]
        structured-block
    [#pragma omp section new-line]
        structured-block]
    ...
}
```

where *clause* can be any of the clauses accepted by the **parallel** or **sections** directives, except the **nowait** clause, with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **parallel sections** construct is as follows:

```
!$omp parallel sections [clause[ [, ] clause] ... ]  
    [!$omp section  
        structured-block  
    [!$omp section  
        structured-block]  
    ...  
!$omp end parallel sections
```

where *clause* can be any of the clauses accepted by the **parallel** or **sections** directives, with identical meanings and restrictions.

The last section ends at the **end parallel sections** directive. **nowait** cannot be specified on an **end parallel sections** directive.

Fortran

Description

C / C++

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **sections** directive.

C / C++

Fortran

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **sections** directive, and an **end sections** directive immediately followed by an **end parallel** directive.

Fortran

Restrictions

The restrictions for the **parallel** construct and the **sections** construct apply.

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- **sections** construct, see Section 2.7.2 on page 63.
- Data attribute clauses, see Section 2.15.3 on page 185.

Fortran

2.11.3 parallel workshare Construct

Summary

The **parallel workshare** construct is a shortcut for specifying a **parallel** construct containing one **workshare** construct and no other statements.

Syntax

The syntax of the **parallel workshare** construct is as follows:

```
!$omp parallel workshare [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end parallel workshare
```

where *clause* can be any of the clauses accepted by the **parallel** directive, with identical meanings and restrictions. **nowait** may not be specified on an **end parallel workshare** directive.

Description

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **workshare** directive, and an **end workshare** directive immediately followed by an **end parallel** directive.

Restrictions

The restrictions for the **parallel** construct and the **workshare** construct apply.

1 **Cross References**

- 2 • **parallel** construct, see Section 2.5 on page 45.
- 3 • **workshare** construct, see Section 2.7.4 on page 68.
- 4 • Data attribute clauses, see Section 2.15.3 on page 185.

Fortran

5 **2.11.4 Parallel Loop SIMD Construct**

6 **Summary**

7 The parallel loop SIMD construct is a shortcut for specifying a **parallel** construct containing

8 one loop SIMD construct and no other statement.

9 **Syntax**

C / C++

10 The syntax of the parallel loop SIMD construct is as follows:

```
#pragma omp parallel for simd [clause[ [, ] clause] ... ] new-line
for-loops
```

11 where *clause* can be any of the clauses accepted by the **parallel**, **for** or **simd** directives,

12 except the **nowait** clause, with identical meanings and restrictions.

C / C++

Fortran

13 The syntax of the parallel loop SIMD construct is as follows:

```
!$omp parallel do simd [clause[ [, ] clause] ... ]
do-loops
!$omp end parallel do simd
```

14 where *clause* can be any of the clauses accepted by the **parallel**, **do** or **simd** directives, with

15 identical meanings and restrictions.

16 If an **end parallel do simd** directive is not specified, an **end parallel do simd** directive

17 is assumed at the end of the *do-loops*. **nowait** may not be specified on an

18 **end parallel do simd** directive.

Fortran

Description

The semantics of the parallel loop SIMD construct are identical to explicitly specifying a **parallel** directive immediately followed by a loop SIMD directive. The effect of any clause that applies to both constructs is as if it were applied to the loop SIMD construct and not to the **parallel** construct.

Restrictions

The restrictions for the **parallel** construct and the loop SIMD construct apply.

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- loop SIMD construct, see Section 2.8.3 on page 79.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.5 target parallel Construct

Summary

The **target parallel** construct is a shortcut for specifying a **target** construct containing a **parallel** construct and no other statements.

Syntax

C / C++

The syntax of the **target parallel** construct is as follows:

```
#pragma omp target parallel [clause[ [, ] clause] ... ]  
    structured-block
```

where *clause* can be any of the clauses accepted by the **target** or **parallel** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **target parallel** construct is as follows:


```
!$omp target parallel [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end target parallel
```

where *clause* can be any of the clauses accepted by the **target** or **parallel** directives with identical meanings and restrictions.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a **parallel** directive.

Restrictions

The restrictions for the **target** and **parallel** constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- **target** construct, see Section 2.10.2 on page 95.
- **if** Clause, see Section 2.12 on page 143.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.6 Target Parallel Loop Construct

Summary

The target parallel loop construct is a shortcut for specifying a **target** construct containing a parallel loop construct and no other statements.

Syntax

C / C++

The syntax of the target parallel loop construct is as follows:

```
#pragma omp target parallel for [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **parallel for** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the target parallel loop construct is as follows:

```
!$omp target parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end target parallel do]
```

where *clause* can be any of the clauses accepted by the **target** or **parallel do** directives with identical meanings and restrictions.

If an **end target parallel do** directive is not specified, an **end target parallel do** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a parallel loop directive.

Restrictions

The restrictions for the **target** and parallel loop constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section 2.10.2 on page 95.
- Parallel loop construct, see Section 2.11.1 on page 121.
- **if** Clause, see Section 2.12 on page 143.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.7 Target Parallel Loop SIMD Construct

Summary

The target parallel loop SIMD construct is a shortcut for specifying a **target** construct containing a parallel loop SIMD construct and no other statements.

Syntax

C / C++

The syntax of the target parallel loop SIMD construct is as follows:

```
#pragma omp target parallel for simd [clause[ [, ] clause] ... ]  
    for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **parallel for simd** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the target parallel loop SIMD construct is as follows:

```
!$omp target parallel do simd [clause[ [, ] clause] ... ]  
    do-loops  
[!$omp end target parallel do simd]
```

where *clause* can be any of the clauses accepted by the **target** or **parallel do simd** directives with identical meanings and restrictions.

If an **end target parallel do simd** directive is not specified, an **end target parallel do simd** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a parallel loop SIMD directive.

Restrictions

The restrictions for the **target** and parallel loop SIMD constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section [2.10.2](#) on page [95](#).
- Parallel loop SIMD construct, see Section [2.11.4](#) on page [125](#).
- **if** Clause, see Section [2.12](#) on page [143](#).
- Data attribute clauses, see Section [2.15.3](#) on page [185](#).

2.11.8 target simd Construct

Summary

The **target simd** construct is a shortcut for specifying a **target** construct containing a **simd** construct and no other statements.

Syntax

The syntax of the **target simd** construct is as follows:

```
#pragma omp target simd [clause[ [, ] clause] ... ]
    for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **simd** directives with identical meanings and restrictions.

The syntax of the **target simd** construct is as follows:

```
!$omp target simd [clause[ [, ] clause] ... ]
    do-loops
[!$omp end target simd]
```

where *clause* can be any of the clauses accepted by the **target** or **simd** directives with identical meanings and restrictions.

If an **end target simd** directive is not specified, an **end target simd** directive is assumed at the end of the *do-loops*.

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a **simd** directive.

Restrictions

The restrictions for the **target** and **simd** constructs apply.

Cross References

- **simd** construct, see Section 2.8.1 on page 71.
- **target** construct, see Section 2.10.2 on page 95.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.9 target teams Construct

Summary

The **target teams** construct is a shortcut for specifying a **target** construct containing a **teams** construct and no other statements.

Syntax

C / C++

The syntax of the **target teams** construct is as follows:

```
#pragma omp target teams [clause[ [, ] clause] ... ]  
    structured-block
```

where *clause* can be any of the clauses accepted by the **target** or **teams** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **target teams** construct is as follows:

```
!$omp target teams [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end target teams
```

where *clause* can be any of the clauses accepted by the **target** or **teams** directives with identical meanings and restrictions.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a **teams** directive.

Restrictions

The restrictions for the **target** and **teams** constructs apply.

Cross References

- **target** construct, see Section 2.10.2 on page 95.
- **teams** construct, see Section 2.10.5 on page 106.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.10 teams distribute Construct

Summary

The **teams distribute** construct is a shortcut for specifying a **teams** construct containing a **distribute** construct and no other statements.

Syntax

C / C++

The syntax of the **teams distribute** construct is as follows:

```
#pragma omp teams distribute [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **teams distribute** construct is as follows:

```
!$omp teams distribute [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end teams distribute]
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute** directives with identical meanings and restrictions.

If an **end teams distribute** directive is not specified, an **end teams distribute** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **teams** directive immediately followed by a **distribute** directive. Some clauses are permitted on both constructs.

Restrictions

The restrictions for the **teams** and **distribute** constructs apply.

Cross References

- **teams** construct, see Section 2.10.5 on page 106.
- **distribute** construct, see Section 2.10.6 on page 109.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.11 teams distribute simd Construct

Summary

The **teams distribute simd** construct is a shortcut for specifying a **teams** construct containing a **distribute simd** construct and no other statements.

Syntax

C / C++

The syntax of the **teams distribute simd** construct is as follows:

```
#pragma omp teams distribute simd [clause[ , ] clause] ... ]  
    for-loops
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute simd** directives with identical meanings and restrictions.

C / C++

The syntax of the **teams distribute simd** construct is as follows:

```
!$omp teams distribute simd [clause[ [, ] clause] ... ]  
    do-loops  
[$omp end teams distribute simd]
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute simd** directives with identical meanings and restrictions.

If an **end teams distribute** directive is not specified, an **end teams distribute** directive is assumed at the end of the *do-loops*.

Description

The semantics are identical to explicitly specifying a **teams** directive immediately followed by a **distribute simd** directive. Some clauses are permitted on both constructs.

Restrictions

The restrictions for the **teams** and **distribute simd** constructs apply.

Cross References

- **teams** construct, see Section 2.10.5 on page 106.
- **distribute simd** construct, see Section 2.10.7 on page 111.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.12 target teams distribute Construct

Summary

The **target teams distribute** construct is a shortcut for specifying a **target** construct containing a **teams distribute** construct and no other statements.

Syntax

C / C++

The syntax of the **target teams distribute** construct is as follows:

```
#pragma omp target teams distribute [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the **target teams distribute** construct is as follows:

```
!$omp target teams distribute [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end target teams distribute]
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute** directives with identical meanings and restrictions.

If an **end target teams distribute** directive is not specified, an **end target teams distribute** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a **teams distribute** directive.

Restrictions

The restrictions for the **target** and **teams distribute** constructs apply.

Cross References

- **target** construct, see Section 2.10.1 on page 93.
- **teams distribute** construct, see Section 2.11.10 on page 133.
- Data attribute clauses, see Section 2.15.3 on page 185.

1 2.11.13 target teams distribute simd Construct

2 Summary

3 The **target teams distribute simd** construct is a shortcut for specifying a **target**
4 construct containing a **teams distribute simd** construct and no other statements.

5 Syntax

▼ C / C++ ▼

6 The syntax of the **target teams distribute simd** construct is as follows:

```
#pragma omp target teams distribute simd [clause[ [, ] clause] ... ]  
    for-loops
```

7 where *clause* can be any of the clauses accepted by the **target** or **teams distribute simd**
8 directives with identical meanings and restrictions.

▲ C / C++ ▲
▼ Fortran ▼

9 The syntax of the **target teams distribute simd** construct is as follows:

```
!$omp target teams distribute simd [clause[ [, ] clause] ... ]  
    do-loops  
[!$omp end target teams distribute simd]
```

10 where *clause* can be any of the clauses accepted by the **target** or **teams distribute simd**
11 directives with identical meanings and restrictions.

12 If an **end target teams distribute simd** directive is not specified, an
13 **end target teams distribute simd** directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

14 Description

15 The semantics are identical to explicitly specifying a **target** directive immediately followed by a
16 **teams distribute simd** directive.

17 Restrictions

18 The restrictions for the **target** and **teams distribute simd** constructs apply.

Cross References

- **target** construct, see Section 2.10.1 on page 93.
- **teams distribute simd** construct, see Section 2.11.11 on page 134.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.14 Teams Distribute Parallel Loop Construct

Summary

The teams distribute parallel loop construct is a shortcut for specifying a **teams** construct containing a distribute parallel loop construct and no other statements.

Syntax

C / C++

The syntax of the teams distribute parallel loop construct is as follows:

```
#pragma omp teams distribute parallel for [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute parallel for** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the teams distribute parallel loop construct is as follows:

```
!$omp teams distribute parallel do [clause[ [, ] clause] ... ]  
do-loops  
[ !$omp end teams distribute parallel do ]
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute parallel do** directives with identical meanings and restrictions.

If an **end teams distribute parallel do** directive is not specified, an **end teams distribute parallel do** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **teams** directive immediately followed by a distribute parallel loop directive. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately.

Restrictions

The restrictions for the **teams** and distribute parallel loop constructs apply.

Cross References

- **teams** construct, see Section 2.10.5 on page 106.
- Distribute parallel loop construct, see Section 2.10.8 on page 113.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.15 Target Teams Distribute Parallel Loop Construct

Summary

The target teams distribute parallel loop construct is a shortcut for specifying a **target** construct containing a teams distribute parallel loop construct and no other statements.

Syntax

C / C++

The syntax of the target teams distribute parallel loop construct is as follows:

```
#pragma omp target teams distribute parallel for [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel for** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the target teams distribute parallel loop construct is as follows:

```
!$omp target teams distribute parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end target teams distribute parallel do]
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel do** directives with identical meanings and restrictions.

If an **end target teams distribute parallel do** directive is not specified, an **end target teams distribute parallel do** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a teams distribute parallel loop directive.

Restrictions

The restrictions for the **target** and teams distribute parallel loop constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section [2.10.2](#) on page [95](#).
- Distribute parallel loop construct, see Section [2.11.14](#) on page [138](#).
- Data attribute clauses, see Section [2.15.3](#) on page [185](#).
- **if** Clause, see Section [2.12](#) on page [143](#).

1 2.11.16 Teams Distribute Parallel Loop SIMD Construct

2 Summary

3 The teams distribute parallel loop SIMD construct is a shortcut for specifying a **teams** construct
4 containing a distribute parallel loop SIMD construct and no other statements.

5 Syntax

▼ C / C++ ▼

6 The syntax of the teams distribute parallel loop construct is as follows:

```
#pragma omp teams distribute parallel for simd [clause[ [, ] clause] ... ]  
for-loops
```

7 where *clause* can be any of the clauses accepted by the **teams** or
8 **distribute parallel for simd** directives with identical meanings and restrictions.

▲ C / C++ ▲
▼ Fortran ▼

9 The syntax of the teams distribute parallel loop construct is as follows:

```
!$omp teams distribute parallel do simd [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end teams distribute parallel do simd]
```

10 where *clause* can be any of the clauses accepted by the **teams** or
11 **distribute parallel do simd** directives with identical meanings and restrictions.

12 If an **end teams distribute parallel do simd** directive is not specified, an
13 **end teams distribute parallel do simd** directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

14 Description

15 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a
16 distribute parallel loop SIMD directive. The effect of any clause that applies to both constructs is as
17 if it were applied to both constructs separately.

18 Restrictions

19 The restrictions for the **teams** and distribute parallel loop SIMD constructs apply.

Cross References

- **teams** construct, see Section 2.10.5 on page 106.
- Distribute parallel loop SIMD construct, see Section 2.10.9 on page 114.
- Data attribute clauses, see Section 2.15.3 on page 185.

2.11.17 Target Teams Distribute Parallel Loop SIMD Construct

Summary

The target teams distribute parallel loop SIMD construct is a shortcut for specifying a **target** construct containing a teams distribute parallel loop SIMD construct and no other statements.

Syntax

C / C++

The syntax of the target teams distribute parallel loop SIMD construct is as follows:

```
#pragma omp target teams distribute parallel for simd [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel for simd** directives with identical meanings and restrictions.

C / C++

Fortran

The syntax of the target teams distribute parallel loop SIMD construct is as follows:

```
!$omp target teams distribute parallel do simd [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end target teams distribute parallel do simd]
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel do simd** directives with identical meanings and restrictions.

If an **end target teams distribute parallel do simd** directive is not specified, an **end target teams distribute parallel do simd** directive is assumed at the end of the *do-loops*.

Fortran

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a teams distribute parallel loop SIMD directive.

Restrictions

The restrictions for the **target** and teams distribute parallel loop SIMD constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section 2.10.2 on page 95.
- Teams distribute parallel loop SIMD construct, see Section 2.11.16 on page 141.
- Data attribute clauses, see Section 2.15.3 on page 185.
- **if** Clause, see Section 2.12 on page 143.

2.12 if Clause

Summary

The semantics of an **if** clause are described in the section on the construct to which it applies. The **if** clause *directive-name-modifier* names the associated construct to which an expression applies, and is particularly useful for composite and combined constructs.

Syntax



The syntax of the **if** clause is as follows:

```
if ([ directive-name-modifier : ] scalar-expression)
```

C / C++

Fortran

The syntax of the **if** clause is as follows:

```
if ([ directive-name-modifier : ] scalar-logical-expression)
```

Fortran

Description

The effect of the **if** clause depends on the construct to which it is applied. For combined or composite constructs, the **if** clause only applies to the semantics of the construct named in the *directive-name-modifier* if one is specified. If no *directive-name-modifier* is specified for a combined or composite construct then the **if** clause applies to all constructs to which an **if** clause can apply.

2.13 Master and Synchronization Constructs and Clauses

OpenMP provides the following synchronization constructs:

- the **master** construct.
- the **critical** construct.
- the **barrier** construct.
- the **taskwait** construct.
- the **taskgroup** construct.
- the **atomic** construct.
- the **flush** construct.
- the **ordered** construct.

1 **2.13.1 master Construct**

2 **Summary**

3 The **master** construct specifies a structured block that is executed by the master thread of the team.

4 **Syntax**

▼──────────────────────────────── C / C++ ─────────────────────────────────▼

5 The syntax of the **master** construct is as follows:

```
#pragma omp master new-line
    structured-block
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲

▼──────────────────────────────── Fortran ─────────────────────────────────▼

6 The syntax of the **master** construct is as follows:

```
!$omp master
    structured-block
!$omp end master
```

▲──────────────────────────────── Fortran ─────────────────────────────────▲

7 **Binding**

8 The binding thread set for a **master** region is the current team. A **master** region binds to the
9 innermost enclosing **parallel** region. Only the master thread of the team executing the binding
10 **parallel** region participates in the execution of the structured block of the **master** region.

11 **Description**

12 Other threads in the team do not execute the associated structured block. There is no implied
13 barrier either on entry to, or exit from, the **master** construct.

14 **Restrictions**

▼──────────────────────────────── C++ ─────────────────────────────────▼

- 15
 - A throw executed inside a **master** region must cause execution to resume within the same

16 **master** region, and the same thread that threw the exception must catch it

▲──────────────────────────────── C++ ─────────────────────────────────▲

1 2.13.2 critical Construct

2 Summary

3 The **critical** construct restricts execution of the associated structured block to a single thread at
4 a time.

5 Syntax

▼ C / C++ ▼

6 The syntax of the **critical** construct is as follows:

```
#pragma omp critical [ (name) [hint (hint-expression) ] ] new-line  
    structured-block
```

7 where *hint-expression* is an integer constant expression which evaluates to a valid lock hint (as
8 described in Section 3.3.2 on page 270).

▼ Fortran ▼

9 The syntax of the **critical** construct is as follows:

```
!$omp critical [ (name) [hint (hint-expression) ] ] new-line  
    structured-block  
!$omp end critical [ (name) ]
```

10 where *hint-expression* is a constant expression that evaluates to a scalar value with kind
11 **omp_lock_hint_kind** with a value which is a valid lock hint (as described in Section 3.3.2 on
12 page 270).

▲ Fortran ▲

13 Binding

14 The binding thread set for a **critical** region is all threads in the contention group. The region is
15 executed as if only a single thread at a time among all threads in the contention group is entering
16 the region for execution, without regard to the team(s) to which the threads belong.

1 **Description**

2 An optional *name* may be used to identify the **critical** construct. All **critical** constructs
3 without a name are considered to have the same unspecified name.

▼ C / C++ ▼

4 Identifiers used to identify a **critical** construct have external linkage and are in a name space
5 that is separate from the name spaces used by labels, tags, members, and ordinary identifiers.

▲ C / C++ ▲

▼ Fortran ▼

6 The names of **critical** constructs are global entities of the program. If a name conflicts with
7 any other entity, the behavior of the program is unspecified.

▲ Fortran ▲

8 The threads of a contention group execute the **critical** region as if only one thread of the
9 contention group is executing the **critical** region at a time. The **critical** construct enforces
10 these execution semantics with respect to all **critical** constructs with the same name in all
11 threads in the contention group, not just those threads in the current team.

12 The presence of a **hint** clause does not affect the isolation guarantees provided by the **critical**
13 construct. If no **hint** clause is specified, the effect is as if **hint (omp_lock_hint_none)** had
14 been specified.

15 **Restrictions**

- 16 • If the **hint** clause is specified, the **critical** construct must have a *name*.
17 • If the **hint** clause is specified, each of the **critical** constructs with the same *name* must
18 have a **hint** clause for which the *hint-expression* evaluates to the same value.

▼ C++ ▼

- 19 • A throw executed inside a **critical** region must cause execution to resume within the same
20 **critical** region, and the same thread that threw the exception must catch it.

▲ C++ ▲

▼ Fortran ▼

21 The following restrictions apply to the critical construct:

- 22 • If a *name* is specified on a **critical** directive, the same *name* must also be specified on the
23 **end critical** directive.
24 • If no *name* appears on the **critical** directive, no *name* can appear on the **end critical**
25 directive.

▲ Fortran ▲

Cross References

- `omp_init_lock_with_hint` and `omp_init_nest_lock_with_hint` routines, see Section 3.3.2 on page 270.

2.13.3 barrier Construct

Summary

The **barrier** construct specifies an explicit barrier at the point at which the construct appears. The **barrier** construct is a stand-alone directive.

Syntax

C / C++

The syntax of the **barrier** construct is as follows:

```
#pragma omp barrier new-line
```

C / C++

Fortran

The syntax of the **barrier** construct is as follows:

```
!$omp barrier
```

Fortran

Binding

The binding thread set for a **barrier** region is the current team. A **barrier** region binds to the innermost enclosing **parallel** region.

Description

All threads of the team executing the binding **parallel** region must execute the **barrier** region and complete execution of all explicit tasks bound to this **parallel** region before any are allowed to continue execution beyond the barrier.

The **barrier** region includes an implicit task scheduling point in the current task region.

1 **Restrictions**

2 The following restrictions apply to the **barrier** construct:

- 3 • Each **barrier** region must be encountered by all threads in a team or by none at all, unless
- 4 cancellation has been requested for the innermost enclosing parallel region.
- 5 • The sequence of worksharing regions and **barrier** regions encountered must be the same for
- 6 every thread in a team.

7 **2.13.4 taskwait Construct**

8 **Summary**

9 The **taskwait** construct specifies a wait on the completion of child tasks of the current task. The

10 **taskwait** construct is a stand-alone directive.

11 **Syntax**

▼ C / C++ ▼

12 The syntax of the **taskwait** construct is as follows:

#pragma omp taskwait *newline*

▲ C / C++ ▲

▼ Fortran ▼

13 The syntax of the **taskwait** construct is as follows:

!\$omp taskwait

▲ Fortran ▲

14 **Binding**

15 The **taskwait** region binds to the current task region. The binding thread set of the **taskwait**

16 region is the current team.

1
2
3
4

Description

The **taskwait** region includes an implicit task scheduling point in the current task region. The current task region is suspended at the task scheduling point until all child tasks that it generated before the **taskwait** region complete execution.

5 **2.13.5 taskgroup Construct**

6 **Summary**

7 The **taskgroup** construct specifies a wait on completion of child tasks of the current task and
8 their descendent tasks.

9 **Syntax**

▼──────────────────────────────── C / C++ ─────────────────────────────────▼

10 The syntax of the **taskgroup** construct is as follows:

```
#pragma omp taskgroup new-line  
                  structured-block
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲
▼──────────────────────────────── Fortran ─────────────────────────────────▼

11 The syntax of the **taskgroup** construct is as follows:

```
!$omp taskgroup  
      structured-block  
!$omp end taskgroup
```

▲──────────────────────────────── Fortran ─────────────────────────────────▲

12 **Binding**

13 A **taskgroup** region binds to the current task region. A **taskgroup** region binds to the
14 innermost enclosing **parallel** region.

Description

When a thread encounters a **taskgroup** construct, it starts executing the region. All child tasks generated in the **taskgroup** region and all of their descendants that bind to the same **parallel** region as the **taskgroup** region are part of the *taskgroup set* associated with the **taskgroup** region.

There is an implicit task scheduling point at the end of the **taskgroup** region. The current task is suspended at the task scheduling point until all tasks in the *taskgroup set* complete execution.

Cross References

- Task scheduling, see Section 2.9.5 on page 92.

2.13.6 atomic Construct

Summary

The **atomic** construct ensures that a specific storage location is accessed atomically, rather than exposing it to the possibility of multiple, simultaneous reading and writing threads that may result in indeterminate values.

Syntax

In the following syntax, *atomic-clause* is a clause that indicates the semantics for which atomicity is enforced and is one of the following:

read
write
update
capture

C / C++

The syntax of the **atomic** construct takes one of the following forms:

```
#pragma omp atomic [seq_cst[,]] atomic-clause [[,]seq_cst] new-line  
expression-stmt
```

or

```
#pragma omp atomic [seq_cst] new-line
    expression-stmt
```

or

```
#pragma omp atomic [seq_cst[,]] capture [[,]seq_cst] new-line
    structured-block
```

where *expression-stmt* is an expression statement with one of the following forms:

- If *atomic-clause* is **read**:
 $v = x;$
- If *atomic-clause* is **write**:
 $x = \text{expr};$
- If *atomic-clause* is **update** or not present:
 $x++;$
 $x--;$
 $++x;$
 $--x;$
 $x \text{ binop} = \text{expr};$
 $x = x \text{ binop } \text{expr};$
 $x = \text{expr binop } x;$
- If *atomic-clause* is **capture**:
 $v = x++;$
 $v = x--;$
 $v = ++x;$
 $v = --x;$
 $v = x \text{ binop} = \text{expr};$
 $v = x = x \text{ binop } \text{expr};$
 $v = x = \text{expr binop } x;$

and where *structured-block* is a structured block with one of the following forms:

```
{v = x; x binop = expr; }
{x binop = expr; v = x; }
{v = x; x = x binop expr; }
{v = x; x = expr binop x; }
{x = x binop expr; v = x; }
{x = expr binop x; v = x; }
{v = x; x = expr; }
```

```

1      {v = x; x++; }
2      {v = x; ++x; }
3      {++x; v = x; }
4      {x++; v = x; }
5      {v = x; x--; }
6      {v = x; --x; }
7      {--x; v = x; }
8      {x--; v = x; }

```

9 In the preceding expressions:

- 10 • x and v (as applicable) are both *l-value* expressions with scalar type.
- 11 • During the execution of an atomic region, multiple syntactic occurrences of x must designate the
- 12 same storage location.
- 13 • Neither of v and $expr$ (as applicable) may access the storage location designated by x .
- 14 • Neither of x and $expr$ (as applicable) may access the storage location designated by v .
- 15 • $expr$ is an expression with scalar type.
- 16 • $binop$ is one of $+$, $*$, $-$, $/$, $\&$, \wedge , $|$, $<<$, or $>>$.
- 17 • $binop$, $binop=$, $++$, and $--$ are not overloaded operators.
- 18 • The expression $x \text{ binop } expr$ must be numerically equivalent to $x \text{ binop } (expr)$. This requirement
- 19 is satisfied if the operators in $expr$ have precedence greater than $binop$, or by using parentheses
- 20 around $expr$ or subexpressions of $expr$.
- 21 • The expression $expr \text{ binop } x$ must be numerically equivalent to $(expr) \text{ binop } x$. This requirement
- 22 is satisfied if the operators in $expr$ have precedence equal to or greater than $binop$, or by using
- 23 parentheses around $expr$ or subexpressions of $expr$.
- 24 • For forms that allow multiple occurrences of x , the number of times that x is evaluated is
- 25 unspecified.



26 The syntax of the **atomic** construct takes any of the following forms:

```

!$omp atomic [seq_cst[,]] read [[,]seq_cst]
    capture-statement
[!$omp end atomic]

```

1 or

```
!$omp atomic [seq_cst[,]] write [[,]seq_cst]
    write-statement
[!$omp end atomic]
```

2 or

```
!$omp atomic [seq_cst[,]] update [[,]seq_cst]
    update-statement
[!$omp end atomic]
```

3 or

```
!$omp atomic [seq_cst]
    update-statement
[!$omp end atomic]
```

4 or

```
!$omp atomic [seq_cst[,]] capture [[,]seq_cst]
    update-statement
    capture-statement
!$omp end atomic
```

5 or

```
!$omp atomic [seq_cst[,]] capture [[,]seq_cst]
    capture-statement
    update-statement
!$omp end atomic
```

6 or

```
!$omp atomic [seq_cst[,]] capture [[,]seq_cst]
    capture-statement
    write-statement
!$omp end atomic
```

7 where *write-statement* has the following form (if *atomic-clause* is **capture** or **write**):

$x = \text{expr}$

where *capture-statement* has the following form (if *atomic-clause* is **capture** or **read**):

$v = x$

and where *update-statement* has one of the following forms (if *atomic-clause* is **update**, **capture**, or not present):

$x = x \text{ operator } \text{expr}$

$x = \text{expr operator } x$

$x = \text{intrinsic_procedure_name } (x, \text{expr_list})$

$x = \text{intrinsic_procedure_name } (\text{expr_list}, x)$

In the preceding statements:

- x and v (as applicable) are both scalar variables of intrinsic type.
- x must not have the **ALLOCATABLE** attribute. .
- During the execution of an atomic region, multiple syntactic occurrences of x must designate the same storage location.
- None of v , expr and expr_list (as applicable) may access the same storage location as x .
- None of x , expr and expr_list (as applicable) may access the same storage location as v .
- expr is a scalar expression.
- expr_list is a comma-separated, non-empty list of scalar expressions. If *intrinsic_procedure_name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must appear in expr_list .
- *intrinsic_procedure_name* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- *operator* is one of **+**, *****, **-**, **/**, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**
- The expression $x \text{ operator } \text{expr}$ must be numerically equivalent to $x \text{ operator } (\text{expr})$. This requirement is satisfied if the operators in expr have precedence greater than *operator*, or by using parentheses around expr or subexpressions of expr .
- The expression $\text{expr operator } x$ must be numerically equivalent to $(\text{expr}) \text{ operator } x$. This requirement is satisfied if the operators in expr have precedence equal to or greater than *operator*, or by using parentheses around expr or subexpressions of expr .
- *intrinsic_procedure_name* must refer to the intrinsic procedure name and not to other program entities.

- *operator* must refer to the intrinsic operator and not to a user-defined operator.
- All assignments must be intrinsic assignments.
- For forms that allow multiple occurrences of *x*, the number of times that *x* is evaluated is unspecified.

Fortran

Binding

The binding thread set for an atomic region is all threads in the contention group. **atomic** regions enforce exclusive access with respect to other **atomic** regions that access the same storage location *x* among all threads in the contention group without regard to the teams to which the threads belong.

Description

The **atomic** construct with the **read** clause forces an atomic read of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **write** clause forces an atomic write of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **update** clause forces an atomic update of the location designated by *x* using the designated operator or intrinsic. Note that when no clause is present, the semantics are equivalent to atomic update. Only the read and write of the location designated by *x* are performed mutually atomically. The evaluation of *expr* or *expr_list* need not be atomic with respect to the read or write of the location designated by *x*. No task scheduling points are allowed between the read and the write of the location designated by *x*.

The **atomic** construct with the **capture** clause forces an atomic update of the location designated by *x* using the designated operator or intrinsic while also capturing the original or final value of the location designated by *x* with respect to the atomic update. The original or final value of the location designated by *x* is written in the location designated by *v* depending on the form of the **atomic** construct structured block or statements following the usual language semantics. Only the read and write of the location designated by *x* are performed mutually atomically. Neither the evaluation of *expr* or *expr_list*, nor the write to the location designated by *v* need be atomic with respect to the read or write of the location designated by *x*. No task scheduling points are allowed between the read and the write of the location designated by *x*.

Any **atomic** construct with a **seq_cst** clause forces the atomically performed operation to include an implicit flush operation without a list.

Note – As with other implicit flush regions, Section 1.4.4 on page 19 reduces the ordering that must be enforced. The intent is that, when the analogous operation exists in C++11 or C11, a sequentially consistent **atomic** construct has the same semantics as a **memory_order_seq_cst** atomic operation in C++11/C11. Similarly, a non-sequentially consistent **atomic** construct has the same semantics as a **memory_order_relaxed** atomic operation in C++11/C11.

Unlike non-sequentially consistent **atomic** constructs, sequentially consistent **atomic** constructs preserve the interleaving (sequentially consistent) behavior of correct, data-race-free programs. However, they are not designed to replace the **flush** directive as a mechanism to enforce ordering for non-sequentially consistent **atomic** constructs, and attempts to do so require extreme caution. For example, a sequentially consistent **atomic write** construct may appear to be reordered with a subsequent non-sequentially consistent **atomic write** construct, since such reordering would not be observable by a correct program if the second write were outside an **atomic** directive.

For all forms of the **atomic** construct, any combination of two or more of these **atomic** constructs enforces mutually exclusive access to the locations designated by x . To avoid race conditions, all accesses of the locations designated by x that could potentially occur in parallel must be protected with an **atomic** construct.

atomic regions do not guarantee exclusive access with respect to any accesses outside of **atomic** regions to the same storage location x even if those accesses occur during a **critical** or **ordered** region, while an OpenMP lock is owned by the executing task, or during the execution of a **reduction** clause.

However, other OpenMP synchronization can ensure the desired exclusive access. For example, a barrier following a series of atomic updates to x guarantees that subsequent accesses do not form a race with the atomic accesses.

A compliant implementation may enforce exclusive access between **atomic** regions that update different storage locations. The circumstances under which this occurs are implementation defined.

If the storage location designated by x is not size-aligned (that is, if the byte alignment of x is not a multiple of the size of x), then the behavior of the **atomic** region is implementation defined.

Restrictions

The following restrictions apply to the **atomic** construct:

- At most one **seq_cst** clause may appear on the construct.

▼ C / C++ ▼

- All atomic accesses to the storage locations designated by *x* throughout the program are required to have a compatible type.

▲ C / C++ ▲

▼ Fortran ▼

- All atomic accesses to the storage locations designated by *x* throughout the program are required to have the same type and type parameters.

▲ Fortran ▲

Cross References

- **critical** construct, see Section 2.13.2 on page 146.
- **barrier** construct, see Section 2.13.3 on page 148.
- **flush** construct, see Section 2.13.7 on page 158.
- **ordered** construct, see Section 2.13.8 on page 163.
- **reduction** clause, see Section 2.15.3.6 on page 198.
- lock routines, see Section 3.3 on page 267.

2.13.7 flush Construct

Summary

The **flush** construct executes the OpenMP flush operation. This operation makes a thread's temporary view of memory consistent with memory, and enforces an order on the memory operations of the variables explicitly specified or implied. See the memory model description in Section 1.4 on page 16 for more details. The **flush** construct is a stand-alone directive.

1 **Syntax**

▼ C / C++ ▼

2 The syntax of the **flush** construct is as follows:

```
#pragma omp flush [ (list) ] new-line
```

▲ C / C++ ▲

▼ Fortran ▼

3 The syntax of the **flush** construct is as follows:

```
!$omp flush [ (list) ]
```

▲ Fortran ▲

4 **Binding**

5 The binding thread set for a **flush** region is the encountering thread. Execution of a **flush**
6 region affects the memory and the temporary view of memory of only the thread that executes the
7 region. It does not affect the temporary view of other threads. Other threads must themselves
8 execute a flush operation in order to be guaranteed to observe the effects of the encountering
9 thread's flush operation

Description

A **flush** construct without a list, executed on a given thread, operates as if the whole thread-visible data state of the program, as defined by the base language, is flushed. A **flush** construct with a list applies the flush operation to the items in the list, and does not return until the operation is complete for all specified list items. An implementation may implement a **flush** with a list by ignoring the list, and treating it the same as a **flush** without a list.

C / C++

If a pointer is present in the list, the pointer itself is flushed, not the memory block to which the pointer refers.

C / C++

Fortran

If the list item or a subobject of the list item has the **POINTER** attribute, the allocation or association status of the **POINTER** item is flushed, but the pointer target is not. If the list item is a Cray pointer, the pointer is flushed, but the object to which it points is not. If the list item is of type **C_PTR**, the variable is flushed, but the storage that corresponds to that address is not flushed. If the list item or the subobject of the list item has the **ALLOCATABLE** attribute and has an allocation status of currently allocated, the allocated variable is flushed; otherwise the allocation status is flushed.

Fortran

Note – Use of a **flush** construct with a list is extremely error prone and users are strongly discouraged from attempting it. The following examples illustrate the ordering properties of the flush operation. In the following incorrect pseudocode example, the programmer intends to prevent simultaneous execution of the protected section by the two threads, but the program does not work properly because it does not enforce the proper ordering of the operations on variables **a** and **b**. Any shared data accessed in the protected section is not guaranteed to be current or consistent during or after the protected section. The atomic notation in the pseudocode in the following two examples indicates that the accesses to **a** and **b** are **ATOMIC** writes and captures. Otherwise both examples would contain data races and automatically result in unspecified behavior.

Incorrect example:

a = b = 0

thread 1

thread 2

atomic(b = 1)

atomic(a = 1)

flush (b)

flush (a)

flush (a)

flush (b)

atomic(tmp = a)

atomic(tmp = b)

if (tmp == 0) then

if (tmp == 0) then

protected section

protected section

end if

end if

The problem with this example is that operations on variables **a** and **b** are not ordered with respect to each other. For instance, nothing prevents the compiler from moving the flush of **b** on thread 1 or the flush of **a** on thread 2 to a position completely after the protected section (assuming that the protected section on thread 1 does not reference **b** and the protected section on thread 2 does not reference **a**). If either re-ordering happens, both threads can simultaneously execute the protected section.

The following pseudocode example correctly ensures that the protected section is executed by not more than one of the two threads at any one time. Notice that execution of the protected section by neither thread is considered correct in this example. This occurs if both flushes complete prior to either thread executing its **if** statement.

Correct example:

a = b = 0

thread 1

thread 2

atomic(b = 1)

atomic(a = 1)

flush (a, b)

flush (a, b)

atomic(tmp = a)

atomic(tmp = b)

if (tmp == 0) then

if (tmp == 0) then

protected section

protected section

end if

end if

1 The compiler is prohibited from moving the flush at all for either thread, ensuring that the
2 respective assignment is complete and the data is flushed before the **if** statement is executed.

3 A **flush** region without a list is implied at the following locations:

- 4 • During a barrier region.
- 5 • At entry to a **target update** region whose corresponding construct has a **to** clause.
- 6 • At exit from a **target update** region whose corresponding construct has a **from** clause.
- 7 • At entry to and exit from **parallel**, **critical**, **ordered**, **target** and **target**
8 **data** regions.
- 9 • At entry to a **target enter data** region.
- 10 • At exit from a **target exit data** region.
- 11 • At exit from worksharing regions unless a **nowait** is present.
- 12 • At entry to and exit from the **atomic** operation (read, write, update, or capture) performed in a
13 sequentially consistent atomic region.
- 14 • During **omp_set_lock** and **omp_unset_lock** regions.
- 15 • During **omp_test_lock**, **omp_set_nest_lock**, **omp_unset_nest_lock** and
16 **omp_test_nest_lock** regions, if the region causes the lock to be set or unset.
- 17 • Immediately before and immediately after every task scheduling point.
- 18 • During a **cancel** or **cancellation point** region, if *cancel-var* is *true* and cancellation
19 has been activated.

20 A **flush** region with a list is implied at the following locations:

- 21 • At entry to and exit from the **atomic** operation (read, write, update, or capture) performed in a
22 non-sequentially consistent **atomic** region, where the list contains only the storage location
23 designated as *x* according to the description of the syntax of the **atomic** construct in
24 Section 2.13.6 on page 151.

25 Note – A **flush** region is not implied at the following locations:

- 26 • At entry to worksharing regions.
 - 27 • At entry to or exit from a **master** region.
-

1 2.13.8 ordered Construct

2 Summary

3 The **ordered** construct specifies a structured block in a loop region that will be executed in the
4 order of the loop iterations. This sequentializes and orders the code within an **ordered** region
5 while allowing code outside the region to run in parallel.

6 Syntax

C / C++

7 The syntax of the **ordered** construct is as follows:

```
#pragma omp ordered [clause[ [, ] clause] ... ] new-line
    structured-block
```

8 where *clause* is one of the following:

9 **threads**

10 **simd**

11 or

```
#pragma omp ordered clause [[[ , ] clause] ... ] new-line
```

12 where *clause* is:

13 **depend** (*dependence-type* [: *vec*])

C / C++

The syntax of the **ordered** construct is as follows:

```
!$omp ordered [clause [ , clause ] ... ]
      structured-block
!$omp end ordered
```

where *clause* is one of the following:

threads

simd

or

```
!$omp ordered clause [[[ , clause ] ... ]
```

where *clause* is:

```
depend (dependence-type [ : vec ])
```

If the **depend** clause is specified, the **ordered** construct is a stand-alone directive.

Binding

The binding thread set for an **ordered** region is the current team. An **ordered** region binds to the innermost enclosing loop region. **ordered** regions that bind to different loop regions execute independently of each other.

Description

If no clause is specified, the **ordered** construct behaves as if the **threads** clause had been specified. If the **threads** clause is specified, the threads in the team executing the loop region execute **ordered** regions sequentially in the order of the loop iterations. If any **depend** clauses are specified then those clauses specify the order in which the threads in the team execute **ordered** regions. When the thread executing the first iteration of the loop encounters an **ordered** construct, it can enter the **ordered** region without waiting. When a thread executing any subsequent iteration encounters an **ordered** region, it waits at the beginning of that ordered region until execution of all **ordered** regions belonging to all previous iterations or iterations specified by the **depend** clauses have completed. If the **simd** clause is specified, the **ordered** regions encountered by any thread will use only a single SIMD lane to execute the **ordered** regions in the order of the loop iterations.

Restrictions

Restrictions to the **ordered** construct are as follows:

- The loop region to which an **ordered** region without any clause or with a **threads** clause binds must have an **ordered** clause without the parameter specified on the corresponding loop directive.
- The loop region to which an **ordered** region with any **depend** clauses binds must have an **ordered** clause with the parameter specified on the corresponding loop directive.
- An **ordered** construct with the **depend** clause specified must be closely nested inside a loop (or parallel loop) construct.
- An **ordered** construct with the **depend** clause specified must not be closely nested in a **simd** construct.
- During execution of an iteration of a loop or a loop nest within a loop region, a thread must not execute more than one ordered region that arises from an **ordered** clause with no parameter and binds to the same loop region.

C++

- A throw executed inside a **ordered** region must cause execution to resume within the same **ordered** region, and the same thread that threw the exception must catch it.

C++

Cross References

- loop construct, see Section [2.7.1](#) on page [55](#).
- parallel loop construct, see Section [2.11.1](#) on page [121](#).
- **depend** Clause, see Section [2.13.9](#) on page [165](#)

2.13.9 depend Clause

Summary

The **depend** clause enforces additional constraints on the scheduling of tasks or loop iterations. These constraints establish dependences only between sibling tasks or between loop iterations.

Syntax

The syntax of the **depend** clause is as follows:

```
depend (dependence-type : list)
```

or

```
depend (dependence-type [: vec])
```

where *vec* is the iteration vector.

Description

Task dependences are derived from the *dependence-type* of a **depend** clause and its list items, where *dependence-type* is one of the following:

The **in** *dependence-type*. The generated task will be a dependent task of all previously generated sibling tasks that reference at least one of the list items in an **out** or **inout** *dependence-type* list.

The **out** and **inout** *dependence-types*. The generated task will be a dependent task of all previously generated sibling tasks that reference at least one of the list items in an **in**, **out**, or **inout** *dependence-type* list.

Fortran

If a list item has the **ALLOCATABLE** attribute and its allocation status is "not currently allocated", the behavior is unspecified. If a list item has the **POINTER** attribute and its association status is disassociated or undefined, the behavior is unspecified.

Fortran

The list items that appear in the **depend** clause may include array sections.

Note – The enforced task dependence establishes a synchronization of memory accesses performed by a dependent task with respect to accesses performed by the predecessor tasks. However, it is the responsibility of the programmer to synchronize properly with respect to other concurrent accesses that occur outside of those tasks.

1 Loop dependences are derived from the *dependence-type* of a **depend** clause and the iteration
2 vector *vec*.

3 The **source** *dependence-type* specifies the completion of cross-iteration dependences that arise
4 from the current iteration.

5 The **sink** *dependence-type* specifies a sink of cross-iteration dependences. The current iteration is
6 blocked until all source iterations specified complete execution.

7 The iteration vector *vec* must have the form of $(x_1 \pm d_1, x_2 \pm d_2, \dots, x_n \pm d_n)$, where *n* is the
8 value specified by the **ordered** clause in the loop directive, x_i denotes the loop iteration variable
9 of the *i*-th nested loop associated with the loop directive, and d_i is a constant non-negative integer.
10 If the iteration vector *vec* indicates a lexicographically later iteration, it can cause a deadlock. If the
11 iteration vector *vec* indicates an invalid iteration, the **ordered** construct with the **depend** clause
12 is ignored.

13 Restrictions

14 Restrictions to the **depend** clause are as follows:

- 15 • List items used in **depend** clauses of the same task or sibling tasks must indicate identical
16 storage or disjoint storage.
- 17 • List items used in **depend** clauses cannot be zero-length array sections.
- 18 • A variable that is part of another variable (such as an element of a structure) but is not an array
19 element or an array section cannot appear in a **depend** clause.
- 20 • If the *dependence-type* is **source**, the iteration vector *vec* must not be specified.

21 Cross References

- 22 • Array sections, see Section [2.4](#) on page [43](#).
- 23 • **task** construct, see Section [2.9.1](#) on page [81](#).
- 24 • Task scheduling constraints, see Section [2.9.5](#) on page [92](#).
- 25 • **ordered** construct, see Section [2.13.8](#) on page [163](#).

1 2.14 Cancellation Constructs

2 2.14.1 `cancel` Construct

3 Summary

4 The **`cancel`** construct activates cancellation of the innermost enclosing region of the type
5 specified. The **`cancel`** construct is a stand-alone directive.

6 Syntax

▼ C / C++ ▼

7 The syntax of the **`cancel`** construct is as follows:

```
#pragma omp cancel construct-type-clause[ , if-clause] new-line
```

8 where *construct-type-clause* is one of the following:

```
parallel  
sections  
for  
taskgroup
```

13 and *if-clause* is

```
if (scalar-expression)
```

▲ C / C++ ▲

The syntax of the **cancel** construct is as follows:

```
!$omp cancel construct-type-clause [ , ] if-clause
```

where *construct-type-clause* is one of the following:

parallel

sections

do

taskgroup

and *if-clause* is

```
if (scalar-logical-expression)
```

Binding

The binding thread set of the **cancel** region is the current team. The binding region of the **cancel** region is the innermost enclosing region of the type corresponding to the *construct-type-clause* specified in the directive (that is, the innermost **parallel**, **sections**, loop, or **taskgroup** region).

Description

The **cancel** construct activates cancellation of the binding region only if *cancel-var* is *true*, in which case the **cancel** construct causes the encountering task to continue execution at the end of the binding region if *construct-type-clause* is **parallel**, **for**, **do**, or **sections**. If *cancel-var* is *true* and *construct-type-clause* is **taskgroup** the encountering task continues execution at the end of the current task region. If *cancel-var* is *false*, the **cancel** construct is ignored.

Threads check for active cancellation only at cancellation points that are implied at the following locations:

- **cancel** regions;
- **cancellation point** regions;
- **barrier** regions;
- implicit barriers regions.

When a thread reaches one of the above cancellation points and if *cancel-var* is *true*, then:

- If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is **parallel**, **for**, **do**, or **sections**, the thread continues execution at the end of the canceled region if cancellation has been activated for the innermost enclosing region of the type specified.
- If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is **taskgroup**, the encountering task checks for active cancellation of one of the taskgroup sets to which the encountering task belongs and continues execution at the end of the current task region if cancellation has been activated.
- If the encountering task is at a barrier region, the encountering task checks for active cancellation of the innermost enclosing **parallel** region. If cancellation has been activated, then the encountering task continues execution at the end of the canceled region.

Note – If one thread activates cancellation and another thread encounters a cancellation point, the absolute order of execution between the two threads is non-deterministic. Whether the thread that encounters a cancellation point detects the activated cancellation depends on the underlying hardware and operating system.

When cancellation of tasks is activated through the **cancel taskgroup** construct, the tasks that belong to *taskgroup set* of the innermost enclosing **taskgroup** region will be canceled. The task that encountered the **cancel taskgroup** construct continues execution at the end of its **task** region, which implies completion of that task. Any task that belongs to the innermost enclosing **taskgroup** and has already begun execution must run to completion or until a cancellation point is reached. Upon reaching a cancellation point and if cancellation is active, the task continues execution at the end of its **task** region, which implies the task's completion. Any task that belongs

to the innermost enclosing **taskgroup** and that has not begun execution may be discarded, which implies its completion.

When cancellation is active for a **parallel**, **sections**, **for**, or **do** region, each thread of the binding thread set resumes execution at the end of the canceled region if a cancellation point is encountered. If the canceled region is a **parallel** region, any tasks that have been created by a **task** construct and their descendent tasks are canceled according to the above **taskgroup** cancellation semantics. If the canceled region is a **sections**, **for**, or **do** region, no task cancellation occurs.

▼ C++ ▼

The usual C++ rules for object destruction are followed when cancellation is performed.

▲ C++ ▲

▼ Fortran ▼

All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside the canceled construct are deallocated.

▲ Fortran ▲

▼

Note – The programmer is responsible for releasing locks and other synchronization data structures that might cause a deadlock when a **cancel** construct is encountered and blocked threads cannot be canceled. The programmer is also responsible for ensuring proper synchronization to avoid deadlocks that might arise from cancellation of OpenMP regions that contain OpenMP synchronization constructs.

▲

If the canceled construct contains a **reduction** or **lastprivate** clause, the final value of the **reduction** or **lastprivate** variable is undefined.

When an **if** clause is present on a **cancel** construct and the **if** expression evaluates to *false*, the **cancel** construct does not activate cancellation. The cancellation point associated with the **cancel** construct is always encountered regardless of the value of the **if** expression.

Restrictions

The restrictions to the **cancel** construct are as follows:

- The behavior for concurrent cancellation of a region and a region nested within it is unspecified.
- If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a **task** construct and the **cancel** region must be strictly closely nested inside a **taskgroup** region. If *construct-type-clause* is **sections**, the **cancel** construct must be closely nested inside a **sections** or **section** construct. Otherwise, the **cancel** construct must be closely nested inside an OpenMP construct that matches the type specified in *construct-type-clause* of the **cancel** construct.
- If *construct-type-clause* is **taskgroup** and the **cancel** region is not strictly closely nested inside a **taskgroup** region, then the behavior is unspecified.
- A worksharing construct that is canceled must not have a **nowait** clause.
- A loop construct that is canceled must not have an **ordered** clause.
- A construct that may be subject to cancellation must not encounter an orphaned cancellation point. That is, a cancellation point must only be encountered within that construct and must not be encountered elsewhere in its region.

Cross References

- *cancel-var*, see Section 2.3.1 on page 36.
- **cancellation point** construct, see Section 2.14.2 on page 172.
- **omp_get_cancellation** routine, see Section 3.2.9 on page 237.

2.14.2 cancellation point Construct

Summary

The **cancellation point** construct introduces a user-defined cancellation point at which implicit or explicit tasks check if cancellation of the innermost enclosing region of the type specified has been activated. The **cancellation point** construct is a stand-alone directive.

Syntax



The syntax of the **cancellation point** construct is as follows:

```
#pragma omp cancellation point construct-type-clause new-line
```

where *construct-type-clause* is one of the following:

```
parallel
sections
for
taskgroup
```

C / C++

Fortran

The syntax of the **cancellation point** construct is as follows:

```
!$omp cancellation point construct-type-clause
```

where *construct-type-clause* is one of the following:

```
parallel
sections
do
taskgroup
```

Fortran

Binding

The binding thread set of the **cancellation point** construct is the current team. The binding task set of the **cancellation point** region is the current task region. The binding region of the **cancellation point** region is the innermost enclosing region of the type corresponding to the *construct-type-clause* specified in the directive (that is, the innermost **parallel**, **sections**, loop, or **taskgroup** region).

Description

This directive introduces a user-defined cancellation point at which an implicit or explicit task must check if cancellation of the innermost enclosing region of the type specified in the clause has been requested. This construct does not implement a synchronization between threads or tasks.

When an implicit or explicit task reaches a user-defined cancellation point and if *cancel-var* is *true*, then:

- If the *construct-type-clause* of the encountered **cancellation point** construct is **parallel**, **for**, **do**, or **sections**, the thread continues execution at the end of the canceled region if cancellation has been activated for the innermost enclosing region of the type specified.
- If the *construct-type-clause* of the encountered **cancellation point** construct is **taskgroup**, the encountering task checks for active cancellation of all taskgroup sets to which the encountering task belongs and continues execution at the end of the current task region if cancellation has been activated for any of them.

Restrictions

- A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be closely nested inside a **task** construct and the **cancellation point** region must be strictly closely nested inside a **taskgroup** region. A **cancellation point** construct for which *construct-type-clause* is **sections** must be closely nested inside a **sections** or **section** construct. Otherwise, a **cancellation point** construct must be closely nested inside an OpenMP construct that matches the type specified in *construct-type-clause*.
- An OpenMP program with orphaned **cancellation point** constructs is non-conforming.

Cross References

- *cancel-var*, see Section 2.3.1 on page 36.
- **cancel** construct, see Section 2.14.1 on page 168.
- **omp_get_cancellation** routine, see Section 3.2.9 on page 237.

2.15 Data Environment

This section presents a directive and several clauses for controlling the data environment during the execution of **target**, **teams**, **parallel**, **task**, **taskloop**, **simd**, and worksharing regions.

- Section 2.15.1 on page 175 describes how the data-sharing attributes of variables referenced in **target**, **teams**, **parallel**, **task**, **taskloop**, **simd**, and worksharing regions are determined.
- The **threadprivate** directive, which is provided to create threadprivate memory, is described in Section 2.15.2 on page 179.
- Clauses that may be specified on directives to control the data-sharing attributes of variables referenced in **target**, **teams**, **parallel**, **task**, **taskloop**, **simd**, or worksharing constructs are described in Section 2.15.3 on page 185
- Clauses that may be specified on directives to copy data values from private or threadprivate variables on one thread to the corresponding variables on other threads in the team are described in Section 2.15.4 on page 208.
- Clauses that may be specified on directives to control the data-mapping of variables to a device data environment are described in Section 2.15.5.1 on page 213.

2.15.1 Data-sharing Attribute Rules

This section describes how the data-sharing attributes of variables referenced in **target**, **parallel**, **task**, **taskloop**, **simd**, and worksharing regions are determined. The following two cases are described separately:

- Section 2.15.1.1 on page 175 describes the data-sharing attribute rules for variables referenced in a construct.
- Section 2.15.1.2 on page 179 describes the data-sharing attribute rules for variables referenced in a region, but outside any construct.

2.15.1.1 Data-sharing Attribute Rules for Variables Referenced in a Construct

The data-sharing attributes of variables that are referenced in a construct can be *predetermined*, *explicitly determined*, or *implicitly determined*, according to the rules outlined in this section.

Specifying a variable on a **firstprivate**, **lastprivate**, **linear**, **reduction**, or **copyprivate** clause of an enclosed construct causes an implicit reference to the variable in the enclosing construct. Specifying a variable on a **map** clause of an enclosed construct may cause an implicit reference to the variable in the enclosing construct. Such implicit references are also subject to the data-sharing attribute rules outlined in this section.

Certain variables and objects have predetermined data-sharing attributes as follows:

C / C++

- Variables appearing in **threadprivate** directives are threadprivate.
- Variables with automatic storage duration that are declared in a scope inside the construct are private.
- Objects with dynamic storage duration are shared.
- Static data members are shared.
- The loop iteration variable(s) in the associated *for-loop(s)* of a **for** or **parallel for** construct is (are) private.
- The loop iteration variable in the associated *for-loop* of a **simd** construct with just one associated *for-loop* is linear with a *constant-linear-step* that is the increment of the associated *for-loop*.
- The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple associated *for-loops* are lastprivate.
- Variables with static storage duration that are declared in a scope inside the construct are shared.
- If an array section is a list item in a **map** clause on the **target** construct, and the array section is derived from a variable for which the type is pointer, reference to array or reference to pointer then that variable is private.

C / C++

Fortran

- Variables and common blocks appearing in **threadprivate** directives are threadprivate.
- The loop iteration variable(s) in the associated *do-loop(s)* of a **do** or **parallel do** construct is (are) private.
- The loop iteration variable in the associated *do-loop* of a **simd** construct with just one associated *do-loop* is linear with a *constant-linear-step* that is the increment of the associated *do-loop*.
- The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple associated *do-loops* are lastprivate.
- A loop iteration variable for a sequential loop in a **parallel**, **task**, or **taskloop** construct is private in the innermost such construct that encloses the loop.
- Implied-do indices and **forall** indices are private.
- Cray pointees have the same the data-sharing attribute as the storage with which their Cray pointers are associated.
- Assumed-size arrays are shared.
- An associate name preserves the association with the selector established at the **ASSOCIATE** statement.

Fortran

Variables with predetermined data-sharing attributes may not be listed in data-sharing attribute clauses, except for the cases listed below. For these exceptions only, listing a predetermined variable in a data-sharing attribute clause is allowed and overrides the variable's predetermined data-sharing attributes.

C / C++

- The loop iteration variable(s) in the associated *for-loop(s)* of a **for**, **parallel for**, or **taskloop** construct may be listed in a **private** or **lastprivate** clause.
- The loop iteration variable in the associated *for-loop* of a **simd** construct with just one associated *for-loop* may be listed in a **linear** clause with a *constant-linear-step* that is the increment of the associated *for-loop*.
- The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple associated *for-loops* may be listed in a **lastprivate** clause.
- Variables with **const**-qualified type having no mutable member may be listed in a **firstprivate** clause, even if they are static data members.

C / C++

Fortran

- The loop iteration variable(s) in the associated *do-loop(s)* of a **do**, **parallel do**, or **taskloop** construct may be listed in a **private** or **lastprivate** clause.
- The loop iteration variable in the associated *do-loop* of a **simd** construct with just one associated *do-loop* may be listed in a **linear** clause with a *constant-linear-step* that is the increment of the associated loop.
- The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple associated *do-loops* may be listed in a **lastprivate** clause.
- Variables used as loop iteration variables in sequential loops in a **parallel**, **task**, or **taskloop** construct may be listed in data-sharing clauses on the construct itself, and on enclosed constructs, subject to other restrictions.
- Assumed-size arrays may be listed in a **shared** clause.

Fortran

Additional restrictions on the variables that may appear in individual clauses are described with each clause in Section 2.15.3 on page 185.

Variables with *explicitly determined* data-sharing attributes are those that are referenced in a given construct and are listed in a data-sharing attribute clause on the construct.

Variables with *implicitly determined* data-sharing attributes are those that are referenced in a given construct, do not have predetermined data-sharing attributes, and are not listed in a data-sharing attribute clause on the construct.

Rules for variables with *implicitly determined* data-sharing attributes are as follows:

- In a **parallel**, **task**, or **taskloop** construct, the data-sharing attributes of these variables are determined by the **default** clause, if present (see Section 2.15.3.1 on page 186).
- In a **parallel** construct, if no **default** clause is present, these variables are shared.
- For constructs other than **task**, **taskloop**, or **target**, if no **default** clause is present, these variables reference the variables with the same names that exist in the enclosing context.
- In a **task** or **taskloop** construct, if no **default** clause is present, a variable that in the enclosing context is determined to be shared by all implicit tasks bound to the current team is shared.
- In a **target** construct, variables that are not mapped after applying data-mapping attribute rules (see Section 2.15.5 on page 212) are firstprivate.

Fortran

- In an orphaned **task** construct or an orphaned **taskloop** construct, if no **default** clause is present, dummy arguments are firstprivate.

Fortran

- In a **task** construct or a **taskloop** construct, if no **default** clause is present, a variable whose data-sharing attribute is not determined by the rules above is firstprivate.

Additional restrictions on the variables for which data-sharing attributes cannot be implicitly determined in a **task** or **taskloop** construct are described in Section 2.15.3.4 on page 193.

1 2.15.1.2 Data-sharing Attribute Rules for Variables Referenced 2 in a Region but not in a Construct

3 The data-sharing attributes of variables that are referenced in a region, but not in a construct, are
4 determined as follows:

▼ C / C++ ▼

- 5 • Variables with static storage duration that are declared in called routines in the region are shared.
- 6 • File-scope or namespace-scope variables referenced in called routines in the region are shared
7 unless they appear in a **threadprivate** directive.
- 8 • Objects with dynamic storage duration are shared.
- 9 • Static data members are shared unless they appear in a **threadprivate** directive.
- 10 • In C++, formal arguments of called routines in the region that are passed by reference have the
11 same data-sharing attributes as the associated actual arguments.
- 12 • Other variables declared in called routines in the region are private.

▲ C / C++ ▲

▼ Fortran ▼

- 13 • Local variables declared in called routines in the region and that have the **save** attribute, or that
14 are data initialized, are shared unless they appear in a **threadprivate** directive.
- 15 • Variables belonging to common blocks, or accessed by host or use association, and referenced in
16 called routines in the region are shared unless they appear in a **threadprivate** directive.
- 17 • Dummy arguments of called routines in the region that are passed by reference have the same
18 data-sharing attributes as the associated actual arguments.
- 19 • Cray pointees have the same data-sharing attribute as the storage with which their Cray pointers
20 are associated.
- 21 • Implied-do indices, **forall** indices, and other local variables declared in called routines in the
22 region are private.

▲ Fortran ▲

23 2.15.2 threadprivate Directive

24 Summary

25 The **threadprivate** directive specifies that variables are replicated, with each thread having its
26 own copy. The **threadprivate** directive is a declarative directive.

Syntax

C / C++

The syntax of the **threadprivate** directive is as follows:

```
#pragma omp threadprivate(list) new-line
```

where *list* is a comma-separated list of file-scope, namespace-scope, or static block-scope variables that do not have incomplete types.

C / C++

Fortran

The syntax of the **threadprivate** directive is as follows:

```
!$omp threadprivate(list)
```

where *list* is a comma-separated list of named variables and named common blocks. Common block names must appear between slashes.

Fortran

Description

Each copy of a threadprivate variable is initialized once, in the manner specified by the program, but at an unspecified point in the program prior to the first reference to that copy. The storage of all copies of a threadprivate variable is freed according to how static variables are handled in the base language, but at an unspecified point in the program.

A program in which a thread references another thread's copy of a threadprivate variable is non-conforming.

The content of a threadprivate variable can change across a task scheduling point if the executing thread switches to another task that modifies the variable. For more details on task scheduling, see Section 1.3 on page 14 and Section 2.9 on page 81.

In **parallel** regions, references by the master thread will be to the copy of the variable in the thread that encountered the **parallel** region.

During a sequential part references will be to the initial thread's copy of the variable. The values of data in the initial thread's copy of a threadprivate variable are guaranteed to persist between any two consecutive references to the variable in the program.

The values of data in the threadprivate variables of non-initial threads are guaranteed to persist between two consecutive active **parallel** regions only if all the following conditions hold:

- Neither **parallel** region is nested inside another explicit **parallel** region.
- The number of threads used to execute both **parallel** regions is the same.
- The thread affinity policies used to execute both **parallel** regions are the same.
- The value of the *dyn-var* internal control variable in the enclosing task region is *false* at entry to both **parallel** regions.

If these conditions all hold, and if a threadprivate variable is referenced in both regions, then threads with the same thread number in their respective regions will reference the same copy of that variable.

▼ C / C++ ▼

If the above conditions hold, the storage duration, lifetime, and value of a thread's copy of a threadprivate variable that does not appear in any **copyin** clause on the second region will be retained. Otherwise, the storage duration, lifetime, and value of a thread's copy of the variable in the second region is unspecified.

If the value of a variable referenced in an explicit initializer of a threadprivate variable is modified prior to the first reference to any instance of the threadprivate variable, then the behavior is unspecified.

▲ C / C++ ▲

C++

The order in which any constructors for different threadprivate variables of class type are called is unspecified. The order in which any destructors for different threadprivate variables of class type are called is unspecified.

C++

Fortran

A variable is affected by a **copyin** clause if the variable appears in the **copyin** clause or it is in a common block that appears in the **copyin** clause.

If the above conditions hold, the definition, association, or allocation status of a thread's copy of a **threadprivate** variable or a variable in a **threadprivate** common block, that is not affected by any **copyin** clause that appears on the second region, will be retained. Otherwise, the definition and association status of a thread's copy of the variable in the second region is undefined, and the allocation status of an allocatable variable will be implementation defined.

If a **threadprivate** variable or a variable in a **threadprivate** common block is not affected by any **copyin** clause that appears on the first **parallel** region in which it is referenced, the variable or any subobject of the variable is initially defined or undefined according to the following rules:

- If it has the **ALLOCATABLE** attribute, each copy created will have an initial allocation status of not currently allocated.
- If it has the **POINTER** attribute:
 - if it has an initial association status of disassociated, either through explicit initialization or default initialization, each copy created will have an association status of disassociated;
 - otherwise, each copy created will have an association status of undefined.
- If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
 - if it is initially defined, either through explicit initialization or default initialization, each copy created is so defined;
 - otherwise, each copy created is undefined.

Fortran

Restrictions

The restrictions to the **threadprivate** directive are as follows:

- A **threadprivate** variable must not appear in any clause except the **copyin**, **copyprivate**, **schedule**, **num_threads**, **thread_limit**, and **if** clauses.
- A program in which an untied task accesses **threadprivate** storage is non-conforming.

C / C++

- A variable that is part of another variable (as an array or structure element) cannot appear in a **threadprivate** clause unless it is a static data member of a C++ class.
- A **threadprivate** directive for file-scope variables must appear outside any definition or declaration, and must lexically precede all references to any of the variables in its list.
- A **threadprivate** directive for namespace-scope variables must appear outside any definition or declaration other than the namespace definition itself, and must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive at file, namespace, or class scope must refer to a variable declaration at file, namespace, or class scope that lexically precedes the directive.
- A **threadprivate** directive for static block-scope variables must appear in the scope of the variable and not in a nested scope. The directive must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive in block scope must refer to a variable declaration in the same scope that lexically precedes the directive. The variable declaration must use the static storage-class specifier.
- If a variable is specified in a **threadprivate** directive in one translation unit, it must be specified in a **threadprivate** directive in every translation unit in which it is declared.
- The address of a **threadprivate** variable is not an address constant.

C / C++

C++

- A **threadprivate** directive for static class member variables must appear in the class definition, in the same scope in which the member variables are declared, and must lexically precede all references to any of the variables in its list.
- A threadprivate variable must not have an incomplete type or a reference type.
- A threadprivate variable with class type must have:
 - an accessible, unambiguous default constructor in case of default initialization without a given initializer;
 - an accessible, unambiguous constructor accepting the given argument in case of direct initialization;
 - an accessible, unambiguous copy constructor in case of copy initialization with an explicit initializer

C++

Fortran

- A variable that is part of another variable (as an array or structure element) cannot appear in a **threadprivate** clause.
- The **threadprivate** directive must appear in the declaration section of a scoping unit in which the common block or variable is declared. Although variables in common blocks can be accessed by use association or host association, common block names cannot. This means that a common block name specified in a **threadprivate** directive must be declared to be a common block in the same scoping unit in which the **threadprivate** directive appears.
- If a **threadprivate** directive specifying a common block name appears in one program unit, then such a directive must also appear in every other program unit that contains a **COMMON** statement specifying the same name. It must appear after the last such **COMMON** statement in the program unit.
- If a **threadprivate** variable or a **threadprivate** common block is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **threadprivate** directive in the C program.
- A blank common block cannot appear in a **threadprivate** directive.
- A variable can only appear in a **threadprivate** directive in the scope in which it is declared. It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- A variable that appears in a **threadprivate** directive must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

1 **Cross References**

- 2 • *dyn-var* ICV, see Section 2.3 on page 35.
- 3 • number of threads used to execute a **parallel** region, see Section 2.5.1 on page 49.
- 4 • **copyin** clause, see Section 2.15.4.1 on page 208.

5 **2.15.3 Data-Sharing Attribute Clauses**

6 Several constructs accept clauses that allow a user to control the data-sharing attributes of variables
7 referenced in the construct. Data-sharing attribute clauses apply only to variables for which the
8 names are visible in the construct on which the clause appears.

9 Not all of the clauses listed in this section are valid on all directives. The set of clauses that is valid
10 on a particular directive is described with the directive.

11 Most of the clauses accept a comma-separated list of list items (see Section 2.1 on page 26). All list
12 items appearing in a clause must be visible, according to the scoping rules of the base language.
13 With the exception of the **default** clause, clauses may be repeated as needed. A list item that
14 specifies a given variable may not appear in more than one clause on the same directive, except that
15 a variable may be specified in both **firstprivate** and **lastprivate** clauses.

▼ C++ ▼

16 If a variable referenced in a data-sharing attribute clause has a type derived from a template, and
17 there are no other references to that variable in the program, then any behavior related to that
18 variable is unspecified.

▲ C++ ▲

▼ Fortran ▼

19 A named common block may be specified in a list by enclosing the name in slashes. When a named
20 common block appears in a list, it has the same meaning as if every explicit member of the common
21 block appeared in the list. An explicit member of a common block is a variable that is named in a
22 **COMMON** statement that specifies the common block name and is declared in the same scoping unit
23 in which the clause appears.

24 Although variables in common blocks can be accessed by use association or host association,
25 common block names cannot. As a result, a common block name specified in a data-sharing
26 attribute clause must be declared to be a common block in the same scoping unit in which the
27 data-sharing attribute clause appears.

28 When a named common block appears in a **private**, **firstprivate**, **lastprivate**, or
29 **shared** clause of a directive, none of its members may be declared in another data-sharing

attribute clause in that directive. When individual members of a common block appear in a **private**, **firstprivate**, **lastprivate**, **reduction**, or **linear** clause of a directive, the storage of the specified variables is no longer Fortran associated with the storage of the common block itself.

Fortran

2.15.3.1 default Clause

Summary

The **default** clause explicitly determines the data-sharing attributes of variables that are referenced in a **parallel**, **task**, **taskloop** or **teams** construct and would otherwise be implicitly determined (see Section 2.15.1.1 on page 175).

Syntax

C / C++

The syntax of the **default** clause is as follows:

```
default (shared | none)
```

C / C++

Fortran

The syntax of the **default** clause is as follows:

```
default (private | firstprivate | shared | none)
```

Fortran

1

2
3

Fortran

4
56
7

Fortran

8
9
10

11

12

- 13
14

15

16

17
18

19

20

shared (*list*)

Description

All references to a list item within a task refer to the storage area of the original variable at the point the directive was encountered.

It is the programmer's responsibility to ensure, by adding proper synchronization, that storage shared by an explicit **task** region does not reach the end of its lifetime before the explicit **task** region completes its execution.

Fortran

The association status of a shared pointer becomes undefined upon entry to and on exit from the **parallel**, **task**, **taskloop** or **teams** construct if it is associated with a target or a subobject of a target that is in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause inside the construct.

Under certain conditions, 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. It is implementation defined when this situation occurs

Note – Use of intervening temporary storage may occur when the following three conditions hold regarding an actual argument in a reference to a non-intrinsic procedure:

a The actual argument is one of the following:

- A shared variable.
- A subobject of a shared variable.
- An object associated with a shared variable.
- An object associated with a subobject of a shared variable.

b The actual argument is also one of the following:

- An array section.
- An array section with a vector subscript.
- An assumed-shape array.
- A pointer array.

c The associated dummy argument for this actual argument is an explicit-shape array or an assumed-size array.

These conditions effectively result in references to, and definitions of, the temporary storage during the procedure reference. Any references to (or definitions of) the shared storage that is associated with the dummy argument by any other task must be synchronized with the procedure reference to avoid possible race conditions.

Fortran

Restrictions

The restrictions for the **shared** clause are as follows:

- A variable that is part of another variable (as an array or structure element) cannot appear in a shared clause.

2.15.3.3 private Clause

Summary

The **private** clause declares one or more list items to be private to a task or to a SIMD lane.

Syntax

The syntax of the private clause is as follows:

```
private (list)
```

Description

Each task that references a list item that appears in a **private** clause in any statement in the construct receives a new list item. Each SIMD lane used in a **simd** construct that references a list item that appears in a private clause in any statement in the construct receives a new list item. Language-specific attributes for new list items are derived from the corresponding original list item. Inside the construct, all references to the original list item are replaced by references to the new list item. In the rest of the region, it is unspecified whether references are to the new list item or the original list item.

C++

If the construct is contained in a member function, it is unspecified anywhere in the region if accesses through the implicit **this** pointer refer to the new list item or the original list item.

C++

Therefore, if an attempt is made to reference the original item, its value after the region is also unspecified. If a SIMD construct or a task does not reference a list item that appears in a **private** clause, it is unspecified whether SIMD lanes or the task receive a new list item.

The value and/or allocation status of the original list item will change only:

- if accessed and modified via pointer,
- if possibly accessed in the region but outside of the construct,
- as a side effect of directives or clauses, or

Fortran

- if accessed and modified via construct association.

Fortran

List items that appear in a **private**, **firstprivate**, or **reduction** clause in a **parallel** construct may also appear in a **private** clause in an enclosed **parallel**, **task**, or worksharing, or **simd** construct.

List items that appear in a **private** or **firstprivate** clause in a **task** construct may also appear in a **private** clause in an enclosed **parallel** or **task** construct.

List items that appear in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause in a worksharing construct may also appear in a **private** clause in an enclosed **parallel** or **task** construct.

C / C++

1 If the type of a list item is a reference to a type T then the type will be considered to be T for all
2 purposes of this clause.

3 A new list item of the same type, with automatic storage duration, is allocated for the construct.
4 The storage and thus lifetime of these list items lasts until the block in which they are created exits.
5 The size and alignment of the new list item are determined by the type of the variable. This
6 allocation occurs once for each task generated by the construct and/or once for each SIMD lane
7 used by the construct.

8 The new list item is initialized, or has an undefined initial value, as if it had been locally declared
9 without an initializer.

C / C++

C++

10 The order in which any default constructors for different private variables of class type are called is
11 unspecified. The order in which any destructors for different private variables of class type are
12 called is unspecified.

C++

Fortran

13 If any statement of the construct references a list item, a new list item of the same type and type
14 parameters is allocated: once for each implicit task in the **parallel** construct; once for each task
15 generated by a **task** construct; and once for each SIMD lane used by a **simd** construct. The initial
16 value of the new list item is undefined. Within a **parallel**, **worksharing**, **task**, **teams**, or
17 **simd** region, the initial status of a private pointer is undefined.

18 For a list item or the subobject of a list item with the **ALLOCATABLE** attribute:

- 19 • if the allocation status is “not currently allocated”, the new list item or the subobject of the new
20 list item will have an initial allocation status of “not currently allocated”;
- 21 • if the allocation status is “currently allocated”, the new list item or the subobject of the new list
22 item will have an initial allocation status of “currently allocated”. If the new list item or the
23 subobject of the new list item is an array, its bounds will be the same as those of the original list
24 item or the subobject of the original list item.

25 A list item that appears in a **private** clause may be storage-associated with other variables when
26 the **private** clause is encountered. Storage association may exist because of constructs such as
27 **EQUIVALENCE** or **COMMON**. If A is a variable appearing in a **private** clause and B is a variable
28 that is storage-associated with A , then:

- 29 • The contents, allocation, and association status of B are undefined on entry to the **parallel**,
30 **task**, **simd**, or **teams** region.

- Any definition of *A*, or of its allocation or association status, causes the contents, allocation, and association status of *B* to become undefined.
- Any definition of *B*, or of its allocation or association status, causes the contents, allocation, and association status of *A* to become undefined.

A list item that appears in a **private** clause may be a selector of an **ASSOCIATE** construct. If the construct association is established prior to a **parallel** region, the association between the associate name and the original list item will be retained in the region.

Finalization of a list item of a finalizable type or subobjects of a list item of a finalizable type occurs at the end of the region. The order in which any final subroutines for different variables of a finalizable type are called is unspecified.

Fortran

Restrictions

The restrictions to the **private** clause are as follows:

C

- A variable that is part of another variable (as an array or structure element) cannot appear in a **private** clause.

C

Fortran

- A variable that is part of another variable (as an array or structure element) cannot appear in a **private** clause.

Fortran

C++

- A variable that is part of another variable (as an array or structure element) cannot appear in a **private** clause except if the **private** clause is associated with a construct within a class member function and the variable is a member of the same class or the variable is an accessible member of a base class.
- A variable of class type (or array thereof) that appears in a **private** clause requires an accessible, unambiguous default constructor for the class type.

C++

C / C++

- A variable that appears in a **private** clause must not have a **const**-qualified type unless it is of class type with a **mutable** member. This restriction does not apply to the **firstprivate** clause.
- A variable that appears in a **private** clause must not have an incomplete type or be a reference to an incomplete type.
- If a list item is a reference type then it must bind to the same object for all threads of the team.

C / C++

Fortran

- A variable that appears in a **private** clause must either be definable, or an allocatable variable. This restriction does not apply to the **firstprivate** clause.
- Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions, may not appear in a **private** clause.
- Pointers with the **INTENT (IN)** attribute may not appear in a **private** clause. This restriction does not apply to the **firstprivate** clause.

Fortran

2.15.3.4 **firstprivate** Clause

Summary

The **firstprivate** clause declares one or more list items to be private to a task, and initializes each of them with the value that the corresponding original item has when the construct is encountered.

Syntax

The syntax of the **firstprivate** clause is as follows:

```
firstprivate (list)
```

Description

The **firstprivate** clause provides a superset of the functionality provided by the **private** clause.

A list item that appears in a **firstprivate** clause is subject to the **private** clause semantics described in Section 2.15.3.3 on page 189, except as noted. In addition, the new list item is initialized from the original list item existing before the construct. The initialization of the new list item is done once for each task that references the list item in any statement in the construct. The initialization is done prior to the execution of the construct.

For a **firstprivate** clause on a **target**, **parallel**, **task**, or **teams** construct, the initial value of the new list item is the value of the original list item that exists immediately prior to the construct in the task region where the construct is encountered. For a **firstprivate** clause on a worksharing construct, the initial value of the new list item for each implicit task of the threads that execute the worksharing construct is the value of the original list item that exists in the implicit task immediately prior to the point in time that the worksharing construct is encountered.

To avoid race conditions, concurrent updates of the original list item must be synchronized with the read of the original list item that occurs as a result of the **firstprivate** clause.

If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required for **lastprivate** occurs after all the initializations for **firstprivate**.

▼ C / C++ ▼

For variables of non-array type, the initialization occurs by copy assignment. For an array of elements of non-array type, each element is initialized as if by assignment from an element of the original array to the corresponding element of the new array.

▲ C / C++ ▲

▼ C++ ▼

For variables of class type, a copy constructor is invoked to perform the initialization. The order in which copy constructors for different variables of class type are called is unspecified.

▲ C++ ▲

▼ Fortran ▼

If the original list item does not have the **POINTER** attribute, initialization of the new list items occurs as if by intrinsic assignment, unless the original list item has the allocation status of not currently allocated, in which case the new list items will have the same status.

If the original list item has the **POINTER** attribute, the new list items receive the same association status of the original list item as if by pointer assignment.

▲ Fortran ▲

1 **Restrictions**

2 The restrictions to the **firstprivate** clause are as follows:

- 3 • A list item that is private within a **parallel** region must not appear in a **firstprivate**
4 clause on a worksharing construct if any of the worksharing regions arising from the worksharing
5 construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
- 6 • A list item that is private within a **teams** region must not appear in a **firstprivate** clause
7 on a **distribute** construct if any of the **distribute** regions arising from the
8 **distribute** construct ever bind to any of the **teams** regions arising from the **teams**
9 construct.
- 10 • A list item that appears in a **reduction** clause of a **parallel** construct must not appear in a
11 **firstprivate** clause on a worksharing or **task** construct if any of the worksharing or **task**
12 regions arising from the worksharing or **task** construct ever bind to any of the **parallel**
13 regions arising from the **parallel** construct.
- 14 • A list item that appears in a **reduction** clause of a **teams** construct must not appear in a
15 **firstprivate** clause on a **distribute** construct if any of the **distribute** regions
16 arising from the **distribute** construct ever bind to any of the **teams** regions arising from the
17 **teams** construct.
- 18 • A list item that appears in a **reduction** clause in a worksharing construct must not appear in a
19 **firstprivate** clause in a task construct encountered during execution of any of the
20 worksharing regions arising from the worksharing construct.

▼ C++ ▼

- 21 • A variable of class type (or array thereof) that appears in a **firstprivate** clause requires an
22 accessible, unambiguous copy constructor for the class type.

▲ C++ ▲

▼ C / C++ ▼

- 23 • A variable that appears in a **firstprivate** clause must not have an incomplete C/C++ type or
24 be a reference to an incomplete type.
- 25 • If a list item is a reference type then it must bind to the same object for all threads of the team.

▲ C / C++ ▲

▼ Fortran ▼

- 26 • Variables that appear in namelist statements, in variable format expressions, and in expressions
27 for statement function definitions, may not appear in a **firstprivate** clause.

▲ Fortran ▲

1 **2.15.3.5 lastprivate Clause**

2 **Summary**

3 The **lastprivate** clause declares one or more list items to be private to an implicit task or to a
4 SIMD lane, and causes the corresponding original list item to be updated after the end of the region.

5 **Syntax**

6 The syntax of the **lastprivate** clause is as follows:

`lastprivate (list)`

7 **Description**

8 The **lastprivate** clause provides a superset of the functionality provided by the **private**
9 clause.

10 A list item that appears in a **lastprivate** clause is subject to the **private** clause semantics
11 described in Section 2.15.3.3 on page 189. In addition, when a **lastprivate** clause appears on
12 the directive that identifies a worksharing construct or a SIMD construct, the value of each new list
13 item from the sequentially last iteration of the associated loops, or the lexically last **section**
14 construct, is assigned to the original list item.



15 For an array of elements of non-array type, each element is assigned to the corresponding element
16 of the original array.



17 If the original list item does not have the **POINTER** attribute, its update occurs as if by intrinsic
18 assignment.

19 If the original list item has the **POINTER** attribute, its update occurs as if by pointer assignment.

List items that are not assigned a value by the sequentially last iteration of the loops, or by the lexically last **section** construct, have unspecified values after the construct. Unassigned subcomponents also have unspecified values after the construct.

The original list item becomes defined at the end of the construct if there is an implicit barrier at that point. To avoid race conditions, concurrent reads or updates of the original list item must be synchronized with the update of the original list item that occurs as a result of the **lastprivate** clause.

If the **lastprivate** clause is used on a construct to which **nowait** is applied, accesses to the original list item may create a data race. To avoid this, synchronization must be inserted to ensure that the sequentially last iteration or lexically last section construct has stored and flushed that list item.

If the **lastprivate** clause is used on a **distribute simd**, distribute parallel loop, or distribute parallel loop SIMD, accesses to the original list item may create a data race. To avoid this, synchronization must be inserted to ensure that the sequentially last iteration has stored and flushed that list item.

If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required for **lastprivate** occurs after all initializations for **firstprivate**.

Restrictions

The restrictions to the **lastprivate** clause are as follows:

- A list item that is private within a **parallel** region, or that appears in the **reduction** clause of a **parallel** construct, must not appear in a **lastprivate** clause on a worksharing construct if any of the corresponding worksharing regions ever binds to any of the corresponding **parallel** regions.

C++

- A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an accessible, unambiguous default constructor for the class type, unless the list item is also specified in a **firstprivate** clause.
- A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an accessible, unambiguous copy assignment operator for the class type. The order in which copy assignment operators for different variables of class type are called is unspecified.

C++

C / C++

- A variable that appears in a **lastprivate** clause must not have a **const**-qualified type unless it is of class type with a **mutable** member.
- A variable that appears in a **lastprivate** clause must not have an incomplete C/C++ type or be a reference to an incomplete type.
- If a list item is a reference type then it must bind to the same object for all threads of the team.

C / C++

Fortran

- A variable that appears in a **lastprivate** clause must be definable.
- If the original list item has the **ALLOCATABLE** attribute, the corresponding list item in the sequentially last iteration or lexically last section must have an allocation status of allocated upon exit from that iteration or section.
- Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions, may not appear in a **lastprivate** clause.

Fortran

2.15.3.6 reduction Clause

Summary

The **reduction** clause specifies a *reduction-identifier* and one or more list items. For each list item, a private copy is created in each implicit task or SIMD lane, and is initialized with the initializer value of the *reduction-identifier*. After the end of the region, the original list item is updated with the values of the private copies using the combiner associated with the *reduction-identifier*.

Syntax

C / C++

The syntax of the **reduction** clause is as follows:

```
reduction (reduction-identifier : list)
```

where:

C

reduction-identifier is either an *identifier* or one of the following operators: +, -, *, &, |, ^, && and ||

C

C++

reduction-identifier is either an *id-expression* or one of the following operators: +, -, *, &, |, ^, && and ||

C++

Table 2.7 lists each *reduction-identifier* that is implicitly declared at every scope for arithmetic types and its semantic initializer value. The actual initializer value is that value as expressed in the data type of the reduction list item.

Identifier	Initializer	Combiner
+	omp_priv = 0	omp_out += omp_in
*	omp_priv = 1	omp_out *= omp_in
-	omp_priv = 0	omp_out += omp_in
&	omp_priv = 0	omp_out &= omp_in
	omp_priv = 0	omp_out = omp_in
^	omp_priv = 0	omp_out ^= omp_in
&&	omp_priv = 1	omp_out = omp_in && omp_out
	omp_priv = 0	omp_out = omp_in omp_out
max	omp_priv = <i>Least representable number in the reduction list item type</i>	omp_out = omp_in > omp_out ? omp_in : omp_out


table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
min	omp_priv = <i>Largest representable number in the reduction list item type</i>	omp_out = omp_in < omp_out ? omp_in : omp_out

Table 2.7.: Implicitly Declared C/C++ *reduction-identifiers*

where **omp_in** and **omp_out** correspond to two identifiers that refer to storage of the type of the list item. **omp_out** holds the final value of the combiner operation.



The syntax of the **reduction** clause is as follows:

reduction(*reduction-identifier* : *list*)

where *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the following operators: **+**, **-**, *****, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of the following intrinsic procedure names: **max**, **min**, **iand**, **ior**, **ieor**.

Table 2.8 lists each *reduction-identifier* that is implicitly declared for numeric and logical types and its semantic initializer value. The actual initializer value is that value as expressed in the data type of the reduction list item.

Identifier	Initializer	Combiner
+	omp_priv = 0	omp_out = omp_in + omp_out
*	omp_priv = 1	omp_out = omp_in * omp_out
-	omp_priv = 0	omp_out = omp_in + omp_out
.and.	omp_priv = .true.	omp_out = omp_in .and. omp_out
.or.	omp_priv = .false.	omp_out = omp_in .or. omp_out
.eqv.	omp_priv = .true.	omp_out = omp_in .eqv. omp_out
.neqv.	omp_priv = .false.	omp_out = omp_in .neqv. omp_out

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
max	omp_priv = <i>Least representable number in the reduction list item type</i>	omp_out = max (omp_in , omp_out)
min	omp_priv = <i>Largest representable number in the reduction list item type</i>	omp_out = min (omp_in , omp_out)
iand	omp_priv = <i>All bits on</i>	omp_out = iand (omp_in , omp_out)
ior	omp_priv = 0	omp_out = ior (omp_in , omp_out)
ieor	omp_priv = 0	omp_out = ieor (omp_in , omp_out)

Table 2.8.: Implicitly Declared Fortran *reduction-identifiers*

Fortran

Any *reduction-identifier* that is defined with the **declare reduction** directive is also valid. In that case, the initializer and combiner of the *reduction-identifier* are specified by the *initializer-clause* and the combiner in the **declare reduction** directive.

Description

The reduction clause can be used to perform some forms of recurrence calculations (involving mathematically associative and commutative operators) in parallel.

For **parallel** and worksharing constructs, a private copy of each list item is created, one for each implicit task, as if the **private** clause had been used. For the **simd** construct, a private copy of each list item is created, one for each SIMD lane as if the **private** clause had been used. For the **teams** construct, a private copy of each list item is created, one for each team in the league as if the **private** clause had been used. The private copy is then initialized as specified above. At the end of the region for which the **reduction** clause was specified, the original list item is updated by combining its original value with the final value of each of the private copies, using the combiner of the specified *reduction-identifier*.

Fortran

If the original list item has the **POINTER** attribute, the private copy of the list item is associated with a private target.

Fortran

The *reduction-identifier* specified in the **reduction** clause must match a previously declared *reduction-identifier* of the same name and type for each of the list items. This match is done by means of a name lookup in the base language.

C / C++

The list items that appear in the **reduction** clause may include array sections.

If the list item is an array or an array section it will be treated as if a **reduction** clause would be applied to each separate element of the array section. The elements of each private array section will be allocated contiguously.

C / C++

C++

If the type is a derived class, then any *reduction-identifier* that matches its base classes are also a match, if there is no specific match for the type.

If the *reduction-identifier* is not an *id-expression* then it is implicitly converted to one by prepending the keyword operator (for example, **+** becomes *operator+*).

If the *reduction-identifier* is qualified then a qualified name lookup is used to find the declaration.

If the *reduction-identifier* is unqualified then an *argument-dependent name lookup* must be performed using the type of each list item.

C++

If **nowait** is not used, the reduction computation will be complete at the end of the construct; however, if the reduction clause is used on a construct to which **nowait** is also applied, accesses to the original list item will create a race and, thus, have unspecified effect unless synchronization ensures that they occur after all threads have executed all of their iterations or **section** constructs, and the reduction computation has completed and stored the computed value of that list item. This can most simply be ensured through a barrier synchronization.

The location in the OpenMP program at which the values are combined and the order in which the values are combined are unspecified. Therefore, when comparing sequential and parallel runs, or when comparing one parallel run to another (even if the number of threads used is the same), there is no guarantee that bit-identical results will be obtained or that side effects (such as floating-point exceptions) will be identical or take place at the same location in the OpenMP program.

To avoid race conditions, concurrent reads or updates of the original list item must be synchronized with the update of the original list item that occurs as a result of the **reduction** computation.

Restrictions

The restrictions to the **reduction** clause are as follows:

- A list item that appears in a **reduction** clause of a worksharing construct must be shared in the **parallel** regions to which any of the worksharing regions arising from the worksharing construct bind.
- A list item that appears in a **reduction** clause of the innermost enclosing worksharing or **parallel** construct may not be accessed in an explicit task.
- Any number of **reduction** clauses can be specified on the directive, but a list item can appear only once in the **reduction** clauses for that directive.
- For a *reduction-identifier* declared with the **declare reduction** construct, the directive must appear before its use in a **reduction** clause.
- If a list item is an array section, it must specify contiguous storage and it cannot be a zero-length array section.
- If a list item is an array section, its lower-bound must be zero.
- If a list item is an array section, accesses to the elements of the array outside the specified array section result in unspecified behavior.

▼ C / C++ ▼

- The type of a list item that appears in a **reduction** clause must be valid for the *reduction-identifier*. For a **max** or **min** reduction in C, the type of the list item must be an allowed arithmetic data type: **char**, **int**, **float**, **double**, or **_Bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**. For a **max** or **min** reduction in C++, the type of the list item must be an allowed arithmetic data type: **char**, **wchar_t**, **int**, **float**, **double**, or **bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.
- A list item that appears in a **reduction** clause must not be **const**-qualified.
- If a list item is a reference type then it must bind to the same object for all threads of the team.
- The *reduction-identifier* for any list item must be unambiguous and accessible.

▲ C / C++ ▲

- 1 • The type and the rank of a list item that appears in a **reduction** clause must be valid for the
- 2 *combiner* and *initializer*.
- 3 • A list item that appears in a **reduction** clause must be definable.
- 4 • A procedure pointer may not appear in a **reduction** clause.
- 5 • A pointer with the **INTENT (IN)** attribute may not appear in the **reduction** clause.
- 6 • An original list item with the **POINTER** attribute or any pointer component of an original list
- 7 item that is referenced in *combiner* must be associated at entry to the construct that contains the
- 8 **reduction** clause. Additionally, the list item or the pointer component of the list item must
- 9 not be deallocated, allocated, or pointer assigned within the region.
- 10 • An original list item with the **ALLOCATABLE** attribute or any allocatable component of an
- 11 original list item that is referenced in *combiner* must be in the allocated state at entry to the
- 12 construct that contains the **reduction** clause. Additionally, the list item or the allocatable
- 13 component of the list item must be neither deallocated nor allocated within the region.
- 14 • If the *reduction-identifier* is defined in a **declare reduction** directive, the
- 15 **declare reduction** directive must be in the same subprogram, or accessible by host or use
- 16 association.
- 17 • If the *reduction-identifier* is a user-defined operator, the same explicit interface for that operator
- 18 must be accessible as at the **declare reduction** directive.
- 19 • If the *reduction-identifier* is defined in a **declare reduction** directive, any subroutine or
- 20 function referenced in the initializer clause or combiner expression must be an intrinsic function,
- 21 or must have an explicit interface where the same explicit interface is accessible as at the
- 22 **declare reduction** directive.

23 2.15.3.7 linear Clause

24 Summary

25 The **linear** clause declares one or more list items to be private to a SIMD lane and to have a
 26 linear relationship with respect to the iteration space of a loop.

1 **Syntax**

C

2 The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

3 where *linear-list* is one of the following

- 4 *list*
- 5 *modifier* (*list*)

6 where *modifier* is one of the following:

7 **val**

C

C++

8 The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

9 where *linear-list* is one of the following

- 10 *list*
- 11 *modifier* (*list*)

12 where *modifier* is one of the following:

- 13 **ref**
- 14 **val**
- 15 **uval**

C++

The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

where *linear-list* is one of the following

list

modifier (*list*)

where *modifier* is one of the following:

ref

val

uval

Description

The **linear** clause provides a superset of the functionality provided by the **private** clause.

A list item that appears in a **linear** clause is subject to the **private** clause semantics described in Section 2.15.3.3 on page 189 except as noted. In addition, if the **val** or **uval** modifier is used, the value of the new list item on each iteration of the associated loop(s) corresponds to the value of the original list item before entering the construct plus the logical number of the iteration times *linear-step*. If the **ref** modifier is used, the value of the new list item on each iteration of the associated loop(s) corresponds to the value of the variable resulting from applying the *linear-step* times the logical number of the iteration as a subscript to the original list item. If no modifier is specified it is assumed to be **val**. If *linear-step* is not specified it is assumed to be 1. The value corresponding to the sequentially last iteration of the associated loops is assigned to the original list item.

The **uval** modifier declares the address of the original list item to have an invariant value for all iterations of the associated loop(s).

Restrictions

- The *linear-step* expression must be invariant during the execution of the region associated with the construct. Otherwise, the execution results in unspecified behavior.
- A *list-item* cannot appear in more than one **linear** clause.
- A *list-item* that appears in a **linear** clause cannot appear in any other data-sharing attribute clause.

C

- A *list-item* that appears in a **linear** clause must be of integral or pointer type.

C

C++

- A *list-item* that appears in a **linear** clause without the **ref** modifier must be of integral or pointer type, or must be a reference to an integral or pointer type.
- The **ref** or **uval** modifier can only be used if the *list-item* is of a reference type.

C++

Fortran

- A *list-item* that appears in a **linear** clause without the **ref** modifier must be of type **integer**.
- The **ref** or **uval** modifier can only be used if the *list-item* is a dummy argument without the **VALUE** attribute.
- Variables that have the **POINTER** attribute and Cray pointers may not appear in a linear clause.
- The list item with the **ALLOCATABLE** attribute in the sequentially last iteration must have an allocation status of allocated upon exit from that iteration.

Fortran

1 2.15.4 Data Copying Clauses

2 This section describes the **copyin** clause (allowed on the **parallel** directive and combined
3 parallel worksharing directives) and the **copyprivate** clause (allowed on the **single** directive).

4 These clauses support the copying of data values from private or threadprivate variables on one
5 implicit task or thread to the corresponding variables on other implicit tasks or threads in the team.

6 The clauses accept a comma-separated list of list items (see Section 2.1 on page 26). All list items
7 appearing in a clause must be visible, according to the scoping rules of the base language. Clauses
8 may be repeated as needed, but a list item that specifies a given variable may not appear in more
9 than one clause on the same directive

▼ Fortran ▼

10 An associate name preserves the association with the selector established at the **ASSOCIATE**
11 statement. A list item that appears in a data copying clause may be a selector of an **ASSOCIATE**
12 construct. If the construct association is established prior to a parallel region, the association
13 between the associate name and the original list item will be retained in the region.

▲ Fortran ▲

14 2.15.4.1 copyin Clause

15 Summary

16 The **copyin** clause provides a mechanism to copy the value of the master thread's threadprivate
17 variable to the threadprivate variable of each other member of the team executing the **parallel**
18 region.

19 Syntax

20 The syntax of the **copyin** clause is as follows:

copyin (<i>list</i>)

Description

C / C++

The copy is done after the team is formed and prior to the start of execution of the associated structured block. For variables of non-array type, the copy occurs by copy assignment. For an array of elements of non-array type, each element is copied as if by assignment from an element of the master thread's array to the corresponding element of the other thread's array.

C / C++

C++

For class types, the copy assignment operator is invoked. The order in which copy assignment operators for different variables of class type are called is unspecified.

C++

Fortran

The copy is done, as if by assignment, after the team is formed and prior to the start of execution of the associated structured block.

On entry to any **parallel** region, each thread's copy of a variable that is affected by a **copyin** clause for the **parallel** region will acquire the allocation, association, and definition status of the master thread's copy, according to the following rules:

- If the original list item has the **POINTER** attribute, each copy receives the same association status of the master thread's copy as if by pointer assignment.
- If the original list item does not have the **POINTER** attribute, each copy becomes defined with the value of the master thread's copy as if by intrinsic assignment, unless it has the allocation status of not currently allocated, in which case each copy will have the same status.

Fortran

Restrictions

The restrictions to the **copyin** clause are as follows:

C / C++

- A list item that appears in a **copyin** clause must be threadprivate.
- A variable of class type (or array thereof) that appears in a **copyin** clause requires an accessible, unambiguous copy assignment operator for the class type.

C / C++

Fortran

- A list item that appears in a **copyin** clause must be threadprivate. Named variables appearing in a threadprivate common block may be specified: it is not necessary to specify the whole common block.
- A common block name that appears in a **copyin** clause must be declared to be a common block in the same scoping unit in which the **copyin** clause appears.

Fortran

2.15.4.2 copyprivate Clause

Summary

The **copyprivate** clause provides a mechanism to use a private variable to broadcast a value from the data environment of one implicit task to the data environments of the other implicit tasks belonging to the **parallel** region.

To avoid race conditions, concurrent reads or updates of the list item must be synchronized with the update of the list item that occurs as a result of the **copyprivate** clause.

Syntax

The syntax of the **copyprivate** clause is as follows:

```
copyprivate (list)
```

Description

The effect of the **copyprivate** clause on the specified list items occurs after the execution of the structured block associated with the **single** construct (see Section 2.7.3 on page 66), and before any of the threads in the team have left the barrier at the end of the construct.

C / C++

In all other implicit tasks belonging to the **parallel** region, each specified list item becomes defined with the value of the corresponding list item in the implicit task whose thread executed the structured block. For variables of non-array type, the definition occurs by copy assignment. For an array of elements of non-array type, each element is copied by copy assignment from an element of the array in the data environment of the implicit task associated with the thread that executed the structured block to the corresponding element of the array in the data environment of the other implicit tasks

C / C++

C++

For class types, a copy assignment operator is invoked. The order in which copy assignment operators for different variables of class type are called is unspecified.

C++

Fortran

If a list item does not have the **POINTER** attribute, then in all other implicit tasks belonging to the **parallel** region, the list item becomes defined as if by intrinsic assignment with the value of the corresponding list item in the implicit task associated with the thread that executed the structured block.

If the list item has the **POINTER** attribute, then, in all other implicit tasks belonging to the **parallel** region, the list item receives, as if by pointer assignment, the same association status of the corresponding list item in the implicit task associated with the thread that executed the structured block.

The order in which any final subroutines for different variables of a finalizable type are called is unspecified.

Fortran

Note – The **copyprivate** clause is an alternative to using a shared variable for the value when providing such a shared variable would be difficult (for example, in a recursion requiring a different variable at each level).

Restrictions

The restrictions to the **copyprivate** clause are as follows:

- All list items that appear in the **copyprivate** clause must be either **threadprivate** or **private** in the enclosing context.
- A list item that appears in a **copyprivate** clause may not appear in a **private** or **firstprivate** clause on the **single** construct.

C++

- A variable of class type (or array thereof) that appears in a **copyprivate** clause requires an accessible unambiguous copy assignment operator for the class type.

C++

Fortran

- A common block that appears in a **copyprivate** clause must be **threadprivate**.
- Pointers with the **INTENT (IN)** attribute may not appear in the **copyprivate** clause.
- The list item with the **ALLOCATABLE** attribute must have the allocation status of **allocated** when the intrinsic assignment is performed.

Fortran

2.15.5 Data-mapping Attribute Rules and Clauses

This section describes how the data mapping attributes of any variable referenced in a **target** region are determined. When specified, explicit **map** clauses on **target data** and **target** directives determine these attributes. Otherwise, the following implicit data-mapping rules apply for variables referenced in a **target** construct that are not declared in the construct and do not appear in data-sharing attribute or **map** clauses:

C / C++

- A variable with a type of pointer, reference to array, or reference to pointer is treated as if it had appeared in a **map** clause as the variable in a zero-length array section.

- If a **defaultmap(tofrom: scalar)** clause is not present then a scalar variable is not mapped, but instead has an implicit data-sharing attribute of firstprivate (see Section 2.15.1.1 on page 175).
- If a **defaultmap(tofrom: scalar)** clause is present then a scalar variable is treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.
- If a variable is not a scalar then it is treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.

2.15.5.1 map Clause

Summary

The **map** clause specifies how an original list item is mapped from the current task's data environment to a corresponding list item in the device data environment of the device identified by the construct.

Syntax

The syntax of the map clause is as follows:

```
map ([ [map-type-modifier[,]] map-type : ] list)
```

Description

The list items that appear in a **map** clause may include array sections and structure elements.

The *map-type* and *map-type-modifier* specify the effect of the **map** clause, as described below.

The original and corresponding list items may share storage such that writes to either item by one task followed by a read or write of the other item by another task without intervening synchronization can result in data races.

If the **map** clause appears on a **target**, **target data**, or **target enter data** construct then on entry to the region the following sequence of steps occurs:

1. If a corresponding list item of the original list item is not present in the device data environment, then:
 - a) A new list item with language-specific attributes is derived from the original list item and created in the device data environment.
 - b) The new list item becomes the corresponding list item to the original list item in the device data environment.
 - c) The corresponding list item has a reference count that is initialized to zero.
2. The corresponding list item's reference count is incremented by one.
3. If the corresponding list item's reference count is one or the **always map-type-modifier** is present, then:
 - a) If the *map-type* is **to** or **tofrom**, then the corresponding list item is assigned the value of the original list item.
4. If the corresponding list item's reference count is one, then:
 - a) If the *map-type* is **from** or **alloc** the value of the corresponding list item is undefined.

If the **map** clause appears on a **target**, **target data**, or **target exit data** construct then on exit from the region the following sequence of steps occurs:

1. If a corresponding list item of the original list item is not present in the device data environment, then the list item is ignored.
2. If a corresponding list item of the original list item is present in the device data environment, then:
 - a) If the corresponding list item's reference count is greater than zero, then:
 - i. the corresponding list item's reference count is decremented by one.
 - ii. If the *map-type* is **delete**, then the corresponding list item's reference count is set to zero.

- b) If the corresponding list item's reference count is zero or the **always** *map-type-modifier* is present, then:
 - i. If the *map-type* is **from** or **tofrom**, then the original list item is assigned the value of the corresponding list item.
- c) If the corresponding list item's reference count is zero, then the corresponding list item is removed from the device data environment

C / C++

If a new list item is created then a new list item of the same type, with automatic storage duration, is allocated for the construct. The size and alignment of the new list item are determined by the type of the variable. This allocation occurs if the region references the list item in any statement.

If a new list item is created for an array section and the type of the variable appearing in that array section is pointer, reference to array, or reference to pointer then if the variable appears in a **to** or **from** clause in a **target update** region during the lifetime of the new list item but not as part of the specification of an array section, the behavior is unspecified.

C / C++

Fortran

If a new list item is created then a new list item of the same type, type parameter, and rank is allocated.

Fortran

The *map-type* determines how the new list item is initialized.

If a *map-type* is not specified, the *map-type* defaults to **tofrom**.

Restrictions

- If a list item is an array section, it must specify contiguous storage.
- At most one list item can be an array item derived from a given variable in **map** clauses of the same construct.
- List items of **map** clauses in the same construct must not share original storage.
- If any part of the original storage of a list item has corresponding storage in the device data environment, all of the original storage must have corresponding storage in the device data environment.
- If a list item is an element of a structure, and a different element of the structure has a corresponding list item in the device data environment prior to a task encountering the construct associated with the **map** clause, then the list item must also have a corresponding list item in the device data environment prior to the task encountering the construct.
- If a list item is an element of a structure, only the rightmost symbol of the variable reference can be an array section.
- If variables that share storage are mapped, the behavior is unspecified.
- A list item must have a mappable type.
- **threadprivate** variables cannot appear in a **map** clause.

C / C++

- Initialization and assignment are through bitwise copy.
- A variable for which the type is pointer, reference to array, or reference to pointer and an array section derived from that variable must not appear as list items of **map** clauses of the same construct.
- A variable for which the type is pointer, reference to array, or reference to pointer must not appear as a list item if the device data environment already contains an array section derived from that variable.
- An array section derived from a variable for which the type is pointer, reference to array, or reference to pointer must not appear as a list item if the device data environment already contains that variable.
- A list item cannot be a variable that is a member of a structure with a union type.

C / C++

Fortran

- The value of the new list item becomes that of the original list item in the map initialization and assignment.

Fortran

1 **2.15.5.2 defaultmap Clause**

2 **Summary**

3 The **defaultmap** clause explicitly determines the data-mapping attributes of variables that are
4 referenced in a **target** construct and would otherwise be implicitly determined.

5 **Syntax**

▼ C / C++ ▼

6 The syntax of the **defaultmap** clause is as follows:

`defaultmap (tofrom: scalar)`

▲ C / C++ ▲

▼ Fortran ▼

7 The syntax of the **defaultmap** clause is as follows:

`defaultmap (tofrom: scalar)`

▲ Fortran ▲

8 **Description**

9 The **defaultmap (tofrom: scalar)** clause causes all scalar variables referenced in the
10 construct that have implicitly determined data-mapping attributes to have the **tofrom** *map-type*.

11 **2.16 declare reduction Directive**

12 **Summary**

13 The following section describes the directive for declaring user-defined reductions. The
14 **declare reduction** directive declares a *reduction-identifier* that can be used in a
15 **reduction** clause. The **declare reduction** directive is a declarative directive.

Syntax

C

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner ) [initializer-clause] new-line
```

where:

- *reduction-identifier* is either a base language identifier or one of the following operators: **+**, **-**, *****, **&**, **|**, **^**, **&&** and **||**
- *typename-list* is list of type names
- *combiner* is an expression
- *initializer-clause* is **initializer**(*initializer-expr*) where *initializer-expr* is **omp_priv** = *initializer* or *function-name* (*argument-list*)

C

C++

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner) [initializer-clause] new-line
```

where:

- *reduction-identifier* is either a base language identifier or one of the following operators: **+**, **-**, *****, **&**, **|**, **^**, **&&** and **||**
- *typename-list* is list of type names
- *combiner* is an expression
- *initializer-clause* is **initializer**(*initializer-expr*) where *initializer-expr* is **omp_priv** *initializer* or *function-name* (*argument-list*)

C++

```
!$omp declare reduction(reduction-identifier : type-list : combiner)  
[initializer-clause]
```

where:

- *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the following operators: **+**, **-**, *****, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of the following intrinsic procedure names: **max**, **min**, **iand**, **ior**, **ieor**.
- *type-list* is a list of type specifiers
- *combiner* is either an assignment statement or a subroutine name followed by an argument list
- *initializer-clause* is **initializer**(*initializer-expr*), where *initializer-expr* is **omp_priv** = *expression* or *subroutine-name* (*argument-list*)

Description

Custom reductions can be defined using the **declare reduction** directive; the *reduction-identifier* and the type identify the **declare reduction** directive. The *reduction-identifier* can later be used in a **reduction** clause using variables of the type or types specified in the **declare reduction** directive. If the directive applies to several types then it is considered as if there were multiple **declare reduction** directives, one for each type.

If a type with deferred or assumed length type parameter is specified in a **declare reduction** directive, the *reduction-identifier* of that directive can be used in a **reduction** clause with any variable of the same type and the same kind parameter, regardless of the length type Fortran parameters with which the variable is declared.

The visibility and accessibility of this declaration are the same as those of a variable declared at the same point in the program. The enclosing context of the *combiner* and of the *initializer-expr* will be that of the **declare reduction** directive. The *combiner* and the *initializer-expr* must be correct in the base language as if they were the body of a function defined at the same point in the program.

Fortran

If the *reduction-identifier* is the same as the name of a user-defined operator or an extended operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the operator or procedure name appears in an accessibility statement in the same module, the accessibility of the corresponding **declare reduction** directive is determined by the accessibility attribute of the statement.

If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic procedures and is accessible, and if it has the same name as a derived type in the same module, the accessibility of the corresponding **declare reduction** directive is determined by the accessibility of the generic name according to the base language.

Fortran

C++

The **declare reduction** directive can also appear at points in the program at which a static data member could be declared. In this case, the visibility and accessibility of the declaration are the same as those of a static data member declared at the same point in the program.

C++

The *combiner* specifies how partial results can be combined into a single value. The *combiner* can use the special variable identifiers **omp_in** and **omp_out** that are of the type of the variables being reduced with this *reduction-identifier*. Each of them will denote one of the values to be combined before executing the *combiner*. It is assumed that the special **omp_out** identifier will refer to the storage that holds the resulting combined value after executing the *combiner*.

The number of times the *combiner* is executed, and the order of these executions, for any **reduction** clause is unspecified.

Fortran

If the *combiner* is a subroutine name with an argument list, the *combiner* is evaluated by calling the subroutine with the specified argument list.

If the *combiner* is an assignment statement, the *combiner* is evaluated by executing the assignment statement.

Fortran

As the *initializer-expr* value of a user-defined reduction is not known *a priori* the *initializer-clause* can be used to specify one. Then the contents of the *initializer-clause* will be used as the initializer for private copies of reduction list items where the **omp_priv** identifier will refer to the storage to be initialized. The special identifier **omp_orig** can also appear in the *initializer-clause* and it will refer to the storage of the original variable to be reduced.

The number of times that the *initializer-expr* is evaluated, and the order of these evaluations, is unspecified.

C / C++

If the *initializer-expr* is a function name with an argument list, the *initializer-expr* is evaluated by calling the function with the specified argument list. Otherwise, the *initializer-expr* specifies how **omp_priv** is declared and initialized.

C / C++

C

If no *initializer-clause* is specified, the private variables will be initialized following the rules for initialization of objects with static storage duration.

C

C++

If no *initializer-expr* is specified, the private variables will be initialized following the rules for *default-initialization*.

C++

Fortran

If the *initializer-expr* is a subroutine name with an argument list, the *initializer-expr* is evaluated by calling the subroutine with the specified argument list.

If the *initializer-expr* is an assignment statement, the *initializer-expr* is evaluated by executing the assignment statement.

If no *initializer-clause* is specified, the private variables will be initialized as follows:

- For **complex**, **real**, or **integer** types, the value 0 will be used.
- For **logical** types, the value **.false.** will be used.
- For derived types for which default initialization is specified, default initialization will be used.
- Otherwise, not specifying an *initializer-clause* results in unspecified behavior.

Fortran

C / C++

If *reduction-identifier* is used in a **target** region then a **declare target** construct must be specified for any function that can be accessed through *combiner* and *initializer-expr*.

C / C++

Fortran

If *reduction-identifier* is used in a **target** region then a **declare target** construct must be specified for any function or subroutine that can be accessed through *combiner* and *initializer-expr*.

Fortran

Restrictions

- Only the variables **omp_in** and **omp_out** are allowed in the *combiner*.
- Only the variables **omp_priv** and **omp_orig** are allowed in the *initializer-clause*.
- If the variable **omp_orig** is modified in the *initializer-clause*, the behavior is unspecified.
- If execution of the *combiner* or the *initializer-expr* results in the execution of an OpenMP construct or an OpenMP API call, then the behavior is unspecified.
- A *reduction-identifier* may not be re-declared in the current scope for the same type or for a type that is compatible according to the base language rules.
- At most one *initializer-clause* can be specified.

C / C++

- A type name in a **declare reduction** directive cannot be a function type, an array type, a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

C

- If the *initializer-expr* is a function name with an argument list, then one of the arguments must be the address of **omp_priv**.

C

C++

- If the *initializer-expr* is a function name with an argument list, then one of the arguments must be **omp_priv** or the address of **omp_priv**.

C++

- If the *initializer-expr* is a subroutine name with an argument list, then one of the arguments must be `omp_priv`.
- If the **declare reduction** directive appears in the specification part of a module and the corresponding reduction clause does not appear in the same module, the *reduction-identifier* must be the same as the name of a user-defined operator, one of the allowed operators that is extended or a generic name that is the same as the name of one of the allowed intrinsic procedures.
- If the **declare reduction** directive appears in the specification of a module, if the corresponding **reduction** clause does not appear in the same module, and if the *reduction-identifier* is the same as the name of a user-defined operator or an extended operator, or the same as a generic name that is the same as one of the allowed intrinsic procedures, the interface for that operator or the generic name must be defined in the specification of the same module, or must be accessible by use association.
- Any subroutine, or function used in the **initializer** clause or *combiner* expression must be an intrinsic function, or must have an accessible interface.
- Any user-defined operator, or extended operator used in the **initializer** clause or *combiner* expression must have an accessible interface.
- If any subroutine, function, user-defined operator or extended operator is used in the **initializer** clause or *combiner* expression, it must be accessible to the subprogram in which the corresponding **reduction** clause is specified.
- If the length type parameter is specified for a character type, it must be a constant, a colon or an `*`.
- If a character type with deferred or assumed length parameter is specified in a **declare reduction** directive, no other **declare reduction** directives with Fortran character type of the same kind parameter and the same *reduction-identifier* are allowed in the same scope.
- Any subroutine used in the **initializer** clause or *combiner* expression must not have any alternate returns appear in the argument list.

Cross References

- **reduction** clause, Section [2.15.3.6](#) on page [198](#).

1 2.17 Nesting of Regions

2 This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are
3 as follows:

- 4 • A worksharing region may not be closely nested inside a worksharing, explicit **task**,
5 **taskloop**, **critical**, **ordered**, **atomic**, or **master** region.
- 6 • A **barrier** region may not be closely nested inside a worksharing, explicit **task**, **taskloop**,
7 **critical**, **ordered**, **atomic**, or **master** region.
- 8 • A **master** region may not be closely nested inside a worksharing, **atomic**, explicit **task**, or
9 **taskloop** region.
- 10 • An **ordered** region may not be closely nested inside a **critical**, **atomic**, explicit **task**,
11 or **taskloop** region.
- 12 • An **ordered** region must be closely nested inside a loop region (or parallel loop region) with an
13 **ordered** clause.
- 14 • A **critical** region may not be nested (closely or otherwise) inside a **critical** region with
15 the same name. Note that this restriction is not sufficient to prevent deadlock.
- 16 • OpenMP constructs may not be nested inside an **atomic** region.
- 17 • OpenMP constructs may not be nested inside a **simd** region.
- 18 • If a **target**, **target update**, **target data**, **target enter data**, or
19 **target exit data** construct appears within a **target** region then the behavior is
20 unspecified.
- 21 • If specified, a **teams** construct must be contained within a **target** construct. That **target**
22 construct must contain no statements or directives outside of the **teams** construct.
- 23 • **distribute**, **parallel**, **parallel sections**, **parallel workshare**, and the
24 parallel loop and parallel loop SIMD constructs are the only OpenMP constructs that can be
25 closely nested in the **teams** region.
- 26 • A **distribute** construct must be closely nested in a **teams** region.
- 27 • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a
28 **task** construct and the **cancel** construct must be nested inside a **taskgroup** region. If
29 *construct-type-clause* is **sections**, the **cancel** construct must be closely nested inside a
30 **sections** or **section** construct. Otherwise, the **cancel** construct must be closely nested
31 inside an OpenMP construct that matches the type specified in *construct-type-clause* of the
32 **cancel** construct.
- 33 • A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be
34 nested inside a **task** construct and the **cancellation point** construct must be nested
35 inside a **taskgroup** region. A **cancellation point** construct for which

1 *construct-type-clause* is **sections** must be closely nested inside a **sections** or **section**
2 construct. Otherwise, a **cancellation point** construct must be closely nested inside an
3 OpenMP construct that matches the type specified in *construct-type-clause*.

2 **Runtime Library Routines**

3 This chapter describes the OpenMP API runtime library routines and is divided into the following
4 sections:

- 5
 - Runtime library definitions (Section 3.1 on page 227).
 - 6 • Execution environment routines that can be used to control and to query the parallel execution
7 environment (Section 3.2 on page 228).
 - 8 • Lock routines that can be used to synchronize access to data (Section 3.3 on page 267).
 - 9 • Portable timer routines (Section 3.4 on page 276).

10 Throughout this chapter, *true* and *false* are used as generic terms to simplify the description of the
11 routines.

▼ C / C++ ▼

12 *true* means a nonzero integer value and *false* means an integer value of zero.

▲ C / C++ ▲

▼ Fortran ▼

13 *true* means a logical value of **.TRUE.** and *false* means a logical value of **.FALSE.**

▲ Fortran ▲

▼ Fortran ▼

14 **Restrictions**

15 The following restriction applies to all OpenMP runtime library routines:

- 16
 - OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.

▲ Fortran ▲

1 3.1 Runtime Library Definitions

2 For each base language, a compliant implementation must supply a set of definitions for the
3 OpenMP API runtime library routines and the special data types of their parameters. The set of
4 definitions must contain a declaration for each OpenMP API runtime library routine and a
5 declaration for the *simple lock*, *nestable lock*, *schedule*, and *thread affinity policy* data types. In
6 addition, each set of definitions may specify other implementation specific values.

C / C++

7 The library routines are external functions with “C” linkage.

8 Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a
9 header file named **omp.h**. This file defines the following:

- 10 • The prototypes of all the routines in the chapter.
- 11 • The type **omp_lock_t**.
- 12 • The type **omp_nest_lock_t**.
- 13 • The type **omp_lock_hint_t**.
- 14 • The type **omp_sched_t**.
- 15 • The type **omp_proc_bind_t**.

16 See Section Section B.1 on page 324 for an example of this file.

C / C++

Fortran

17 The OpenMP Fortran API runtime library routines are external procedures. The return values of
18 these routines are of default kind, unless otherwise specified.

19 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter
20 shall be provided in the form of a Fortran **include** file named **omp_lib.h** or a Fortran 90
21 **module** named **omp_lib**. It is implementation defined whether the **include** file or the
22 **module** file (or both) is provided.

23 These files define the following:

- 24 • The interfaces of all of the routines in this chapter.
- 25 • The **integer parameter** **omp_lock_kind**.
- 26 • The **integer parameter** **omp_nest_lock_kind**.
- 27 • The **integer parameter** **omp_lock_hint_kind**.
- 28 • The **integer parameter** **omp_sched_kind**.
- 29 • The **integer parameter** **omp_proc_bind_kind**.

- The **integer parameter `openmp_version`** with a value `yyyymm` where `yyyy` and `mm` are the year and month designations of the version of the OpenMP Fortran API that the implementation supports. This value matches that of the C preprocessor macro `_OPENMP`, when a macro preprocessor is supported (see Section 2.2 on page 33).

See Section B.1 on page 328 and Section B.3 on page 332 for examples of these files.

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 accommodated. See Appendix B.4 for an example of such an extension.

Fortran

3.2 Execution Environment Routines

This section describes routines that affect and monitor threads, processors, and the parallel environment.

3.2.1 `omp_set_num_threads`

Summary

The `omp_set_num_threads` routine affects the number of threads to be used for subsequent parallel regions that do not specify a **`num_threads`** clause, by setting the value of the first element of the *nthreads-var* ICV of the current task.

Format

C / C++

```
void omp_set_num_threads(int num_threads);
```

C / C++

Fortran

```
subroutine omp_set_num_threads(num_threads)
integer num_threads
```

Fortran

1 **Constraints on Arguments**

2 The value of the argument passed to this routine must evaluate to a positive integer, or else the
3 behavior of this routine is implementation defined.

4 **Binding**

5 The binding task set for an `omp_set_num_threads` region is the generating task.

6 **Effect**

7 The effect of this routine is to set the value of the first element of the *nthreads-var* ICV of the
8 current task to the value specified in the argument.

9 See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a
10 `parallel` region.

11 **Cross References**

- 12 • *nthreads-var* ICV, see Section 2.3 on page 35.
- 13 • `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 290.
- 14 • `omp_get_max_threads` routine, see Section 3.2.3 on page 230.
- 15 • `parallel` construct, see Section 2.5 on page 45.
- 16 • `num_threads` clause, see Section 2.5 on page 45.

17 **3.2.2 omp_get_num_threads**

18 **Summary**

19 The `omp_get_num_threads` routine returns the number of threads in the current team.

20 **Format**

▼ C / C++ ▼

```
int omp_get_num_threads(void);
```

▲ C / C++ ▲

```
integer function omp_get_num_threads()
```

Binding

The binding region for an **omp_get_num_threads** region is the innermost enclosing **parallel** region.

Effect

The **omp_get_num_threads** routine returns the number of threads in the team executing the **parallel** region to which the routine region binds. If called from the sequential part of a program, this routine returns 1.

See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a **parallel** region.

Cross References

- **parallel** construct, see Section 2.5 on page 45.
- **omp_set_num_threads** routine, see Section 3.2.1 on page 228.
- **OMP_NUM_THREADS** environment variable, see Section 4.2 on page 290.

3.2.3 omp_get_max_threads

Summary

The **omp_get_max_threads** routine returns an upper bound on the number of threads that could be used to form a new team if a **parallel** construct without a **num_threads** clause were encountered after execution returns from this routine.

Format

C / C++

```
int omp_get_max_threads(void);
```

C / C++

Fortran

```
integer function omp_get_max_threads()
```

Fortran

Binding

The binding task set for an **omp_get_max_threads** region is the generating task.

Effect

The value returned by **omp_get_max_threads** is the value of the first element of the *nthreads-var* ICV of the current task. This value is also an upper bound on the number of threads that could be used to form a new team if a parallel region without a **num_threads** clause were encountered after execution returns from this routine.

See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a **parallel** region.

Note – The return value of the **omp_get_max_threads** routine can be used to dynamically allocate sufficient storage for all threads in the team formed at the subsequent active **parallel** region.

Cross References

- *nthreads-var* ICV, see Section 2.3 on page 35.
- **parallel** construct, see Section 2.5 on page 45.
- **num_threads** clause, see Section 2.5 on page 45.
- **omp_set_num_threads** routine, see Section 3.2.1 on page 228.
- **OMP_NUM_THREADS** environment variable, see Section 4.2 on page 290.

1 3.2.4 omp_get_thread_num

2 Summary

3 The **omp_get_thread_num** routine returns the thread number, within the current team, of the
4 calling thread.

5 Format

C / C++	
<pre>int omp_get_thread_num(void);</pre>	
C / C++	
Fortran	
<pre>integer function omp_get_thread_num()</pre>	
Fortran	

6 Binding

7 The binding thread set for an **omp_get_thread_num** region is the current team. The binding
8 region for an **omp_get_thread_num** region is the innermost enclosing **parallel** region.

9 Effect

10 The **omp_get_thread_num** routine returns the thread number of the calling thread, within the
11 team executing the **parallel** region to which the routine region binds. The thread number is an
12 integer between 0 and one less than the value returned by **omp_get_num_threads**, inclusive.
13 The thread number of the master thread of the team is 0. The routine returns 0 if it is called from
14 the sequential part of a program.

15 Note – The thread number may change during the execution of an untied task. The value returned
16 by **omp_get_thread_num** is not generally useful during the execution of such a task region.

17 Cross References


- 18
- **omp_get_num_threads** routine, see Section [3.2.2](#) on page [229](#).

1 3.2.5 `omp_get_num_procs`

2 Summary

3 The `omp_get_num_procs` routine returns the number of processors available to the device.

4 Format



5 Binding

6 The binding thread set for an `omp_get_num_procs` region is all threads on a device. The effect
7 of executing this routine is not related to any specific region corresponding to any construct or API
8 routine.

9 Effect

10 The `omp_get_num_procs` routine returns the number of processors that are available to the
11 device at the time the routine is called. Note that this value may change between the time that it is
12 determined by the `omp_get_num_procs` routine and the time that it is read in the calling
13 context due to system actions outside the control of the OpenMP implementation.

14 3.2.6 `omp_in_parallel`

15 Summary

16 The `omp_in_parallel` routine returns *true* if the *active-levels-var* ICV is greater than zero;
17 otherwise, it returns *false*.

1

C/C++

```
int omp_in_parallel(void);
```

C / C++

Fortran

```
logical function omp_in_parallel()
```

Fortran

2

3

4

5

6

7

8

- 9

- 10

11

12

13

14

15

Format

C / C++

```
void omp_set_dynamic(int dynamic_threads);
```

C / C++

Fortran

```
subroutine omp_set_dynamic(dynamic_threads)  
  logical dynamic_threads
```

Fortran

Binding

The binding task set for an **omp_set_dynamic** region is the generating task.

Effect

For implementations that support dynamic adjustment of the number of threads, if the argument to **omp_set_dynamic** evaluates to *true*, dynamic adjustment is enabled for the current task; otherwise, dynamic adjustment is disabled for the current task. For implementations that do not support dynamic adjustment of the number of threads this routine has no effect: the value of *dyn-var* remains *false*.

See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a **parallel** region.

Cross References


- *dyn-var* ICV, see Section 2.3 on page 35.
- **omp_get_num_threads** routine, see Section 3.2.2 on page 229.
- **omp_get_dynamic** routine, see Section 3.2.8 on page 236.
- **OMP_DYNAMIC** environment variable, see Section 4.3 on page 291.

1 3.2.8 omp_get_dynamic

2 Summary

3 The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether
4 dynamic adjustment of the number of threads is enabled or disabled.

5 Format



6 Binding

7 The binding task set for an `omp_get_dynamic` region is the generating task.

8 Effect

9 This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current
10 task; it returns *false*, otherwise. If an implementation does not support dynamic adjustment of the
11 number of threads, then this routine always returns *false*.

12 See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a
13 `parallel` region.

14 Cross References

- 15 • *dyn-var* ICV, see Section 2.3 on page 35.
- 16 • `omp_set_dynamic` routine, see Section 3.2.7 on page 234.
- 17 • `OMP_DYNAMIC` environment variable, see Section 4.3 on page 291.

1 3.2.9 `omp_get_cancellation`

2 Summary

3 The `omp_get_cancellation` routine returns *true* if cancellation of the active region has been
4 requested; it returns *false* otherwise.

5 Format

▼ C / C++ ▼	
<pre>int omp_get_cancellation(void);</pre>	
▲ C / C++ ▲	
▼ Fortran ▼	
<pre>logical function omp_get_cancellation()</pre>	
▲ Fortran ▲	

6 Binding

7 The binding task set for an `omp_get_cancellation` region is the whole program.

8 Effect

9 This routine returns *true* if cancellation is activated. It returns *false* otherwise.

10 Cross References

- 11 • `cancel` construct, see Section [2.14.1](#) on page [168](#)
- 12 • `cancel-var` ICV, see Section [2.3.1](#) on page [36](#).
- 13 • `OMP_CANCELLATION` environment variable, see Section [4.11](#) on page [297](#)

14 3.2.10 `omp_set_nested`

15 Summary

16 The `omp_set_nested` routine enables or disables nested parallelism, by setting the *nest-var*
17 ICV.

Format

C / C++

```
void omp_set_nested(int nested);
```

C / C++

Fortran

```
subroutine omp_set_nested(nested)  
logical nested
```

Fortran

Binding

The binding task set for an **omp_set_nested** region is the generating task.

Effect

For implementations that support nested parallelism, if the argument to **omp_set_nested** evaluates to *true*, nested parallelism is enabled for the current task; otherwise, nested parallelism is disabled for the current task. For implementations that do not support nested parallelism, this routine has no effect: the value of *nest-var* remains *false*.

See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a **parallel** region.

Cross References

- *nest-var* ICV, see Section 2.3 on page 35.
- **omp_set_max_active_levels** routine, see Section 3.2.15 on page 243.
- **omp_get_max_active_levels** routine, see Section 3.2.16 on page 245.
- **omp_get_nested** routine, see Section 3.2.11 on page 239.
- **OMP_NESTED** environment variable, see Section 4.6 on page 294.

1 3.2.11 omp_get_nested

2 Summary

3 The **omp_get_nested** routine returns the value of the *nest-var* ICV, which determines if nested
4 parallelism is enabled or disabled.

▼ C / C++ ▼

```
int omp_get_nested(void);
```

▲ C / C++ ▲

▼ Fortran ▼

```
logical function omp_get_nested()
```

▲ Fortran ▲

5 Binding

6 The binding task set for an **omp_get_nested** region is the generating task.

7 Effect

8 This routine returns *true* if nested parallelism is enabled for the current task; it returns *false*,
9 otherwise. If an implementation does not support nested parallelism, this routine always returns
10 *false*.

11 See Section 2.5.1 on page 49 for the rules governing the number of threads used to execute a
12 **parallel** region.

13 Cross References

- 14 • *nest-var* ICV, see Section 2.3 on page 35.
- 15 • **omp_set_nested** routine, see Section 3.2.10 on page 237.
- 16 • **OMP_NESTED** environment variable, see Section 4.6 on page 294.

1 3.2.12 omp_set_schedule

2 Summary

3 The **omp_set_schedule** routine affects the schedule that is applied when **runtime** is used as
4 schedule kind, by setting the value of the *run-sched-var* ICV.

5 Format

C / C++

```
void omp_set_schedule(omp_sched_t kind, int chunk_size);
```

C / C++

Fortran

```
subroutine omp_set_schedule(kind, chunk_size)  
integer (kind=omp_sched_kind) kind  
integer chunk_size
```

Fortran

6 Constraints on Arguments

7 The first argument passed to this routine can be one of the valid OpenMP schedule kinds (except for
8 **runtime**) or any implementation specific schedule. The C/C++ header file (**omp.h**) and the
9 Fortran include file (**omp_lib.h**) and/or Fortran 90 module file (**omp_lib**) define the valid
10 constants. The valid constants must include the following, which can be extended with
11 implementation specific values:

```
typedef enum omp_sched_t {
    omp_sched_static = 1,
    omp_sched_dynamic = 2,
    omp_sched_guided = 3,
    omp_sched_auto = 4
} omp_sched_t;
```

```
integer(kind=omp_sched_kind), parameter :: omp_sched_static = 1
integer(kind=omp_sched_kind), parameter :: omp_sched_dynamic = 2
integer(kind=omp_sched_kind), parameter :: omp_sched_guided = 3
integer(kind=omp_sched_kind), parameter :: omp_sched_auto = 4
```

Binding

The binding task set for an **omp_set_schedule** region is the generating task.

Effect

The effect of this routine is to set the value of the *run-sched-var* ICV of the current task to the values specified in the two arguments. The schedule is set to the schedule type specified by the first argument *kind*. It can be any of the standard schedule types or any other implementation specific one. For the schedule types **static**, **dynamic**, and **guided** the *chunk_size* is set to the value of the second argument, or to the default *chunk_size* if the value of the second argument is less than 1; for the schedule type **auto** the second argument has no meaning; for implementation specific schedule types, the values and associated meanings of the second argument are implementation defined.

Cross References

- *run-sched-var* ICV, see Section 2.3 on page 35.
- **omp_get_schedule** routine, see Section 3.2.13 on page 242.
- **OMP_SCHEDULE** environment variable, see Section 4.1 on page 289.
- Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 62.

1 3.2.13 omp_get_schedule

2 Summary

3 The **omp_get_schedule** routine returns the schedule that is applied when the runtime schedule
4 is used.

5 Format

C / C++	
<pre>void omp_get_schedule(omp_sched_t * kind, int * chunk_size);</pre>	
C / C++	
Fortran	
<pre>subroutine omp_get_schedule(kind, chunk_size) integer (kind=omp_sched_kind) kind integer chunk_size</pre>	
Fortran	

6 Binding

7 The binding task set for an **omp_get_schedule** region is the generating task.

8 Effect

9 This routine returns the *run-sched-var* ICV in the task to which the routine binds. The first
10 argument *kind* returns the schedule to be used. It can be any of the standard schedule types as
11 defined in Section 3.2.12 on page 240, or any implementation specific schedule type. The second
12 argument is interpreted as in the **omp_set_schedule** call, defined in Section 3.2.12 on
13 page 240.

14 Cross References

- 15 • *run-sched-var* ICV, see Section 2.3 on page 35.
- 16 • **omp_set_schedule** routine, see Section 3.2.12 on page 240.
- 17 • **OMP_SCHEDULE** environment variable, see Section 4.1 on page 289.
- 18 • Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 62.

1 3.2.14 `omp_get_thread_limit`

2 Summary

3 The `omp_get_thread_limit` routine returns the maximum number of OpenMP threads
4 available to participate in the current contention group.

5 Format

▼ C / C++ ▼	
<pre>int omp_get_thread_limit(void);</pre>	
▲ C / C++ ▲	▼ Fortran ▼
<pre>integer function omp_get_thread_limit()</pre>	
▲ Fortran ▲	

6 Binding

7 The binding thread set for an `omp_get_thread_limit` region is all threads on the device. The
8 effect of executing this routine is not related to any specific region corresponding to any construct
9 or API routine.

10 Effect

11 The `omp_get_thread_limit` routine returns the value of the *thread-limit-var* ICV.

12 Cross References

- 13 • *thread-limit-var* ICV, see Section [2.3](#) on page [35](#).
- 14 • `OMP_THREAD_LIMIT` environment variable, see Section [4.10](#) on page [297](#).

15 3.2.15 `omp_set_max_active_levels`

16 Summary

17 The `omp_set_max_active_levels` routine limits the number of nested active parallel
18 regions on the device, by setting the *max-active-levels-var* ICV

Format

C / C++

```
void omp_set_max_active_levels(int max_levels);
```

C / C++

Fortran

```
subroutine omp_set_max_active_levels(max_levels)  
integer max_levels
```

Fortran

Constraints on Arguments

The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise the behavior of this routine is implementation defined.

Binding

When called from a sequential part of the program, the binding thread set for an **omp_set_max_active_levels** region is the encountering thread. When called from within any explicit parallel region, the binding thread set (and binding region, if required) for the **omp_set_max_active_levels** region is implementation defined.

Effect

The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified in the argument.

If the number of parallel levels requested exceeds the number of levels of parallelism supported by the implementation, the value of the *max-active-levels-var* ICV will be set to the number of parallel levels supported by the implementation.

This routine has the described effect only when called from a sequential part of the program. When called from within an explicit **parallel** region, the effect of this routine is implementation defined.

Cross References


- *max-active-levels-var* ICV, see Section 2.3 on page 35.
- **omp_get_max_active_levels** routine, see Section 3.2.16 on page 245.
- **OMP_MAX_ACTIVE_LEVELS** environment variable, see Section 4.9 on page 297.

1 3.2.16 `omp_get_max_active_levels`

2 Summary

3 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*
4 ICV, which determines the maximum number of nested active parallel regions on the device.

5 Format



6 Binding

7 When called from a sequential part of the program, the binding thread set for an
8 `omp_get_max_active_levels` region is the encountering thread. When called from within
9 any explicit parallel region, the binding thread set (and binding region, if required) for the
10 `omp_get_max_active_levels` region is implementation defined.

11 Effect

12 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*
13 ICV, which determines the maximum number of nested active parallel regions on the device.

14 Cross References

- 15 • *max-active-levels-var* ICV, see Section 2.3 on page 35.
- 16 • `omp_set_max_active_levels` routine, see Section 3.2.15 on page 243.
- 17 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 297.

1 3.2.17 omp_get_level

2 Summary

3 The `omp_get_level` routine returns the value of the *levels-var* ICV.

4 Format

▼ C / C++ ▼	
<pre>int omp_get_level(void);</pre>	
▲ C / C++ ▲	
▼ Fortran ▼	
<pre>integer function omp_get_level()</pre>	
▲ Fortran ▲	

5 Binding

6 The binding task set for an `omp_get_level` region is the generating task.

7 Effect

8 The effect of the `omp_get_level` routine is to return the number of nested **parallel** regions
9 (whether active or inactive) enclosing the current task such that all of the **parallel** regions are
10 enclosed by the outermost initial task region on the current device.

11 Cross References

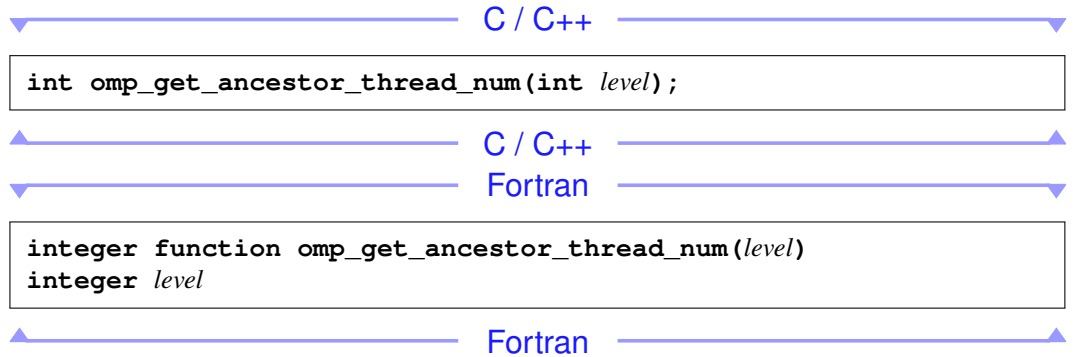
- 12 • *levels-var* ICV, see Section 2.3 on page 35.
- 13 • `omp_get_active_level` routine, see Section 3.2.20 on page 249.
- 14 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 297.

1 3.2.18 omp_get_ancestor_thread_num

2 Summary

3 The **omp_get_ancestor_thread_num** routine returns, for a given nested level of the current
4 thread, the thread number of the ancestor of the current thread.

5 Format



6 Binding

7 The binding thread set for an **omp_get_ancestor_thread_num** region is the encountering
8 thread. The binding region for an **omp_get_ancestor_thread_num** region is the innermost
9 enclosing **parallel** region.

10 Effect

11 The **omp_get_ancestor_thread_num** routine returns the thread number of the ancestor at a
12 given nest level of the current thread or the thread number of the current thread. If the requested
13 nest level is outside the range of 0 and the nest level of the current thread, as returned by the
14 **omp_get_level** routine, the routine returns -1.

15 Note – When the **omp_get_ancestor_thread_num** routine is called with a value of
16 **level=0**, the routine always returns 0. If **level=omp_get_level()**, the routine has the
17 same effect as the **omp_get_thread_num** routine.

Cross References

- `omp_get_level` routine, see Section 3.2.17 on page 246.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 232.
- `omp_get_team_size` routine, see Section 3.2.19 on page 248.

3.2.19 `omp_get_team_size`

Summary

The `omp_get_team_size` routine returns, for a given nested level of the current thread, the size of the thread team to which the ancestor or the current thread belongs.

Format

C / C++

```
int omp_get_team_size(int level);
```

C / C++

Fortran

```
integer function omp_get_team_size(level)  
integer level
```

Fortran

Binding

The binding thread set for an `omp_get_team_size` region is the encountering thread. The binding region for an `omp_get_team_size` region is the innermost enclosing `parallel` region.

Effect

The `omp_get_team_size` routine returns the size of the thread team to which the ancestor or the current thread belongs. If the requested nested level is outside the range of 0 and the nested level of the current thread, as returned by the `omp_get_level` routine, the routine returns -1. Inactive parallel regions are regarded like active parallel regions executed with one thread.

Note – When the `omp_get_team_size` routine is called with a value of `level=0`, the routine always returns 1. If `level=omp_get_level()`, the routine has the same effect as the `omp_get_num_threads` routine.

Cross References

- `omp_get_num_threads` routine, see Section 3.2.2 on page 229.
- `omp_get_level` routine, see Section 3.2.17 on page 246.
- `omp_get_ancestor_thread_num` routine, see Section 3.2.18 on page 247.

3.2.20 `omp_get_active_level`

Summary

The `omp_get_active_level` routine returns the value of the *active-level-vars* ICV..

Format

C / C++

```
int omp_get_active_level(void);
```

C / C++

Fortran

integer function omp_get_active_level()

Fortran

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- Binding

The binding task set for the an `omp_get_active_level` region is the generating task.

Effect

The effect of the `omp_get_active_level` routine is to return the number of nested, active **parallel** regions enclosing the current task such that all of the **parallel** regions are enclosed by the outermost initial task region on the current device.

Cross References

- *active-levels-var* ICV, see Section [2.3](#) on page [35](#).
 - `omp_get_level` routine, see Section [3.2.17](#) on page [246](#).

10

3.2.21 omp_in_final

- 11
- 12
- 13
- Summary

The `omp_in_final` routine returns *true* if the routine is executed in a final task region; otherwise, it returns *false*.

14

Format

C / C++

int omp_in_final(void);

C / C++

Fortran

logical function omp_in_final()

Fortran

Binding

The binding task set for an `omp_in_final` region is the generating task.

Effect

`omp_in_final` returns *true* if the enclosing task region is final. Otherwise, it returns *false*.

3.2.22 `omp_get_proc_bind`

Summary

The `omp_get_proc_bind` routine returns the thread affinity policy to be used for the subsequent nested `parallel` regions that do not specify a `proc_bind` clause.

Format

C / C++

```
omp_proc_bind_t omp_get_proc_bind(void);
```

C / C++

Fortran

```
integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()
```

Fortran

Constraints on Arguments

The value returned by this routine must be one of the valid affinity policy kinds. The C/ C++ header file (**omp.h**) and the Fortran include file (**omp_lib.h**) and/or Fortran 90 module file (**omp_lib**) define the valid constants. The valid constants must include the following:

C / C++

```
typedef enum omp_proc_bind_t {  
    omp_proc_bind_false = 0,  
    omp_proc_bind_true = 1,  
    omp_proc_bind_master = 2,  
    omp_proc_bind_close = 3,  
    omp_proc_bind_spread = 4  
} omp_proc_bind_t;
```

C / C++

Fortran

```
integer (kind=omp_proc_bind_kind), &  
    parameter :: omp_proc_bind_false = 0  
integer (kind=omp_proc_bind_kind), &  
    parameter :: omp_proc_bind_true = 1  
integer (kind=omp_proc_bind_kind), &  
    parameter :: omp_proc_bind_master = 2  
integer (kind=omp_proc_bind_kind), &  
    parameter :: omp_proc_bind_close = 3  
integer (kind=omp_proc_bind_kind), &  
    parameter :: omp_proc_bind_spread = 4
```

Fortran

Binding

The binding task set for an **omp_get_proc_bind** region is the generating task

Effect

The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current task. See Section 2.5.2 on page 51 for the rules governing the thread affinity policy.

1 **Cross References**

- 2 • *bind-var* ICV, see Section 2.3 on page 35.
- 3 • **OMP_PROC_BIND** environment variable, see Section 4.4 on page 291.
- 4 • Controlling OpenMP thread affinity, see Section 2.5.2 on page 51.

5 **3.2.23 omp_get_num_places**

6 **Summary**

7 The **omp_get_num_places** routine returns the number of places available to the execution

8 environment in the place list.

9 **Format**

▼ C / C++ ▼

```
int omp_get_num_places(void);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer function omp_get_num_places()
```

▲ Fortran ▲

10 **Binding**

11 The binding thread set for an **omp_get_num_places** region is all threads on a device. The

12 effect of executing this routine is not related to any specific region corresponding to any construct

13 or API routine.

14 **Effect**

15 The **omp_get_num_places** routine returns the number of places in the place list. This value is

16 equivalent to the number of places in the *place-partition-var* ICV in the execution environment of

17 the initial task.

Cross References

- *place-partition-var* ICV, see Section 2.3 on page 35.
- **OMP_PLACES** environment variable, see Section 4.5 on page 292.

3.2.24 omp_get_place_num_procs

Summary

The **omp_get_place_num_procs** routine returns the number of processors available to the execution environment in the specified place.

Format

C / C++

```
int omp_get_place_num_procs(int place_num);
```

C / C++

Fortran

```
integer function omp_get_place_num_procs(place_num)  
integer place_num
```

Fortran

Binding

The binding thread set for an **omp_get_place_num_procs** region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The **omp_get_place_num_procs** routine returns the number of processors associated with the place numbered **place_num**. The routine returns zero when **place_num** is negative, or is equal to or larger than the value returned by **omp_get_num_places()**.

Cross References

- `OMP_PLACES` environment variable, see Section 4.5 on page 292.

3.2.25 `omp_get_place_proc_ids`

Summary

The `omp_get_place_proc_ids` routine returns the numerical identifiers of the processors available to the execution environment in the specified place.

Format

C / C++

```
void omp_get_place_proc_ids(int place_num, int *ids);
```

C / C++

Fortran

```
subroutine omp_get_place_proc_ids(place_num, ids)
integer place_num
integer ids(*)
```

Fortran

Binding

The binding thread set for an `omp_get_place_proc_ids` region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The `omp_get_place_proc_ids` routine returns the numerical identifiers of each processor associated with the place numbered `place_num`. The numerical identifiers are nonnegative, and their meaning is implementation defined. The numerical identifiers are returned in the array `ids` and their order in the array is implementation defined. The array must be sufficiently large to contain `omp_get_place_num_procs(place_num)` integers; otherwise, the behavior is unspecified. The routine has no effect when `place_num` has a negative value, or a value equal or larger than `omp_get_num_places()`.

Cross References

- **OMP_PLACES** environment variable, see Section 4.5 on page 292.

3.2.26 `omp_get_place_num`

Summary

The `omp_get_place_num` routine returns the place number of the place to which the encountering thread is bound.

Format

C / C++

```
int omp_get_place_num(void);
```

C / C++

Fortran

```
integer function omp_get_place_num()
```

Fortran

Binding

The binding thread set for an `omp_get_place_num` region is the encountering thread.

Effect

When the encountering thread is bound to a place, the `omp_get_place_num` routine returns the place number associated with the thread. The returned value is between 0 and one less than the value returned by `omp_get_num_places()`, inclusive. When the encountering thread is not bound to a place, the routine returns negative one.

Cross References

- **OMP_PLACES** environment variable, see Section 4.5 on page 292.
- Controlling OpenMP thread affinity, see Section 2.5.2 on page 51.

1 3.2.27 `omp_get_partition_num_places`

2 Summary

3 The `omp_get_partition_num_places` routine returns the number of places in the place
4 partition of the innermost implicit task.

5 Format

C / C++	
<pre>int omp_get_partition_num_places(void);</pre>	
C / C++	
Fortran	
<pre>integer function omp_get_partition_num_places()</pre>	
Fortran	

6 Binding

7 The binding task set for an `omp_get_partition_num_places` region is the encountering
8 implicit task.

9 Effect

10 The `omp_get_partition_num_places` routine returns the number of places in the
11 *place-partition-var* ICV.

12 Cross References

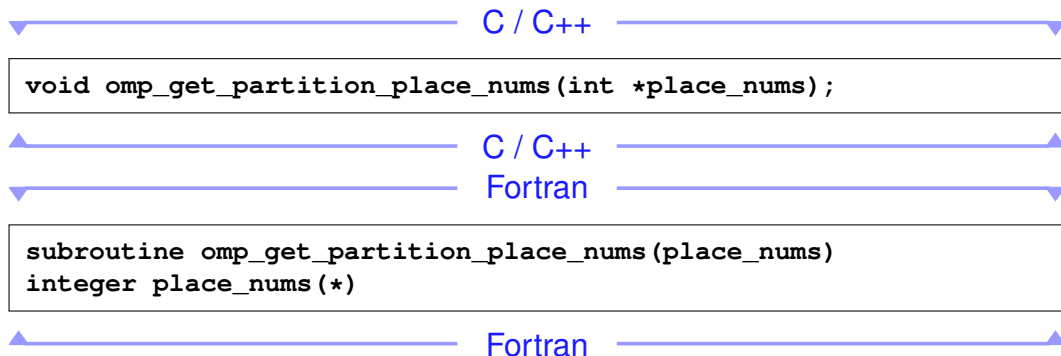
- 13 • *place-partition-var* ICV, see Section 2.3 on page 35.
- 14 • `OMP_PLACES` environment variable, see Section 4.5 on page 292.
- 15 • Controlling OpenMP thread affinity, see Section 2.5.2 on page 51.

1 3.2.28 omp_get_partition_place_nums

2 Summary

3 The `omp_get_partition_place_nums` routine returns the list of place numbers
4 corresponding to the places in the *place-partition-var* ICV of the innermost implicit task.

5 Format



6 Binding

7 The binding task set for an `omp_get_partition_place_nums` region is the encountering
8 implicit task.

9 Effect

10 The `omp_get_partition_place_nums` routine returns the list of place numbers
11 corresponding to the places in the *place-partition-var* ICV of the innermost implicit task. The array
12 must be sufficiently large to contain `omp_get_partition_num_places()` integers;
13 otherwise, the behavior is unspecified.

14 Cross References

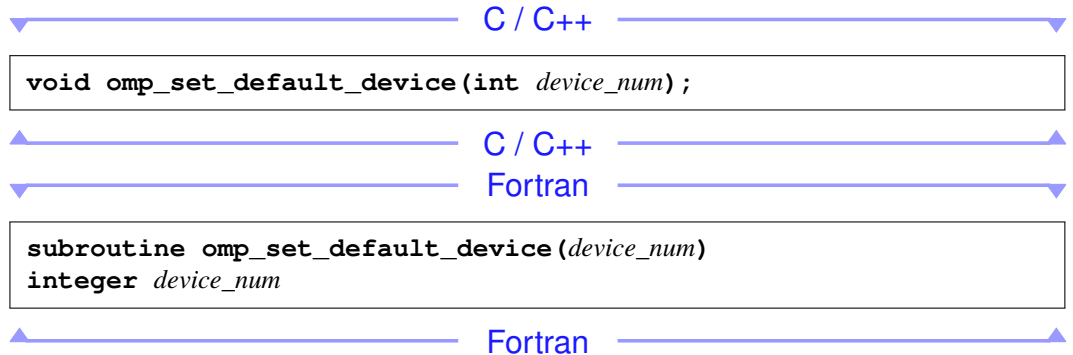
- 15 • *place-partition-var* ICV, see Section 2.3 on page 35.
- 16 • `OMP_PLACES` environment variable, see Section 4.5 on page 292.
- 17 • Controlling OpenMP thread affinity, see Section 2.5.2 on page 51.

1 3.2.29 `omp_set_default_device`

2 Summary

3 The `omp_set_default_device` routine controls the default target device by assigning the
4 value of the *default-device-var* ICV.

5 Format



6 Binding

7 The binding task set for an `omp_set_default_device` region is the generating task.

8 Effect

9 The effect of this routine is to set the value of the *default-device-var* ICV of the current task to the
10 value specified in the argument. When called from within a **target** region the effect of this
11 routine is unspecified.

12 Cross References

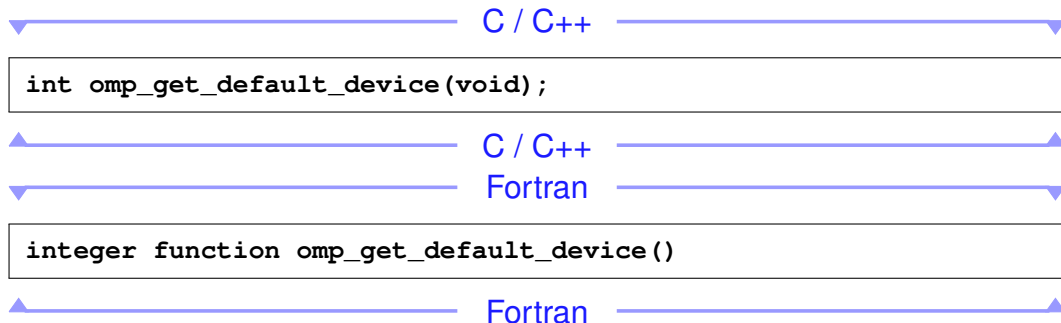
- 13 • *default-device-var*, see Section 2.3 on page 35.
- 14 • `omp_get_default_device`, see Section 3.2.30 on page 260.
- 15 • `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 299

1 3.2.30 omp_get_default_device

2 Summary

3 The `omp_get_default_device` routine returns the default target device.

4 Format



5 Binding

6 The binding task set for an `omp_get_default_device` region is the generating task.

7 Effect

8 The `omp_get_default_device` routine returns the value of the *default-device-var* ICV of the
9 current task. When called from within a **target** region the effect of this routine is unspecified.

10 Cross References

- 11 • *default-device-var*, see Section 2.3 on page 35.
- 12 • `omp_set_default_device`, see Section 3.2.29 on page 259.
- 13 • `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 299.

14 3.2.31 omp_get_num_devices

15 Summary

16 The `omp_get_num_devices` routine returns the number of target devices.

Format

C / C++

```
int omp_get_num_devices(void);
```

C / C++

Fortran

```
integer function omp_get_num_devices()
```

Fortran

Binding

The binding task set for an **omp_get_num_devices** region is the generating task.

Effect

The **omp_get_num_devices** routine returns the number of available target devices. When called from within a **target** region the effect of this routine is unspecified.

Cross References

None.

3.2.32 omp_get_num_teams

Summary

The **omp_get_num_teams** routine returns the number of teams in the current **teams** region.

1

Format

▼ C / C++ ▼

```
int omp_get_num_teams(void);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer function omp_get_num_teams()
```

▲ Fortran ▲

2

Binding

3

The binding task set for an **omp_get_num_teams** region is the generating task

4

Effect

5

The effect of this routine is to return the number of teams in the current **teams** region. The routine returns 1 if it is called from outside of a **teams** region.

6

7

Cross References

8

- **teams** construct, see Section [2.10.5](#) on page [106](#).

1 3.2.33 `omp_get_team_num`

2 Summary

3 The `omp_get_team_num` routine returns the team number of the calling thread.

4 Format

▼ C / C++ ▼	
<pre>int omp_get_team_num(void);</pre>	
▲ C / C++ ▲	
▼ Fortran ▼	
<pre>integer function omp_get_team_num()</pre>	
▲ Fortran ▲	

5 Binding

6 The binding task set for an `omp_get_team_num` region is the generating task.

7 Effect

8 The `omp_get_team_num` routine returns the team number of the calling thread. The team
9 number is an integer between 0 and one less than the value returned by `omp_get_num_teams`,
10 inclusive. The routine returns 0 if it is called outside of a `teams` region.

11 Cross References

- 12 • `teams` construct, see Section [2.10.5](#) on page [106](#).
- 13 • `omp_get_num_teams` routine, see Section [3.2.32](#) on page [261](#).

1 **3.2.34 omp_is_initial_device**

2 **Summary**

3 The `omp_is_initial_device` routine returns *true* if the current task is executing on the host
4 device; otherwise, it returns *false*.

5 **Format**

C / C++

```
int omp_is_initial_device(void);
```

C / C++

Fortran

```
logical function omp_is_initial_device()
```

Fortran

6 **Binding**

7 The binding task set for an `omp_is_initial_device` region is the generating task.

8 **Effect**

9 The effect of this routine is to return *true* if the current task is executing on the host device;
10 otherwise, it returns *false*.

11 **Cross References**

- 12
 - `target` construct, see Section [2.10.2](#) on page [95](#)

13 **3.2.35 omp_get_initial_device**

14 **Summary**

15 The `omp_get_initial_device` routine returns a device number representing the host device.

16 **Format**

C / C++

```
int omp_get_initial_device(void);
```

C / C++

Fortran

```
integer function omp_get_initial_device()
```

Fortran

Binding

The binding task set for an **omp_get_initial_device** region is the generating task.

Effect

The effect of this routine is to return the device number of the host device. The value of the device number is implementation defined. If it is between 0 and one less than **omp_get_num_devices()** then it is valid for use with all device constructs and routines; if it is outside that range, then it is only valid for use with the device memory routines (Section 3.5 on page 279) and not in the **device** clause. When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** construct, see Section 2.10.2 on page 95

3.2.36 omp_get_max_task_priority

Summary

The **omp_get_max_task_priority** routine returns the maximum value that can be specified in the **priority** clause.

Format

C / C++

```
int omp_get_max_task_priority(void);
```

C / C++

Fortran

```
integer function omp_get_max_task_priority()
```

Fortran

Binding

The binding thread set for an **omp_get_max_task_priority** region is all threads on the device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The **omp_get_max_task_priority** routine returns the value of the *max-task-priority-var* ICV, which determines the maximum value that can be specified in the **priority** clause.

3.3 Lock Routines

The OpenMP runtime library includes a set of general-purpose lock routines that can be used for synchronization. These general-purpose lock routines operate on OpenMP locks that are represented by OpenMP lock variables. OpenMP lock variables must be accessed only through the routines described in this section; programs that otherwise access OpenMP lock variables are non-conforming.

An OpenMP lock can be in one of the following states: *uninitialized*, *unlocked*, or *locked*. If a lock is in the unlocked state, a task can *set* the lock, which changes its state to *locked*. The task that sets the lock is then said to *own* the lock. A task that owns a lock can *unset* that lock, returning it to the *unlocked* state. A program in which a task unsets a lock that is owned by another task is non-conforming.

Two types of locks are supported: *simple locks* and *nestable locks*. A *nestable lock* can be set multiple times by the same task before being unset; a *simple lock* cannot be set if it is already owned by the task trying to set it. *Simple lock* variables are associated with *simple locks* and can only be passed to *simple lock* routines. *Nestable lock* variables are associated with *nestable locks* and can only be passed to *nestable lock* routines.

Each type of lock can also have a *lock hint* that contains information about the intended usage of the lock by the application code. The effect of the lock hint is implementation defined, that is, the OpenMP implementation can use this hint to select a proper usage-specific lock if supported by the implementation. Lock hints do not change the mutual exclusion semantics of locks, but allow the user to suggest the use of a specific implementation to the runtime. A conforming implementation can safely ignore the lock hint and instead use locks without hints.

Constraints on the state and ownership of the lock accessed by each of the lock routines are described with the routine. If these constraints are not met, the behavior of the routine is unspecified.

The OpenMP lock routines access a lock variable in such a way that they always read and update the most current value of the lock variable. It is not necessary for an OpenMP program to include explicit **flush** directives to ensure that the lock variable's value is consistent among different tasks.

Binding

The binding thread set for all lock routine regions is all threads in the contention group. As a consequence, for each OpenMP lock, the lock routine effects relate to all tasks that call the routines, without regard to which teams the threads in the contention group executing the tasks belong.

Simple Lock Routines

C / C++

The type **omp_lock_t** is a data type capable of representing a simple lock. For the following routines, a simple lock variable must be of **omp_lock_t** type. All simple lock routines require an argument that is a pointer to a variable of type **omp_lock_t**.

C / C++

Fortran

For the following routines, a simple lock variable must be an integer variable of **kind=omp_lock_kind**.

Fortran

The simple lock routines are as follows:

- The **omp_init_lock** routine initializes a simple lock.
- The **omp_init_lock_with_hint** routine initializes a simple lock and attaches a hint to it.
- The **omp_destroy_lock** routine uninitializes a simple lock.
- The **omp_set_lock** routine waits until a simple lock is available, and then sets it.
- The **omp_unset_lock** routine unsets a simple lock.
- The **omp_test_lock** routine tests a simple lock, and sets it if it is available.

Nestable Lock Routines

C / C++

The type **omp_nest_lock_t** is a data type capable of representing a nestable lock. For the following routines, a nested lock variable must be of **omp_nest_lock_t** type. All nestable lock routines require an argument that is a pointer to a variable of type **omp_nest_lock_t**.

C / C++

Fortran

For the following routines, a nested lock variable must be an integer variable of **kind=omp_nest_lock_kind**.

The nestable lock routines are as follows:

- The `omp_init_nest_lock` routine initializes a nestable lock.
- The `omp_init_nest_lock_with_hint` routine initializes a nestable lock and attaches a hint to it.
- The `omp_destroy_nest_lock` routine uninitializes a nestable lock.
- The `omp_set_nest_lock` routine waits until a nestable lock is available, and then sets it.
- The `omp_unset_nest_lock` routine unsets a nestable lock.
- The `omp_test_nest_lock` routine tests a nestable lock, and sets it if it is available

Restrictions

OpenMP lock routines have the following restrictions:

- The use of the same OpenMP lock in different contention groups results in unspecified behavior.

3.3.1 `omp_init_lock` and `omp_init_nest_lock`

Summary

These routines initialize an OpenMP lock without a hint.

Format

C / C++

```
void omp_init_lock(omp_lock_t *lock);
void omp_init_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_init_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_init_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

Constraints on Arguments

A program that accesses a lock that is not in the uninitialized state through either routine is non-conforming.

Effect

The effect of these routines is to initialize the lock to the unlocked state; that is, no task owns the lock. In addition, the nesting count for a nestable lock is set to zero.

3.3.2 `omp_init_lock_with_hint` and `omp_init_nest_lock_with_hint`

Summary

These routines initialize an OpenMP lock with a hint. The effect of the hint is implementation-defined. The OpenMP implementation can ignore the hint without changing program semantics.

Format

C / C++

```
void omp_init_lock_with_hint(omp_lock_t *lock,  
                             omp_lock_hint_t hint);  
void omp_init_nest_lock_with_hint(omp_nest_lock_t *lock,  
                                  omp_lock_hint_t hint);
```

C / C++

Fortran

```
subroutine omp_init_lock_with_hint(svar, hint)  
integer (kind=omp_lock_kind) svar  
integer (kind=omp_lock_hint_kind) hint  
  
subroutine omp_init_nest_lock_with_hint(nvar, hint)  
integer (kind=omp_nest_lock_kind) nvar  
integer (kind=omp_lock_hint_kind) hint
```

Fortran

Constraints on Arguments

A program that accesses a lock that is not in the uninitialized state through either routine is non-conforming.

The second argument passed to this routine (*hint*) can be one of the valid OpenMP lock hints below or any implementation-defined hint. The C/C++ header file (**omp.h**) and the Fortran include file (**omp_lib.h**) and/or Fortran 90 module file (**omp_lib**) define the valid lock hint constants. The valid constants must include the following, which can be extended with implementation-defined values:

C / C++

```
typedef enum omp_lock_hint_t {  
    omp_lock_hint_none = 0,  
    omp_lock_hint_uncontended = 1,  
    omp_lock_hint_contended = 2,  
    omp_lock_hint_nonspeculative = 4,  
    omp_lock_hint_speculative = 8,  
} omp_lock_hint_t;
```

C / C++

Fortran

```
integer (kind=omp_lock_hint_kind), &  
    parameter :: omp_lock_hint_none = 0  
integer (kind=omp_lock_hint_kind), &  
    parameter :: omp_lock_hint_uncontended = 1  
integer (kind=omp_lock_hint_kind), &  
    parameter :: omp_lock_hint_contended = 2  
integer (kind=omp_lock_hint_kind), &  
    parameter :: omp_lock_hint_nonspeculative = 4  
integer (kind=omp_lock_hint_kind), &  
    parameter :: omp_lock_hint_speculative = 8
```

Fortran

The hints can be combined by using the + or | operators in C/C++ or the + operator in Fortran. The effect of the combined hint is implementation defined and can be ignored by the implementation. Combining **omp_lock_hint_none** with any other hint is equivalent to specifying the other hint. The following restrictions apply to combined hints; violating these restrictions results in unspecified behavior:

- the hints **omp_lock_hint_uncontended** and **omp_lock_hint_contended** cannot be combined,
- the hints **omp_lock_hint_nonspeculative** and **omp_lock_hint_speculative** cannot be combined.

Note – Future OpenMP specifications may add additional hints to the `omp_lock_hint_t` type and the `omp_lock_hint_kind` kind. Implementers are advised to add implementation-defined hints starting from the most significant bit of the `omp_lock_hint_t` type and `omp_lock_hint_kind` kind and to include the name of the implementation in the name of the added hint to avoid name conflicts with other OpenMP implementations.

Effect

The effect of these routines is to initialize the lock to the unlocked state and, optionally, to choose a specific lock implementation based on the hint. After initialization no task owns the lock. In addition, the nesting count for a nestable lock is set to zero.

3.3.3 `omp_destroy_lock` and `omp_destroy_nest_lock`

Summary

These routines ensure that the OpenMP lock is uninitialized

Format

C / C++

```
void omp_destroy_lock(omp_lock_t *lock);
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_destroy_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_destroy_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

1 **Constraints on Arguments**

2 A program that accesses a lock that is not in the unlocked state through either routine is
3 non-conforming.

4 **Effect**

5 The effect of these routines is to change the state of the lock to uninitialized.

6 **3.3.4 omp_set_lock and omp_set_nest_lock**

7 **Summary**

8 These routines provide a means of setting an OpenMP lock. The calling task region behaves as if it
9 were suspended until the lock can be set by this thread.

10 **Format**

▼────────────────── C / C++ ───────────────────►

```
void omp_set_lock(omp_lock_t *lock);  
void omp_set_nest_lock(omp_nest_lock_t *lock);
```

▲────────────────── C / C++ ───────────────────►

▼────────────────── Fortran ───────────────────►

```
subroutine omp_set_lock(svar)  
integer (kind=omp_lock_kind) svar  
  
subroutine omp_set_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

▲────────────────── Fortran ───────────────────►

11 **Constraints on Arguments**

12 A program that accesses a lock that is in the uninitialized state through either routine is
13 non-conforming. A simple lock accessed by **omp_set_lock** that is in the locked state must not
14 be owned by the task that contains the call or deadlock will result.

Effect

Each of these routines has an effect equivalent to suspension of the task executing the routine until the specified lock is available.

Note – The semantics of these routines is specified *as if* they serialize execution of the region guarded by the lock. However, implementations may implement them in other ways provided that the isolation properties are respected so that the actual execution delivers a result that could arise from some serialization.

A simple lock is available if it is unlocked. Ownership of the lock is granted to the task executing the routine.

A nestable lock is available if it is unlocked or if it is already owned by the task executing the routine. The task executing the routine is granted, or retains, ownership of the lock, and the nesting count for the lock is incremented.

3.3.5 `omp_unset_lock` and `omp_unset_nest_lock`

Summary

These routines provide the means of unsetting an OpenMP lock.

Format

C / C++

```
void omp_unset_lock(omp_lock_t *lock);  
void omp_unset_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_unset_lock(svar)  
integer (kind=omp_lock_kind) svar  
  
subroutine omp_unset_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

Constraints on Arguments

A program that accesses a lock that is not in the locked state or that is not owned by the task that contains the call through either routine is non-conforming.

Effect

For a simple lock, the `omp_unset_lock` routine causes the lock to become unlocked.

For a nestable lock, the `omp_unset_nest_lock` routine decrements the nesting count, and causes the lock to become unlocked if the resulting nesting count is zero.

For either routine, if the lock becomes unlocked, and if one or more task regions were effectively suspended because the lock was unavailable, the effect is that one task is chosen and given ownership of the lock.

3.3.6 `omp_test_lock` and `omp_test_nest_lock`

Summary

These routines attempt to set an OpenMP lock but do not suspend execution of the task executing the routine.

Format

C / C++

```
int omp_test_lock(omp_lock_t *lock);  
int omp_test_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
logical function omp_test_lock(svar)  
integer (kind=omp_lock_kind) svar  
integer function omp_test_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

Constraints on Arguments

A program that accesses a lock that is in the uninitialized state through either routine is non-conforming. The behavior is unspecified if a simple lock accessed by `omp_test_lock` is in the locked state and is owned by the task that contains the call.

Effect

These routines attempt to set a lock in the same manner as `omp_set_lock` and `omp_set_nest_lock`, except that they do not suspend execution of the task executing the routine.

For a simple lock, the `omp_test_lock` routine returns *true* if the lock is successfully set; otherwise, it returns *false*.

For a nestable lock, the `omp_test_nest_lock` routine returns the new nesting count if the lock is successfully set; otherwise, it returns zero.

3.4 Timing Routines

This section describes routines that support a portable wall clock timer.

3.4.1 `omp_get_wtime`

Summary

The `omp_get_wtime` routine returns elapsed wall clock time in seconds.

Format

C / C++	
<pre>double omp_get_wtime(void);</pre>	
C / C++	
Fortran	

```
double precision function omp_get_wtime()
```

Fortran

Binding

The binding thread set for an **omp_get_wtime** region is the encountering thread. The routine's return value is not guaranteed to be consistent across any set of threads.

Effect

The **omp_get_wtime** routine returns a value equal to the elapsed wall clock time in seconds since some "time in the past". The actual "time in the past" is arbitrary, but it is guaranteed not to change during the execution of the application program. The time returned is a "per-thread time", so it is not required to be globally consistent across all the threads participating in an application.

Note – It is anticipated that the routine will be used to measure elapsed times as shown in the following example:

C / C++

```
double start;  
double end;  
start = omp_get_wtime();  
... work to be timed ...  
end = omp_get_wtime();  
printf("Work took %f seconds\n", end - start);
```

C / C++

Fortran

```
DOUBLE PRECISION START, END  
START = omp_get_wtime()  
... work to be timed ...  
END = omp_get_wtime()  
PRINT *, "Work took", END - START, "seconds"
```


Fortran

1 3.4.2 omp_get_wtick

2 Summary

3 The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

4 Format



5 Binding

6 The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's
7 return value is not guaranteed to be consistent across any set of threads.

8 Effect

9 The `omp_get_wtick` routine returns a value equal to the number of seconds between successive
10 clock ticks of the timer used by `omp_get_wtime`.

3.5 Device Memory Routines

This section describes routines that support allocation and management of pointers in the data environments of target devices.

3.5.1 `omp_target_alloc`

Summary

The `omp_target_alloc` routine allocates memory in a device data environment.

Format

```
void* omp_target_alloc(size_t size, int device_num);
```

Effect

The `omp_target_alloc` routine returns the device address of a storage location of *size* bytes. The storage location is dynamically allocated in the device data environment of the device specified by *device_num*, which must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`. When called from within a **target** region the effect of this routine is unspecified.

The `omp_target_alloc` routine returns **NULL** if it cannot dynamically allocate the memory in the device data environment.

The device address returned by `omp_target_alloc` can be used in an `is_device_ptr` clause, Section 2.10.2 on page 95.

Pointer arithmetic is not supported on the device address returned by `omp_target_alloc`.

Cross References

- **target** construct, see Section 2.10.2 on page 95
- `omp_target_free` routine, see Section 3.5.2 on page 280
- `omp_get_initial_device` routine, see Section 3.2.35 on page 264
- `omp_get_num_devices` routine, see Section 3.2.31 on page 260

3.5.2 `omp_target_free`

Summary

The `omp_target_free` routine frees the device memory allocated by the `omp_target_alloc` routine.

Format

```
void omp_target_free(void * device_ptr, int device_num);
```

Constraints on Arguments

A program that calls `omp_target_free` with a non-`NULL` pointer that does not have a value returned from `omp_target_alloc` is non-conforming. When called from within a `target` region the effect of this routine is unspecified. The `device_num` must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`.

Effect

The `omp_target_free` routine frees the memory in the device data environment associated with `device_ptr`. If `device_ptr` is `NULL`, the operation is ignored.

Synchronization must be inserted to ensure that all accesses to `device_ptr` are completed before the call to `omp_target_free`.

Cross References

- `target` construct, see Section 2.10.2 on page 95
- `omp_target_alloc` routine, see Section 3.5.1 on page 279
- `omp_get_initial_device` routine, see Section 3.2.35 on page 264
- `omp_get_num_devices` routine, see Section 3.2.31 on page 260

3.5.3 `omp_target_is_present`

Summary

The `omp_target_is_present` routine validates whether a host pointer has an associated device buffer on a given device.

Format

```
int omp_target_is_present(void * ptr, size_t offset,
                          int device_num);
```

Constraints on Arguments

The value of *ptr* must be a valid host pointer or **NULL**. The device specified by *device_num* must be valid. When called from within a **target** region the effect of this routine is unspecified. The *device_num* must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`.

Effect

This routine returns *true* if the specified pointer and offset would be found present on device *device_num* by a **map** clause; otherwise, it returns *false*.

Cross References

- **target** construct, see Section [2.10.2](#) on page [95](#)
- **map** clause, see Section [2.15.5.1](#) on page [213](#).
- `omp_get_initial_device` routine, see Section [3.2.35](#) on page [264](#)
- `omp_get_num_devices` routine, see Section [3.2.31](#) on page [260](#)

3.5.4 `omp_target_memcpy`

Summary

The `omp_target_memcpy` routine copies memory between any combination of host and device pointers.

Format

```
int omp_target_memcpy(void * dst, void * src, size_t length,
                      size_t dst_offset, size_t src_offset,
                      int dst_device_num, int src_device_num);
```

Constraints on Arguments

Each device must be compatible with the device pointer specified on the same side of the copy. The `dst_device_num` and `src_device_num` must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`.

Effect

length bytes of memory at offset *src_offset* from *src* are copied to *dst* starting at offset *dst_offset*. The return value is zero on success and non-zero on failure. The host device and host device data environment can be referenced with the device number returned by `omp_get_initial_device`. This routine contains a task scheduling point.

When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** construct, see Section [2.10.2](#) on page [95](#)
- `omp_target_alloc` routine, see Section [3.5.1](#) on page [279](#)
- `omp_get_initial_device` routine, see Section [3.2.35](#) on page [264](#)

3.5.5 `omp_target_memcpy_rect`

Summary

The `omp_target_memcpy_rect` routine copies a rectangular subvolume from a multi-dimensional array to another multi-dimensional array. The copies can use any combination of host and device pointers.

Format

```
int omp_target_memcpy_rect(
    void * dst, void * src,
    size_t element_size,
    int num_dims,
    const size_t* volume,
    const size_t* dst_offsets,
    const size_t* src_offsets,
    const size_t* dst_dimensions,
    const size_t* src_dimensions,
    int dst_device_num, int src_device_num);
```

Constraints on Arguments

The length of the offset and dimension arrays must be at least the value of `num_dims`. The `dst_device_num` and `src_device_num` must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`.

The value of `num_dims` must be between 1 and the implementation-defined limit, which must be at least three.

Effect

This routine copies a multi-dimensional rectangular subvolume of *src* to *dst*. The volume is specified in terms of the size of the element, number of dimensions, and constant arrays of length *num_dims*. The maximum number of dimensions supported is at least three, support for higher dimensionality is implementation defined. The volume array specifies the length, in number of elements, to copy in each dimension from *src* to *dst*. The *dst_offsets* (*src_offsets*) parameter specifies number of elements from the origin of *dst* (*src*) in elements. The *dst_dimensions* (*src_dimensions*) parameter specifies the length of each dimension of *dst* (*src*)

The routine returns zero if successful. If both *dst* and *src* are **NULL** pointers, the routine returns the number of dimensions supported by the implementation for the specified device numbers. The host device and host device data environment can be referenced with the device number returned by **omp_get_initial_device**. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** construct, see Section 2.10.2 on page 95
- **omp_target_alloc** routine, see Section 3.5.1 on page 279
- **omp_get_initial_device** routine, see Section 3.2.35 on page 264

3.5.6 omp_target_associate_ptr

Summary

The **omp_target_associate_ptr** routine maps a device pointer, which may be returned from **omp_target_alloc** or implementation-defined runtime routines, to a host pointer.

Format

```
int omp_target_associate_ptr(void * host_ptr, void * device_ptr,
                             size_t size, size_t device_offset,
                             int device_num);
```

Constraints on Arguments

The value of *device_ptr* value must be a valid pointer to device memory for the device denoted by the value of *device_num*. The *device_num* argument must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`.

Effect

The `omp_target_associate_ptr` routine associates a device pointer with a host pointer such that when the host pointer appears in a subsequent **map** clause, the associated device pointer is used as the target for data motion associated with that host pointer. The *device_offset* parameter specifies what offset into *device_ptr* will be used as the base address for the device side of the mapping. The reference count of the resulting mapping will be infinite. After being successfully associated, the buffer pointed to by the device pointer is invalidated and accessing data directly through the device pointer results in unspecified behavior. The pointer can be retrieved for other uses by disassociating it. When called from within a **target** region the effect of this routine is unspecified.

The routine returns zero if successful. Otherwise it returns a non-zero value.

Only one device buffer can be associated with a given host pointer value and device number pair. Attempting to associate a second buffer will return non-zero. Associating the same pair of pointers on the same device with the same offset has no effect and returns zero. The `omp_target_is_present` function can be used to test whether a given host pointer has a corresponding variable in the device data environment.

Cross References

- **target** construct, see Section 2.10.2 on page 95
- `omp_target_disassociate_ptr` routine, see Section 3.5.6 on page 284
- **map** clause, see Section 2.15.5.1 on page 213.
- `omp_target_alloc` routine, see Section 3.5.1 on page 279

3.5.7 `omp_target_disassociate_ptr`

Summary

The `omp_target_disassociate_ptr` removes the associated pointer for a given device from a host pointer.

Format

```
int omp_target_disassociate_ptr(void * ptr, int device_num);
```

Constraints on Arguments

When called from within a **target** region the effect of this routine is unspecified. The *device_num* must be greater than or equal to zero and less than the result of **omp_get_num_devices()** or the result of a call to **omp_get_initial_device()**.

Effect

The **omp_target_disassociate_ptr** removes the associated device data on device *device_num* from the presence table for host pointer *ptr*. A call to this routine on a pointer that is not **NULL** and does not have associated data on the given device results in unspecified behavior. The reference count of the mapping is reduced to zero, regardless of its current value.

After a call to **omp_target_disassociate_ptr**, the contents of the device buffer are invalidated.

Cross References

- **target** construct, see Section [2.10.2](#) on page [95](#)
- **omp_target_associate_ptr** routine, see Section [3.5.6](#) on page [284](#)

▲ C / C++ ▲

Environment Variables

This chapter describes the OpenMP environment variables that specify the settings of the ICVs that affect the execution of OpenMP programs (see Section 2.3 on page 35). The names of the environment variables must be upper case. The values assigned to the environment variables are case insensitive and may have leading and trailing white space. Modifications to the environment variables after the program has started, even if modified by the program itself, are ignored by the OpenMP implementation. However, the settings of some of the ICVs can be modified during the execution of the OpenMP program by the use of the appropriate directive clauses or OpenMP API routines.

The environment variables are as follows:

- **OMP_SCHEDULE** sets the *run-sched-var* ICV that specifies the runtime schedule type and chunk size. It can be set to any of the valid OpenMP schedule types.
- **OMP_NUM_THREADS** sets the *nthreads-var* ICV that specifies the number of threads to use for parallel regions.
- **OMP_DYNAMIC** sets the *dyn-var* ICV that specifies the dynamic adjustment of threads to use for **parallel** regions.
- **OMP_PROC_BIND** sets the *bind-var* ICV that controls the OpenMP thread affinity policy.
- **OMP_PLACES** sets the *place-partition-var* ICV that defines the OpenMP places that are available to the execution environment.
- **OMP_NESTED** sets the *nest-var* ICV that enables or disables nested parallelism.
- **OMP_STACKSIZE** sets the *stacksize-var* ICV that specifies the size of the stack for threads created by the OpenMP implementation.
- **OMP_WAIT_POLICY** sets the *wait-policy-var* ICV that controls the desired behavior of waiting threads.
- **OMP_MAX_ACTIVE_LEVELS** sets the *max-active-levels-var* ICV that controls the maximum number of nested active **parallel** regions.

- **OMP_THREAD_LIMIT** sets the *thread-limit-var* ICV that controls the maximum number of threads participating in a contention group.
- **OMP_CANCELLATION** sets the *cancel-var* ICV that enables or disables cancellation.
- **OMP_DISPLAY_ENV** instructs the runtime to display the OpenMP version number and the initial values of the ICVs, once, during initialization of the runtime.
- **OMP_DEFAULT_DEVICE** sets the *default-device-var* ICV that controls the default device number.
- **OMP_MAX_TASK_PRIORITY** sets the *max-task-priority-var* ICV that specifies the maximum value that can be specified in the **priority** clause of the **task** construct. A value of zero (default value) specifies that the runtime will ignore the **priority** clause.

The examples in this chapter only demonstrate how these variables might be set in Unix C shell (csh) environments. In Korn shell (ksh) and DOS environments the actions are similar, as follows:

- csh:

```
setenv OMP_SCHEDULE "dynamic"
```

- ksh:

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

- DOS:

```
set OMP_SCHEDULE=dynamic
```

1 4.1 OMP_SCHEDULE

2 The **OMP_SCHEDULE** environment variable controls the schedule type and chunk size of all loop
3 directives that have the schedule type **runtime**, by setting the value of the *run-sched-var* ICV.

4 The value of this environment variable takes the form:

5 *type*[, *chunk*]

6 where

- 7 • *type* is one of **static**, **dynamic**, **guided**, or **auto**
- 8 • *chunk* is an optional positive integer that specifies the chunk size

9 If *chunk* is present, there may be white space on either side of the “,”. See Section 2.7.1 on
10 page 55 for a detailed description of the schedule types.

11 The behavior of the program is implementation defined if the value of **OMP_SCHEDULE** does not
12 conform to the above format.

13 Implementation specific schedules cannot be specified in **OMP_SCHEDULE**. They can only be
14 specified by calling **omp_set_schedule**, described in Section 3.2.12 on page 240.

15 Example:

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

16 Cross References

- 17 • *run-sched-var* ICV, see Section 2.3 on page 35.
- 18 • Loop construct, see Section 2.7.1 on page 55.
- 19 • Parallel loop construct, see Section 2.11.1 on page 121.
- 20 • **omp_set_schedule** routine, see Section 3.2.12 on page 240.
- 21 • **omp_get_schedule** routine, see Section 3.2.13 on page 242.

1 4.2 OMP_NUM_THREADS

2 The **OMP_NUM_THREADS** environment variable sets the number of threads to use for **parallel**
3 regions by setting the initial value of the *nthreads-var* ICV. See Section 2.3 on page 35 for a
4 comprehensive set of rules about the interaction between the **OMP_NUM_THREADS** environment
5 variable, the **num_threads** clause, the **omp_set_num_threads** library routine and dynamic
6 adjustment of threads, and Section 2.5.1 on page 49 for a complete algorithm that describes how the
7 number of threads for a **parallel** region is determined.

8 The value of this environment variable must be a list of positive integer values. The values of the
9 list set the number of threads to use for **parallel** regions at the corresponding nested levels.

10 The behavior of the program is implementation defined if any value of the list specified in the
11 **OMP_NUM_THREADS** environment variable leads to a number of threads which is greater than an
12 implementation can support, or if any value is not a positive integer.

13 Example:

```
setenv OMP_NUM_THREADS 4,3,2
```

14 Cross References

- 15 • *nthreads-var* ICV, see Section 2.3 on page 35.
- 16 • **num_threads** clause, Section 2.5 on page 45.
- 17 • **omp_set_num_threads** routine, see Section 3.2.1 on page 228.
- 18 • **omp_get_num_threads** routine, see Section 3.2.2 on page 229.
- 19 • **omp_get_max_threads** routine, see Section 3.2.3 on page 230.
- 20 • **omp_get_team_size** routine, see Section 3.2.19 on page 248.

1 4.3 OMP_DYNAMIC

2 The **OMP_DYNAMIC** environment variable controls dynamic adjustment of the number of threads
3 to use for executing **parallel** regions by setting the initial value of the *dyn-var* ICV. The value of
4 this environment variable must be **true** or **false**. If the environment variable is set to **true**, the
5 OpenMP implementation may adjust the number of threads to use for executing **parallel**
6 regions in order to optimize the use of system resources. If the environment variable is set to
7 **false**, the dynamic adjustment of the number of threads is disabled. The behavior of the program
8 is implementation defined if the value of **OMP_DYNAMIC** is neither **true** nor **false**.

9 Example:

```
setenv OMP_DYNAMIC true
```

10 Cross References

- 11 • *dyn-var* ICV, see Section 2.3 on page 35.
- 12 • **omp_set_dynamic** routine, see Section 3.2.7 on page 234.
- 13 • **omp_get_dynamic** routine, see Section 3.2.8 on page 236.

14 4.4 OMP_PROC_BIND

15 The **OMP_PROC_BIND** environment variable sets the initial value of the *bind-var* ICV. The value
16 of this environment variable is either **true**, **false**, or a comma separated list of **master**,
17 **close**, or **spread**. The values of the list set the thread affinity policy to be used for parallel
18 regions at the corresponding nested level.

19 If the environment variable is set to **false**, the execution environment may move OpenMP threads
20 between OpenMP places, thread affinity is disabled, and **proc_bind** clauses on **parallel**
21 constructs are ignored.

22 Otherwise, the execution environment should not move OpenMP threads between OpenMP places,
23 thread affinity is enabled, and the initial thread is bound to the first place in the OpenMP place list.

24 The behavior of the program is implementation defined if any of the values in the
25 **OMP_PROC_BIND** environment variable is not **true**, **false**, or a comma separated list of
26 **master**, **close**, or **spread**. The behavior is also implementation defined if an initial thread
27 cannot be bound to the first place in the OpenMP place list.

28 Example:

```
setenv OMP_PROC_BIND false
setenv OMP_PROC_BIND "spread, spread, close"
```

Cross References

- *bind-var* ICV, see Section 2.3 on page 35.
- `proc_bind` clause, see Section 2.5.2 on page 51.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 251.

4.5 OMP_PLACES

A list of places can be specified in the **OMP_PLACES** environment variable. The *place-partition-var* ICV obtains its initial value from the **OMP_PLACES** value, and makes the list available to the execution environment. The value of **OMP_PLACES** can be one of two types of values: either an abstract name describing a set of places or an explicit list of places described by non-negative numbers.

The **OMP_PLACES** environment variable can be defined using an explicit ordered list of comma-separated places. A place is defined by an unordered set of comma-separated non-negative numbers enclosed by braces. The meaning of the numbers and how the numbering is done are implementation defined. Generally, the numbers represent the smallest unit of execution exposed by the execution environment, typically a hardware thread.

Intervals may also be used to define places. Intervals can be specified using the *<lower-bound> : <length> : <stride>* notation to represent the following list of numbers: “*<lower-bound>*, *<lower-bound> + <stride>*, ..., *<lower-bound> + (<length>- 1)*<stride>*.” When *<stride>* is omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences of places.

An exclusion operator “!” can also be used to exclude the number or place immediately following the operator.

Alternatively, the abstract names listed in Table 4.1 should be understood by the execution and runtime environment. The precise definitions of the abstract names are implementation defined. An implementation may also add abstract names as appropriate for the target platform.

The abstract name may be appended by a positive number in parentheses to denote the length of the place list to be created, that is *abstract_name(num-places)*. When requesting fewer places than available on the system, the determination of which resources of type *abstract_name* are to be

included in the place list is implementation defined. When requesting more resources than available, the length of the place list is implementation defined.

Abstract Name	Meaning
threads	Each place corresponds to a single hardware thread on the target machine.
cores	Each place corresponds to a single core (having one or more hardware threads) on the target machine.
sockets	Each place corresponds to a single socket (consisting of one or more cores) on the target machine.

Table 4.1.: Defined Abstract Names for **OMP_PLACES**

The behavior of the program is implementation defined when the execution environment cannot map a numerical value (either explicitly defined or implicitly derived from an interval) within the **OMP_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor. The behavior is also implementation defined when the **OMP_PLACES** environment variable is defined using an abstract name.

The following grammar describes the values of *list* accepted for the **OMP_PLACES** environment variable.

$$\begin{aligned}
\langle \text{list} \rangle & \models \langle \text{p-list} \rangle \mid \langle \text{aname-list} \rangle \mid \langle \text{p-list} \rangle, \langle \text{list} \rangle \mid \langle \text{aname-list} \rangle, \langle \text{list} \rangle \\
\langle \text{p-list} \rangle & \models \langle \text{p-interval} \rangle \mid \langle \text{p-list} \rangle, \langle \text{p-interval} \rangle \\
\langle \text{p-interval} \rangle & \models \langle \text{place} \rangle : \langle \text{len} \rangle : \langle \text{stride} \rangle \mid \langle \text{place} \rangle : \langle \text{len} \rangle \mid \langle \text{place} \rangle \mid ! \langle \text{place} \rangle \\
\langle \text{place} \rangle & \models \{ \langle \text{res-list} \rangle \} \\
\langle \text{res-list} \rangle & \models \langle \text{res-interval} \rangle \mid \langle \text{res-list} \rangle, \langle \text{res-interval} \rangle \\
\langle \text{res-interval} \rangle & \models \langle \text{res} \rangle : \langle \text{num-places} \rangle : \langle \text{stride} \rangle \mid \langle \text{res} \rangle : \langle \text{num-places} \rangle \mid \langle \text{res} \rangle \mid ! \langle \text{res} \rangle \\
\langle \text{aname-list} \rangle & \models \langle \text{aname} \rangle \mid \langle \text{aname-list} \rangle, \langle \text{aname} \rangle \\
\langle \text{aname} \rangle & \models \langle \text{word} \rangle (\langle \text{num-places} \rangle) \mid \langle \text{word} \rangle \\
\langle \text{word} \rangle & \models \text{sockets} \mid \text{cores} \mid \text{threads} \\
\langle \text{res} \rangle & \models \text{non-negative integer} \\
\langle \text{num-places} \rangle & \models \text{positive integer (not null)} \\
\langle \text{stride} \rangle & \models \text{integer} \\
\langle \text{len} \rangle & \models \text{positive integer (not null)}
\end{aligned}$$

Example:

```
setenv OMP_PLACES threads
setenv OMP_PLACES "threads(4)"
setenv OMP_PLACES "{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
setenv OMP_PLACES "{0:4},{4:4},{8:4},{12:4}"
setenv OMP_PLACES "{0:4}:4:4"
```

where each of the last three definitions corresponds to the same 4 places including the smallest units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11, and 12 to 15.

Cross References

- *place-partition-var*, Section 2.3 on page 35.
- Controlling OpenMP thread affinity, Section 2.5.2 on page 51.
- `omp_get_num_places` routine, see Section 3.2.23 on page 253.
- `omp_get_place_num_procs` routine, see Section 3.2.24 on page 254.
- `omp_get_place_proc_ids` routine, see Section 3.2.25 on page 255.
- `omp_get_place_num` routine, see Section 3.2.26 on page 256.
- `omp_get_partition_num_places` routine, see Section 3.2.27 on page 257.
- `omp_get_partition_place_nums` routine, see Section 3.2.28 on page 258.

4.6 OMP_NESTED

The `OMP_NESTED` environment variable controls nested parallelism by setting the initial value of the *nest-var* ICV. The value of this environment variable must be **true** or **false**. If the environment variable is set to **true**, nested parallelism is enabled; if set to **false**, nested parallelism is disabled. The behavior of the program is implementation defined if the value of `OMP_NESTED` is neither **true** nor **false**.

Example:

```
setenv OMP_NESTED false
```


Cross References

- *nest-var* ICV, see Section 2.3 on page 35.
- `omp_set_nested` routine, see Section 3.2.10 on page 237.
- `omp_get_team_size` routine, see Section 3.2.19 on page 248.

4.7 OMP_STACKSIZE

The **OMP_STACKSIZE** environment variable controls the size of the stack for threads created by the OpenMP implementation, by setting the value of the *stacksize-var* ICV. The environment variable does not control the size of the stack for an initial thread.

The value of this environment variable takes the form:

size | *size***B** | *size***K** | *size***M** | *size***G**

where:

- *size* is a positive integer that specifies the size of the stack for threads that are created by the OpenMP implementation.
- **B**, **K**, **M**, and **G** are letters that specify whether the given size is in Bytes, Kilobytes (1024 Bytes), Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If one of these letters is present, there may be white space between *size* and the letter.

If only *size* is specified and none of **B**, **K**, **M**, or **G** is specified, then *size* is assumed to be in Kilobytes.

The behavior of the program is implementation defined if **OMP_STACKSIZE** does not conform to the above format, or if the implementation cannot provide a stack with the requested size.

Examples:

```
setenv OMP_STACKSIZE 2000500B
setenv OMP_STACKSIZE "3000 k "
setenv OMP_STACKSIZE 10M
setenv OMP_STACKSIZE " 10 M "
setenv OMP_STACKSIZE "20 m "
setenv OMP_STACKSIZE " 1G"
setenv OMP_STACKSIZE 20000
```

Cross References

- *stacksize-var* ICV, see Section 2.3 on page 35.

4.8 OMP_WAIT_POLICY

The **OMP_WAIT_POLICY** environment variable 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 value of this environment variable takes the form:

ACTIVE | **PASSIVE**

The **ACTIVE** value specifies that waiting threads should mostly be active, consuming processor cycles, while waiting. An OpenMP implementation may, for example, make waiting threads spin.

The **PASSIVE** value specifies that waiting threads should mostly be passive, not consuming processor cycles, while waiting. For example, an OpenMP implementation may make waiting threads yield the processor to other threads or go to sleep.

The details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined.

Examples:

```
setenv OMP_WAIT_POLICY ACTIVE
setenv OMP_WAIT_POLICY active
setenv OMP_WAIT_POLICY PASSIVE
setenv OMP_WAIT_POLICY passive
```

Cross References

- *wait-policy-var* ICV, see Section 2.3 on page 35.

1 4.9 OMP_MAX_ACTIVE_LEVELS

2 The **OMP_MAX_ACTIVE_LEVELS** environment variable controls the maximum number of nested
3 active **parallel** regions by setting the initial value of the *max-active-levels-var* ICV.

4 The value of this environment variable must be a non-negative integer. The behavior of the
5 program is implementation defined if the requested value of **OMP_MAX_ACTIVE_LEVELS** is
6 greater than the maximum number of nested active parallel levels an implementation can support,
7 or if the value is not a non-negative integer.

8 Cross References

- 9 • *max-active-levels-var* ICV, see Section 2.3 on page 35.
- 10 • **omp_set_max_active_levels** routine, see Section 3.2.15 on page 243.
- 11 • **omp_get_max_active_levels** routine, see Section 3.2.16 on page 245.

12 4.10 OMP_THREAD_LIMIT

13 The **OMP_THREAD_LIMIT** environment variable sets the maximum number of OpenMP threads
14 to use in a contention group by setting the *thread-limit-var* ICV.

15 The value of this environment variable must be a positive integer. The behavior of the program is
16 implementation defined if the requested value of **OMP_THREAD_LIMIT** is greater than the
17 number of threads an implementation can support, or if the value is not a positive integer.

18 Cross References

- 19 • *thread-limit-var* ICV, see Section 2.3 on page 35.
- 20 • **omp_get_thread_limit** routine, see Section 3.2.14 on page 243.

21 4.11 OMP_CANCELLATION

22 The **OMP_CANCELLATION** environment variable sets the initial value of the *cancel-var* ICV.

The value of this environment variable must be **true** or **false**. If set to **true**, the effects of the **cancel** construct and of cancellation points are enabled and cancellation is activated. If set to **false**, cancellation is disabled and the **cancel** construct and cancellation points are effectively ignored.

Cross References

- *cancel-var*, see Section 2.3.1 on page 36.
- **cancel** construct, see Section 2.14.1 on page 168.
- **cancellation point** construct, see Section 2.14.2 on page 172.
- **omp_get_cancellation** routine, see Section 3.2.9 on page 237.

4.12 OMP_DISPLAY_ENV

The **OMP_DISPLAY_ENV** environment variable instructs the runtime to display the OpenMP version number and the value of the ICVs associated with the environment variables described in Chapter 4, as *name = value* pairs. The runtime displays this information once, after processing the environment variables and before any user calls to change the ICV values by runtime routines defined in Chapter 3.

The value of the **OMP_DISPLAY_ENV** environment variable may be set to one of these values:

TRUE | FALSE | VERBOSE

The **TRUE** value instructs the runtime to display the OpenMP version number defined by the **_OPENMP** version macro (or the **openmp_version** Fortran parameter) value and the initial ICV values for the environment variables listed in Chapter 4. The **VERBOSE** value indicates that the runtime may also display the values of runtime variables that may be modified by vendor-specific environment variables. The runtime does not display any information when the **OMP_DISPLAY_ENV** environment variable is **FALSE** or undefined. For all values of the environment variable other than **TRUE**, **FALSE**, and **VERBOSE**, the displayed information is unspecified.

The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the **_OPENMP** version macro (or the **openmp_version** Fortran parameter) value and ICV values, in the format *NAME '=' VALUE*. *NAME* corresponds to the macro or environment variable name, optionally prepended by a bracketed *device-type*. *VALUE* corresponds to the value of the macro or ICV associated with this environment variable. Values should be enclosed in single quotes. The display is terminated with "OPENMP DISPLAY ENVIRONMENT END".

Example:

```
% setenv OMP_DISPLAY_ENV TRUE
```

The above example causes an OpenMP implementation to generate output of the following form:

```
OPENMP DISPLAY ENVIRONMENT BEGIN
_OPENMP='201307'
[host] OMP_SCHEDULE='GUIDED,4'
[host] OMP_NUM_THREADS='4,3,2'
[device] OMP_NUM_THREADS='2'
[host,device] OMP_DYNAMIC='TRUE'
[host] OMP_PLACES='0:4,4:4,8:4,12:4'
...
OPENMP DISPLAY ENVIRONMENT END
```

4.13 OMP_DEFAULT_DEVICE

The **OMP_DEFAULT_DEVICE** environment variable sets the device number to use in device constructs by setting the initial value of the *default-device-var* ICV.

The value of this environment variable must be a non-negative integer value.

Cross References

- *default-device-var* ICV, see Section 2.3 on page 35.
- device constructs, Section 2.10 on page 93.

1 4.14 OMP_MAX_TASK_PRIORITY

2 The **OMP_TASK_MAX_PRIORITY** environment variable controls the use of task priorities by
3 setting the initial value of the *max-task-priority-var* ICV. The value of this environment variable
4 must be a non-negative integer. If the environment variable is set to zero, the OpenMP
5 implementation will ignore the **priority** clause in any **task** directive. Setting the variable to a
6 non-zero integer value will inform the implementation that values specified in **priority** clauses
7 can be capped at this maximum value. The default value for this environment variable is zero.

8 Example:

```
% setenv OMP_MAX_TASK_PRIORITY 20
```

9 Cross References

- 10 • *max-task-priority-var* ICV, see Section [2.3](#) on page [35](#).
- 11 • Tasking Constructs, see Section [2.9](#) on page [81](#).
- 12 • **omp_get_max_task_priority** routine, see Section [3.2.36](#) on page [265](#).

2 **Stubs for Runtime Library Routines**

3 This section provides stubs for the runtime library routines defined in the OpenMP API. The stubs
4 are provided to enable portability to platforms that do not support the OpenMP API. On these
5 platforms, OpenMP programs must be linked with a library containing these stub routines. The stub
6 routines assume that the directives in the OpenMP program are ignored. As such, they emulate
7 serial semantics.

8 Note that the lock variable that appears in the lock routines must be accessed exclusively through
9 these routines. It should not be initialized or otherwise modified in the user program.

10 In an actual implementation the lock variable might be used to hold the address of an allocated
11 memory block, but here it is used to hold an integer value. Users should not make assumptions
12 about mechanisms used by OpenMP implementations to implement locks based on the scheme
13 used by the stub procedures.

▼────────────────── Fortran ───────────────────▼

14 **Note** – In order to be able to compile the Fortran stubs file, the include file **omp_lib.h** was split
15 into two files: **omp_lib_kinds.h** and **omp_lib.h** and the **omp_lib_kinds.h** file included
16 where needed. There is no requirement for the implementation to provide separate files.

▲────────────────── Fortran ───────────────────▲

1 A.1 C/C++ Stub Routines

```
2      #include <stdio.h>
3      #include <stdlib.h>
4      #include "omp.h"
5
6      void omp_set_num_threads(int num_threads)
7      {
8      }
9
10     int omp_get_num_threads(void)
11     {
12         return 1;
13     }
14
15     int omp_get_max_threads(void)
16     {
17         return 1;
18     }
19
20     int omp_get_thread_num(void)
21     {
22         return 0;
23     }
24
25     int omp_get_num_procs(void)
26     {
27         return 1;
28     }
29
30     int omp_in_parallel(void)
31     {
32         return 0;
33     }
34
35     void omp_set_dynamic(int dynamic_threads)
36     {
37     }
38
39     int omp_get_dynamic(void)
40     {
41         return 0;
42     }
43
44     int omp_get_cancellation(void)
45     {
46         return 0;
```



```

1      }
2
3      void omp_set_nested(int nested)
4      {
5      }
6
7      int omp_get_nested(void)
8      {
9          return 0;
10     }
11
12     void omp_set_schedule(omp_sched_t kind, int chunk_size)
13     {
14     }
15
16     void omp_get_schedule(omp_sched_t *kind, int *chunk_size)
17     {
18         *kind = omp_sched_static;
19         *chunk_size = 0;
20     }
21
22     int omp_get_thread_limit(void)
23     {
24         return 1;
25     }
26
27     void omp_set_max_active_levels(int max_active_levels)
28     {
29     }
30
31     int omp_get_max_active_levels(void)
32     {
33         return 0;
34     }
35
36     int omp_get_level(void)
37     {
38         return 0;
39     }
40
41     int omp_get_ancestor_thread_num(int level)
42     {
43         if (level == 0)
44         {
45             return 0;
46         }
47         else

```

```

1      {
2          return -1;
3      }
4  }
5
6  int omp_get_team_size(int level)
7  {
8      if (level == 0)
9      {
10         return 1;
11     }
12     else
13     {
14         return -1;
15     }
16 }
17
18 int omp_get_active_level(void)
19 {
20     return 0;
21 }
22
23 int omp_in_final(void)
24 {
25     return 1;
26 }
27
28 omp_proc_bind_t omp_get_proc_bind(void)
29 {
30     return omp_proc_bind_false;
31 }
32
33 int omp_get_num_places(void)
34 {
35     return 0;
36 }
37
38 int omp_get_place_num_procs(int place_num)
39 {
40     return 0;
41 }
42
43 void omp_get_place_proc_ids(int place_num)
44 {
45 }
46
47 int omp_get_place_num(void)

```

```

1      {
2          return -1;
3      }
4
5      int omp_get_partition_num_places(void)
6      {
7          return 0;
8      }
9
10     void omp_get_partition_place_nums(int *place_nums)
11     {
12     }
13
14     void omp_set_default_device(int device_num)
15     {
16     }
17
18     int omp_get_default_device(void)
19     {
20         return 0;
21     }
22
23     int omp_get_num_devices(void)
24     {
25         return 0;
26     }
27
28     int omp_get_num_teams(void)
29     {
30         return 1;
31     }
32
33     int omp_get_team_num(void)
34     {
35         return 0;
36     }
37
38     int omp_is_initial_device(void)
39     {
40         return 1;
41     }
42
43     int omp_get_initial_device(void)
44     {
45         return -10;
46     }
47

```

```

1      int omp_get_max_task_priority(void)
2      {
3          return 0;
4      }
5
6      struct __omp_lock
7      {
8          int lock;
9      };
10
11     enum { UNLOCKED = -1, INIT, LOCKED };
12
13     void omp_init_lock(omp_lock_t *arg)
14     {
15         struct __omp_lock *lock = (struct __omp_lock *)arg;
16         lock->lock = UNLOCKED;
17     }
18
19     void omp_init_lock_with_hint(omp_lock_t *arg, omp_lock_hint_t hint)
20     {
21         omp_init_lock(arg);
22     }
23
24     void omp_destroy_lock(omp_lock_t *arg)
25     {
26         struct __omp_lock *lock = (struct __omp_lock *)arg;
27         lock->lock = INIT;
28     }
29
30     void omp_set_lock(omp_lock_t *arg)
31     {
32         struct __omp_lock *lock = (struct __omp_lock *)arg;
33         if (lock->lock == UNLOCKED)
34         {
35             lock->lock = LOCKED;
36         }
37         else if (lock->lock == LOCKED)
38         {
39             fprintf(stderr, "error: deadlock in using lock variable\n");
40             exit(1);
41         }
42
43         else
44         {
45             fprintf(stderr, "error: lock not initialized\n");
46             exit(1);
47         }

```

```

1      }
2
3      void omp_unset_lock(omp_lock_t *arg)
4      {
5          struct __omp_lock *lock = (struct __omp_lock *)arg;
6              if (lock->lock == LOCKED)
7              {
8                  lock->lock = UNLOCKED;
9              }
10             else if (lock->lock == UNLOCKED)
11             {
12                 fprintf(stderr, "error: lock not set\n");
13                 exit(1);
14             }
15             else
16             {
17                 fprintf(stderr, "error: lock not initialized\n");
18                 exit(1);
19             }
20     }
21
22     int omp_test_lock(omp_lock_t *arg)
23     {
24         struct __omp_lock *lock = (struct __omp_lock *)arg;
25         if (lock->lock == UNLOCKED)
26         {
27             lock->lock = LOCKED;
28             return 1;
29         }
30         else if (lock->lock == LOCKED)
31         {
32             return 0;
33         }
34         else
35         {
36             fprintf(stderr, "error: lock not initialized\n");
37             exit(1);
38         }
39     }
40
41     struct __omp_nest_lock
42     {
43         short owner;
44         short count;
45     };
46
47     enum { NOOWNER = -1, MASTER = 0 };

```

```

1
2 void omp_init_nest_lock(omp_nest_lock_t *arg)
3 {
4     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
5     nlock->owner = NOOWNER;
6     nlock->count = 0;
7 }
8
9 void omp_init_nest_lock_with_hint(omp_nest_lock_t *arg,
10                                  omp_lock_hint_t hint)
11 {
12     omp_init_nest_lock(arg);
13 }
14
15 void omp_destroy_nest_lock(omp_nest_lock_t *arg)
16 {
17     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
18     nlock->owner = NOOWNER;
19     nlock->count = UNLOCKED;
20 }
21
22 void omp_set_nest_lock(omp_nest_lock_t *arg)
23 {
24     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
25     if (nlock->owner == MASTER && nlock->count >= 1)
26     {
27         nlock->count++;
28     }
29     else if (nlock->owner == NOOWNER && nlock->count == 0)
30     {
31         nlock->owner = MASTER;
32         nlock->count = 1;
33     }
34     else
35     {
36         fprintf(stderr, "error: lock corrupted or not initialized\n");
37         exit(1);
38     }
39 }
40
41 void omp_unset_nest_lock(omp_nest_lock_t *arg)
42 {
43     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
44     if (nlock->owner == MASTER && nlock->count >= 1)
45     {
46         nlock->count--;
47         if (nlock->count == 0)

```

```

1         {
2             nlock->owner = NOOWNER;
3         }
4     }
5     else if (nlock->owner == NOOWNER && nlock->count == 0)
6     {
7         fprintf(stderr, "error: lock not set\n");
8         exit(1);
9     }
10    else
11    {
12        fprintf(stderr, "error: lock corrupted or not initialized\n");
13        exit(1);
14    }
15 }
16
17 int omp_test_nest_lock(omp_nest_lock_t *arg)
18 {
19     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
20     omp_set_nest_lock(arg);
21     return nlock->count;
22 }
23
24 double omp_get_wtime(void)
25 {
26     /* This function does not provide a working
27      * wallclock timer. Replace it with a version
28      * customized for the target machine.
29      */
30     return 0.0;
31 }
32
33 double omp_get_wtick(void)
34 {
35     /* This function does not provide a working
36      * clock tick function. Replace it with
37      * a version customized for the target machine.
38      */
39     return 365. * 86400.;
40 }
41
42 void * omp_target_alloc(size_t size, int device_num)
43 {
44     if(device_num != -10)
45         return NULL;
46     return malloc(size)
47 }

```

```

1
2 void omp_target_free(void * device_ptr, int device_num)
3 {
4     free(device_ptr);
5 }
6
7 int omp_target_is_present(void * ptr, size_t offset,
8                             int device_num)
9 {
10     return 1;
11 }
12
13 int omp_target_memcpy(void *dst, void *src, size_t length,
14                       size_t dst_offset, size_t src_offset,
15                       int dst_device, int src_device)
16 {
17     //only the default device is valid in a stub
18     if(dst_device != -10 || src_device != -10
19         || ! dst || ! src )
20         return EINVAL;
21     memcpy((char *)dst + dst_offset,
22           (char *)src + src_offset,
23           length);
24     return 0;
25 }
26
27 int omp_target_memcpy_rect(
28     void *dst, void *src,
29     size_t element_size,
30     int num_dims,
31     const size_t *volume,
32     const size_t *dst_offsets,
33     const size_t *src_offsets,
34     const size_t *dst_dimensions,
35     const size_t *src_dimensions,
36     int dst_device_num, int src_device_num)
37 {
38     int ret=0;
39     // Both null, return number of dimensions supported,
40     // this stub supports an arbitrary number
41     if(dst == NULL && src == NULL) return INT_MAX;
42
43     if(!volume || !dst_offsets || !src_offsets
44         || !dst_dimensions || !src_dimensions
45         || num_dims < 1 ) {
46         ret = EINVAL;
47         goto done;

```



```

1      }
2      if(num_dims == 1) {
3          ret = omp_target_memcpy(dst, src,
4                                  element_size * volume[0],
5                                  dst_offsets[0] * element_size,
6                                  src_offsets[0] * element_size,
7                                  dst_device_num, src_device_num);
8          if(ret) goto done;
9      } else {
10         size_t dst_slice_size = element_size;
11         size_t src_slice_size = element_size;
12         for(int i=1; i < num_dims; i++) {
13             dst_slice_size *= dst_dimensions[i];
14             src_slice_size *= src_dimensions[i];
15         }
16         size_t dst_off = dst_offsets[0] * dst_slice_size;
17         size_t src_off = src_offsets[0] * src_slice_size;
18         for(size_t i=0; i < volume[0]; i++) {
19             ret = omp_target_memcpy_rect((char *)dst + dst_off + dst_slice_size*i,
20                                          (char *)src + src_off + src_slice_size*i,
21                                          element_size,
22                                          num_dims - 1,
23                                          volume + 1,
24                                          dst_offsets + 1,
25                                          src_offsets + 1,
26                                          dst_dimensions + 1,
27                                          src_dimensions + 1,
28                                          dst_device_num,
29                                          src_device_num);
30             if(ret) goto done;
31         }
32     }
33 done:
34     return ret;
35 }
36
37 int omp_target_associate_ptr(void * host_ptr, void * device_ptr,
38                             size_t size, size_t device_offset,
39                             int device_num)
40 {
41     // No association is possible because all host pointers are considered present
42     return EINVAL;
43 }
44
45 int omp_target_disassociate_ptr(void * ptr, int device_num)
46 {
47     return EINVAL;

```

```
1  
2 }
```

1 A.2 Fortran Stub Routines

```
2      subroutine omp_set_num_threads(num_threads)
3          integer num_threads
4          return
5      end subroutine
6
7      integer function omp_get_num_threads()
8          omp_get_num_threads = 1
9          return
10     end function
11
12     integer function omp_get_max_threads()
13         omp_get_max_threads = 1
14         return
15     end function
16
17     integer function omp_get_thread_num()
18         omp_get_thread_num = 0
19         return
20     end function
21
22     integer function omp_get_num_procs()
23         omp_get_num_procs = 1
24         return
25     end function
26
27     logical function omp_in_parallel()
28         omp_in_parallel = .false.
29         return
30     end function
31
32     subroutine omp_set_dynamic(dynamic_threads)
33         logical dynamic_threads
34         return
35     end subroutine
36
37     logical function omp_get_dynamic()
38         omp_get_dynamic = .false.
39         return
40     end function
41
42     logical function omp_get_cancellation()
43         omp_get_cancellation = .false.
44         return
45     end function
46
```

```

1      subroutine omp_set_nested(nested)
2          logical nested
3          return
4      end subroutine
5
6      logical function omp_get_nested()
7          omp_get_nested = .false.
8          return
9      end function
10
11     subroutine omp_set_schedule(kind, chunk_size)
12         include 'omp_lib_kinds.h'
13         integer (kind=omp_sched_kind) kind
14         integer chunk_size
15         return
16     end subroutine
17
18     subroutine omp_get_schedule(kind, chunk_size)
19         include 'omp_lib_kinds.h'
20         integer (kind=omp_sched_kind) kind
21         integer chunk_size
22         kind = omp_sched_static
23         chunk_size = 0
24         return
25     end subroutine
26
27     integer function omp_get_thread_limit()
28         omp_get_thread_limit = 1
29         return
30     end function
31
32     subroutine omp_set_max_active_levels( level )
33         integer level
34     end subroutine
35
36     integer function omp_get_max_active_levels()
37         omp_get_max_active_levels = 0
38         return
39     end function
40
41     integer function omp_get_level()
42         omp_get_level = 0
43         return
44     end function
45
46     integer function omp_get_ancestor_thread_num( level )
47         integer level

```

```

1      if ( level .eq. 0 ) then
2          omp_get_ancestor_thread_num = 0
3      else
4          omp_get_ancestor_thread_num = -1
5      end if
6      return
7  end function
8
9  integer function omp_get_team_size( level )
10     integer level
11     if ( level .eq. 0 ) then
12         omp_get_team_size = 1
13     else
14         omp_get_team_size = -1
15     end if
16     return
17 end function
18
19 integer function omp_get_active_level()
20     omp_get_active_level = 0
21     return
22 end function
23
24 logical function omp_in_final()
25     omp_in_final = .true.
26     return
27 end function
28
29 function omp_get_proc_bind()
30     use omp_lib_kinds
31     integer (kind=omp_proc_bind_kind) omp_get_proc_bind
32     omp_get_proc_bind = omp_proc_bind_false
33 end function
34
35 integer function omp_get_num_places()
36     return 0
37 end function
38
39 integer function omp_get_place_num_procs(place_num)
40     integer place_num
41     return 0
42 end function
43
44 subroutine omp_get_place_proc_ids(place_num, ids)
45     integer place_num
46     integer ids(*)
47     return

```

```

1      end subroutine
2
3      integer function omp_get_place_num()
4          return -1
5      end function
6
7      integer function omp_get_partition_num_places()
8          return 0
9      end function
10
11     subroutine omp_get_partition_place_nums(place_nums)
12         integer place_nums(*)
13         return
14     end subroutine
15
16     subroutine omp_set_default_device(device_num)
17         integer device_num
18         return
19     end subroutine
20
21     integer function omp_get_default_device()
22         omp_get_default_device = 0
23         return
24     end function
25
26     integer function omp_get_num_devices()
27         omp_get_num_devices = 0
28         return
29     end function
30
31     integer function omp_get_num_teams()
32         omp_get_num_teams = 1
33         return
34     end function
35
36     integer function omp_get_team_num()
37         omp_get_team_num = 0
38         return
39     end function
40
41     logical function omp_is_initial_device()
42         omp_is_initial_device = .true.
43         return
44     end function
45
46     integer function omp_get_initial_device()
47         omp_get_initial_device = -10

```

```

1      return
2  end function
3
4  integer function omp_get_max_task_priority()
5      omp_get_max_task_priority = 0
6      return
7  end function
8
9  subroutine omp_init_lock(lock)
10     ! lock is 0 if the simple lock is not initialized
11     !      -1 if the simple lock is initialized but not set
12     !      1 if the simple lock is set
13     include 'omp_lib_kinds.h'
14     integer(kind=omp_lock_kind) lock
15
16     lock = -1
17     return
18 end subroutine
19
20 subroutine omp_init_lock_with_hint(lock, hint)
21     include 'omp_lib_kinds.h'
22     integer(kind=omp_lock_kind) lock
23     integer(kind=omp_lock_hint_kind) hint
24
25     call omp_init_lock(lock)
26     return
27 end subroutine
28
29 subroutine omp_destroy_lock(lock)
30     include 'omp_lib_kinds.h'
31     integer(kind=omp_lock_kind) lock
32
33     lock = 0
34     return
35 end subroutine
36
37 subroutine omp_set_lock(lock)
38     include 'omp_lib_kinds.h'
39     integer(kind=omp_lock_kind) lock
40
41     if (lock .eq. -1) then
42         lock = 1
43     elseif (lock .eq. 1) then
44         print *, 'error: deadlock in using lock variable'
45         stop
46     else
47         print *, 'error: lock not initialized'

```

```

1      stop
2  endif
3  return
4  end subroutine
5
6  subroutine omp_unset_lock(lock)
7      include 'omp_lib_kinds.h'
8      integer(kind=omp_lock_kind) lock
9
10     if (lock .eq. 1) then
11         lock = -1
12     elseif (lock .eq. -1) then
13         print *, 'error: lock not set'
14         stop
15     else
16         print *, 'error: lock not initialized'
17         stop
18     endif
19     return
20 end subroutine
21
22 logical function omp_test_lock(lock)
23     include 'omp_lib_kinds.h'
24     integer(kind=omp_lock_kind) lock
25
26     if (lock .eq. -1) then
27         lock = 1
28         omp_test_lock = .true.
29     elseif (lock .eq. 1) then
30         omp_test_lock = .false.
31     else
32         print *, 'error: lock not initialized'
33         stop
34     endif
35
36     return
37 end function
38
39 subroutine omp_init_nest_lock(nlock)
40     ! nlock is
41     ! 0 if the nestable lock is not initialized
42     ! -1 if the nestable lock is initialized but not set
43     ! 1 if the nestable lock is set
44     ! no use count is maintained
45     include 'omp_lib_kinds.h'
46     integer(kind=omp_nest_lock_kind) nlock
47

```



```

1      nlock = -1
2
3      return
4  end subroutine
5
6  subroutine omp_init_nest_lock_with_hint(nlock, hint)
7      include 'omp_lib_kinds.h'
8      integer(kind=omp_nest_lock_kind) nlock
9      integer(kind=omp_lock_hint_kind) hint
10
11     call omp_init_nest_lock(nlock)
12     return
13 end subroutine
14
15 subroutine omp_destroy_nest_lock(nlock)
16     include 'omp_lib_kinds.h'
17     integer(kind=omp_nest_lock_kind) nlock
18
19     nlock = 0
20
21     return
22 end subroutine
23
24 subroutine omp_set_nest_lock(nlock)
25     include 'omp_lib_kinds.h'
26     integer(kind=omp_nest_lock_kind) nlock
27
28     if (nlock .eq. -1) then
29         nlock = 1
30     elseif (nlock .eq. 0) then
31         print *, 'error: nested lock not initialized'
32         stop
33     else
34         print *, 'error: deadlock using nested lock variable'
35         stop
36     endif
37
38     return
39 end subroutine
40
41 subroutine omp_unset_nest_lock(nlock)
42     include 'omp_lib_kinds.h'
43     integer(kind=omp_nest_lock_kind) nlock
44
45     if (nlock .eq. 1) then
46         nlock = -1
47     elseif (nlock .eq. 0) then

```

```

1      print *, 'error: nested lock not initialized'
2      stop
3  else
4      print *, 'error: nested lock not set'
5      stop
6  endif
7
8      return
9  end subroutine
10
11 integer function omp_test_nest_lock(nlock)
12     include 'omp_lib_kinds.h'
13     integer(kind=omp_nest_lock_kind) nlock
14
15     if (nlock .eq. -1) then
16         nlock = 1
17         omp_test_nest_lock = 1
18     elseif (nlock .eq. 1) then
19         omp_test_nest_lock = 0
20     else
21         print *, 'error: nested lock not initialized'
22         stop
23     endif
24
25     return
26 end function
27
28 double precision function omp_get_wtime()
29     ! this function does not provide a working
30     ! wall clock timer. replace it with a version
31     ! customized for the target machine.
32
33     omp_get_wtime = 0.0d0
34
35     return
36 end function
37
38 double precision function omp_get_wtick()
39     ! this function does not provide a working
40     ! clock tick function. replace it with
41     ! a version customized for the target machine.
42     double precision one_year
43     parameter (one_year=365.d0*86400.d0)
44
45     omp_get_wtick = one_year
46
47     return

```

```
1      end function
```


1 APPENDIX B

2 Interface Declarations

3 This appendix gives examples of the C/C++ header file, the Fortran **include** file and Fortran
4 **module** that shall be provided by implementations as specified in Chapter 3. It also includes an
5 example of a Fortran 90 generic interface for a library routine. This is a non-normative section,
6 implementation files may differ.

1 B.1 Example of the omp.h Header File

```
2      #ifndef _OMP_H_DEF
3      #define _OMP_H_DEF
4
5      /*
6       * define the lock data types
7       */
8      typedef void *omp_lock_t;
9
10     typedef void *omp_nest_lock_t;
11
12     /*
13      * define the lock hints
14      */
15     typedef enum omp_lock_hint_t
16     {
17         omp_lock_hint_none = 0,
18         omp_lock_hint_uncontended = 1,
19         omp_lock_hint_contended = 2,
20         omp_lock_hint_nonspeculative = 4,
21         omp_lock_hint_speculative = 8
22     } /* , Add vendor specific constants for lock hints here,
23        starting from the most-significant bit. */
24     omp_lock_hint_t;
25
26     /*
27      * define the schedule kinds
28      */
29     typedef enum omp_sched_t
30     {
31         omp_sched_static = 1,
32         omp_sched_dynamic = 2,
33         omp_sched_guided = 3,
34         omp_sched_auto = 4
35     } /* , Add vendor specific schedule constants here */
36     omp_sched_t;
37
38     /*
39      * define the proc bind values
40      */
41     typedef enum omp_proc_bind_t
42     {
43         omp_proc_bind_false = 0,
44         omp_proc_bind_true = 1,
45         omp_proc_bind_master = 2,
46         omp_proc_bind_close = 3,
```

```

1      omp_proc_bind_spread = 4
2  } omp_proc_bind_t;
3
4  /*
5   * exported OpenMP functions
6   */
7  #ifdef __cplusplus
8  extern "C"
9  {
10 #endif
11
12     extern void omp_set_num_threads(int num_threads);
13     extern int omp_get_num_threads(void);
14     extern int omp_get_max_threads(void);
15     extern int omp_get_thread_num(void);
16     extern int omp_get_num_procs(void);
17     extern int omp_in_parallel(void);
18     extern void omp_set_dynamic(int dynamic_threads);
19     extern int omp_get_dynamic(void);
20     extern void omp_set_nested(int nested);
21     extern int omp_get_cancellation(void);
22     extern int omp_get_nested(void);
23     extern void omp_set_schedule(omp_sched_t kind, int chunk_size);
24     extern void omp_get_schedule(omp_sched_t *kind, int *chunk_size);
25     extern int omp_get_thread_limit(void);
26     extern void omp_set_max_active_levels(int max_active_levels);
27     extern int omp_get_max_active_levels(void);
28     extern int omp_get_level(void);
29     extern int omp_get_ancestor_thread_num(int level);
30     extern int omp_get_team_size(int level);
31     extern int omp_get_active_level(void);
32     extern int omp_in_final(void);
33     extern omp_proc_bind_t omp_get_proc_bind(void);
34     extern int omp_get_num_places(void);
35     extern int omp_get_place_num_procs(int place_num);
36     extern void omp_get_place_proc_ids(int place_num, int *ids);
37     extern int omp_get_place_num(void);
38     extern int omp_get_partition_num_places(void);
39     extern void omp_get_partition_place_nums(int *place_nums);
40     extern void omp_set_default_device(int device_num);
41     extern int omp_get_default_device(void);
42     extern int omp_get_num_devices(void);
43     extern int omp_get_num_teams(void);
44     extern int omp_get_team_num(void);
45     extern int omp_is_initial_device(void);
46     extern int omp_get_initial_device(void);
47     extern int omp_get_max_task_priority(void);

```

```

1
2     extern void omp_init_lock(omp_lock_t *lock);
3     extern void omp_init_lock_with_hint(omp_lock_t *lock,
4                                         omp_lock_hint_t hint);
5     extern void omp_destroy_lock(omp_lock_t *lock);
6     extern void omp_set_lock(omp_lock_t *lock);
7     extern void omp_unset_lock(omp_lock_t *lock);
8     extern int omp_test_lock(omp_lock_t *lock);
9
10    extern void omp_init_nest_lock(omp_nest_lock_t *lock);
11    extern void omp_init_nest_lock_with_hint(omp_nest_lock_t *lock,
12                                              omp_lock_hint_t hint);
13    extern void omp_destroy_nest_lock(omp_nest_lock_t *lock);
14    extern void omp_set_nest_lock(omp_nest_lock_t *lock);
15    extern void omp_unset_nest_lock(omp_nest_lock_t *lock);
16    extern int omp_test_nest_lock(omp_nest_lock_t *lock);
17
18    extern double omp_get_wtime(void);
19    extern double omp_get_wtick(void);
20
21    extern void * omp_target_alloc(size_t size, int device_num);
22    extern void omp_target_free(void * device_ptr, int device_num);
23    extern int omp_target_is_present(void * ptr, size_t offset,
24                                    int device_num);
25    extern int omp_target_memcpy(void *dst, void *src, size_t length,
26                                size_t dst_offset, size_t src_offset,
27                                int dst_device_num, int src_device_num);
28    extern int omp_target_memcpy_rect(
29        void *dst, void *src,
30        size_t element_size,
31        int num_dims,
32        const size_t *volume,
33        const size_t *dst_offsets,
34        const size_t *src_offsets,
35        const size_t *dst_dimensions,
36        const size_t *src_dimensions,
37        int dst_device_num, int src_device_num);
38    extern int omp_target_associate_ptr(void * host_ptr,
39                                       void * device_ptr,
40                                       size_t size,
41                                       size_t device_offset,
42                                       int device_num);
43    extern int omp_target_disassociate_ptr(void * ptr,
44                                           int device_num);
45
46    #ifdef __cplusplus
47    }

```



```
1      #endif
2
3      #endif
```

1 B.2 Example of an Interface Declaration include 2 File

```
3      omp_lib_kinds.h:  
4      integer omp_lock_kind  
5          integer omp_nest_lock_kind  
6          integer omp_lock_hint_kind  
7      ! this selects an integer that is large enough to hold a 64 bit integer  
8          parameter ( omp_lock_kind = selected_int_kind( 10 ) )  
9          parameter ( omp_nest_lock_kind = selected_int_kind( 10 ) )  
10         parameter ( omp_lock_hint_kind = selected_int_kind( 10 ) )  
11  
12         integer omp_sched_kind  
13     ! this selects an integer that is large enough to hold a 32 bit integer  
14         parameter ( omp_sched_kind = selected_int_kind( 8 ) )  
15         integer ( omp_sched_kind ) omp_sched_static  
16         parameter ( omp_sched_static = 1 )  
17         integer ( omp_sched_kind ) omp_sched_dynamic  
18         parameter ( omp_sched_dynamic = 2 )  
19         integer ( omp_sched_kind ) omp_sched_guided  
20         parameter ( omp_sched_guided = 3 )  
21         integer ( omp_sched_kind ) omp_sched_auto  
22         parameter ( omp_sched_auto = 4 )  
23  
24         integer omp_proc_bind_kind  
25         parameter ( omp_proc_bind_kind = selected_int_kind( 8 ) )  
26         integer ( omp_proc_bind_kind ) omp_proc_bind_false  
27         parameter ( omp_proc_bind_false = 0 )  
28         integer ( omp_proc_bind_kind ) omp_proc_bind_true  
29         parameter ( omp_proc_bind_true = 1 )  
30         integer ( omp_proc_bind_kind ) omp_proc_bind_master  
31         parameter ( omp_proc_bind_master = 2 )  
32         integer ( omp_proc_bind_kind ) omp_proc_bind_close  
33         parameter ( omp_proc_bind_close = 3 )  
34         integer ( omp_proc_bind_kind ) omp_proc_bind_spread  
35         parameter ( omp_proc_bind_spread = 4 )  
36  
37         integer ( omp_lock_hint_kind ) omp_lock_hint_none  
38         parameter ( omp_lock_hint_none = 0 )  
39         integer ( omp_lock_hint_kind ) omp_lock_hint_uncontended  
40         parameter ( omp_lock_hint_uncontended = 1 )  
41         integer ( omp_lock_hint_kind ) omp_lock_hint_contended  
42         parameter ( omp_lock_hint_contended = 2 )  
43         integer ( omp_lock_hint_kind ) omp_lock_hint_nonspeculative  
44         parameter ( omp_lock_hint_nonspeculative = 4 )
```

```

1      integer ( omp_lock_hint_kind ) omp_lock_hint_speculative
2      parameter ( omp_lock_hint_speculative = 8 )

3      omp_lib.h:

4      ! default integer type assumed below
5      ! default logical type assumed below
6      ! OpenMP API v4.0

7
8      include 'omp_lib_kinds.h'
9      integer openmp_version
10     parameter ( openmp_version = 201307 )

11
12     external omp_set_num_threads
13     external omp_get_num_threads
14     integer omp_get_num_threads
15     external omp_get_max_threads
16     integer omp_get_max_threads
17     external omp_get_thread_num
18     integer omp_get_thread_num
19     external omp_get_num_procs
20     integer omp_get_num_procs
21     external omp_in_parallel
22     logical omp_in_parallel
23     external omp_set_dynamic
24     external omp_get_dynamic
25     logical omp_get_dynamic
26     external omp_get_cancellation
27     logical omp_get_cancellation
28     external omp_set_nested
29     external omp_get_nested
30     logical omp_get_nested
31     external omp_set_schedule
32     external omp_get_schedule
33     external omp_get_thread_limit
34     integer omp_get_thread_limit
35     external omp_set_max_active_levels
36     external omp_get_max_active_levels
37     integer omp_get_max_active_levels
38     external omp_get_level
39     integer omp_get_level
40     external omp_get_ancestor_thread_num
41     integer omp_get_ancestor_thread_num
42     external omp_get_team_size
43     integer omp_get_team_size
44     external omp_get_active_level
45     integer omp_get_active_level
46     external omp_set_default_device

```

```

1      external omp_get_default_device
2      integer omp_get_default_device
3      external omp_get_num_devices
4      integer omp_get_num_devices
5      external omp_get_num_teams
6      integer omp_get_num_teams
7      external omp_get_team_num
8      integer omp_get_team_num
9      external omp_is_initial_device
10     logical omp_is_initial_device
11     external omp_get_initial_device
12     integer omp_get_initial_device
13     external omp_get_max_task_priority
14     integer omp_get_max_task_priority
15
16     external omp_in_final
17     logical omp_in_final
18
19     integer ( omp_proc_bind_kind ) omp_get_proc_bind
20     external omp_get_proc_bind
21     integer omp_get_num_places
22     external omp_get_num_places
23     integer omp_get_place_num_procs
24     external omp_get_place_num_procs
25     external omp_get_place_proc_ids
26     integer omp_get_place_num
27     external omp_get_place_num
28     integer omp_get_partition_num_places
29     external omp_get_partition_num_places
30     external omp_get_partition_place_nums
31
32     external omp_init_lock
33     external omp_init_lock_with_hint
34     external omp_destroy_lock
35     external omp_set_lock
36     external omp_unset_lock
37     external omp_test_lock
38     logical omp_test_lock
39
40     external omp_init_nest_lock
41     external omp_init_nest_lock_with_hint
42     external omp_destroy_nest_lock
43     external omp_set_nest_lock
44     external omp_unset_nest_lock
45     external omp_test_nest_lock
46     integer omp_test_nest_lock
47

```

```
1      external omp_get_wtick
2      double precision omp_get_wtick
3      external omp_get_wtime
4      double precision omp_get_wtime
```

1 B.3 Example of a Fortran Interface Declaration 2 module

```
3      !      the "!" of this comment starts in column 1
4      !23456
5
6      module omp_lib_kinds
7          integer, parameter :: omp_lock_kind = selected_int_kind( 10 )
8          integer, parameter :: omp_nest_lock_kind = selected_int_kind( 10 )
9          integer, parameter :: omp_lock_hint_kind = selected_int_kind( 10 )
10         integer (kind=omp_lock_hint_kind), parameter ::
11         &    omp_lock_hint_none = 0
12         integer (kind=omp_lock_hint_kind), parameter ::
13         &    omp_lock_hint_uncontended = 1
14         integer (kind=omp_lock_hint_kind), parameter ::
15         &    omp_lock_hint_contended = 2
16         integer (kind=omp_lock_hint_kind), parameter ::
17         &    omp_lock_hint_nonspeculative = 4
18         integer (kind=omp_lock_hint_kind), parameter ::
19         &    omp_lock_hint_speculative = 8
20
21         integer, parameter :: omp_sched_kind = selected_int_kind( 8 )
22         integer(kind=omp_sched_kind), parameter ::
23         &    omp_sched_static = 1
24         integer(kind=omp_sched_kind), parameter ::
25         &    omp_sched_dynamic = 2
26         integer(kind=omp_sched_kind), parameter ::
27         &    omp_sched_guided = 3
28         integer(kind=omp_sched_kind), parameter ::
29         &    omp_sched_auto = 4
30
31         integer, parameter :: omp_proc_bind_kind = selected_int_kind( 8 )
32         integer (kind=omp_proc_bind_kind), parameter ::
33         &    omp_proc_bind_false = 0
34         integer (kind=omp_proc_bind_kind), parameter ::
35         &    omp_proc_bind_true = 1
36         integer (kind=omp_proc_bind_kind), parameter ::
37         &    omp_proc_bind_master = 2
38         integer (kind=omp_proc_bind_kind), parameter ::
39         &    omp_proc_bind_close = 3
40         integer (kind=omp_proc_bind_kind), parameter ::
41         &    omp_proc_bind_spread = 4
42         end module omp_lib_kinds
43
44         module omp_lib
45
```

```

1      use omp_lib_kinds
2
3      !
4      OpenMP API v4.0
5      integer, parameter :: openmp_version = 201307
6
7      interface
8
9          subroutine omp_set_num_threads (num_threads)
10             integer, intent(in) :: num_threads
11         end subroutine omp_set_num_threads
12
13         function omp_get_num_threads ()
14             integer :: omp_get_num_threads
15         end function omp_get_num_threads
16
17         function omp_get_max_threads ()
18             integer :: omp_get_max_threads
19         end function omp_get_max_threads
20
21         function omp_get_thread_num ()
22             integer :: omp_get_thread_num
23         end function omp_get_thread_num
24
25         function omp_get_num_procs ()
26             integer :: omp_get_num_procs
27         end function omp_get_num_procs
28
29         function omp_in_parallel ()
30             logical :: omp_in_parallel
31         end function omp_in_parallel
32
33         subroutine omp_set_dynamic (dynamic_threads)
34             logical, intent(in) :: dynamic_threads
35         end subroutine omp_set_dynamic
36
37         function omp_get_dynamic ()
38             logical :: omp_get_dynamic
39         end function omp_get_dynamic
40
41         function omp_get_cancellation ()
42             logical :: omp_get_cancellation
43         end function omp_get_cancellation
44
45         subroutine omp_set_nested (nested)
46             logical, intent(in) :: nested
47         end subroutine omp_set_nested

```

```

1      function omp_get_nested ()
2          logical :: omp_get_nested
3      end function omp_get_nested
4
5      subroutine omp_set_schedule (kind, chunk_size)
6          use omp_lib_kinds
7          integer(kind=omp_sched_kind), intent(in) :: kind
8          integer, intent(in) :: chunk_size
9      end subroutine omp_set_schedule
10
11     subroutine omp_get_schedule (kind, chunk_size)
12         use omp_lib_kinds
13         integer(kind=omp_sched_kind), intent(out) :: kind
14         integer, intent(out)::chunk_size
15     end subroutine omp_get_schedule
16
17     function omp_get_thread_limit()
18         integer :: omp_get_thread_limit
19     end function omp_get_thread_limit
20
21     subroutine omp_set_max_active_levels(max_levels)
22         integer, intent(in) :: max_levels
23     end subroutine omp_set_max_active_levels
24
25     function omp_get_max_active_levels()
26         integer :: omp_get_max_active_levels
27     end function omp_get_max_active_levels
28
29     function omp_get_level()
30         integer :: omp_get_level
31     end function omp_get_level
32
33     function omp_get_ancestor_thread_num(level)
34         integer, intent(in) :: level
35         integer :: omp_get_ancestor_thread_num
36     end function omp_get_ancestor_thread_num
37
38     function omp_get_team_size(level)
39         integer, intent(in) :: level
40         integer :: omp_get_team_size
41     end function omp_get_team_size
42
43     function omp_get_active_level()
44         integer :: omp_get_active_level
45     end function omp_get_active_level
46
47     function omp_in_final()

```



```

1      logical omp_in_final
2      end function omp_in_final
3
4      function omp_get_proc_bind( )
5          use omp_lib_kinds
6          integer (kind=omp_proc_bind_kind) omp_get_proc_bind
7          omp_get_proc_bind = omp_proc_bind_false
8      end function omp_get_proc_bind
9
10     function omp_get_num_places()
11         integer :: omp_get_num_places
12     end function omp_get_num_places
13
14     function omp_get_place_num_procs(place_num)
15         integer, intent(in) :: place_num
16         integer :: omp_get_place_num_procs
17     end function omp_get_place_num_procs
18
19     subroutine omp_get_place_proc_ids(place_num, ids)
20         integer, intent(in) :: place_num
21         integer, intent(out) :: ids(*)
22     end subroutine omp_get_place_proc_ids
23
24     function omp_get_place_num()
25         integer :: omp_get_place_num
26     end function omp_get_place_num
27
28     function omp_get_partition_num_places()
29         integer :: omp_get_partition_num_places
30     end function omp_get_partition_num_places
31
32     subroutine omp_get_partition_place_nums(place_nums)
33         integer, intent(out) :: place_nums(*)
34     end subroutine omp_get_partition_place_nums
35
36     subroutine omp_set_default_device (device_num)
37         integer :: device_num
38     end subroutine omp_set_default_device
39
40     function omp_get_default_device ( )
41         integer :: omp_get_default_device
42     end function omp_get_default_device
43
44     function omp_get_num_devices ( )
45         integer :: omp_get_num_devices
46     end function omp_get_num_devices
47

```

```

1      function omp_get_num_teams ()
2          integer :: omp_get_num_teams
3      end function omp_get_num_teams
4
5      function omp_get_team_num ()
6          integer :: omp_get_team_num
7      end function omp_get_team_num
8
9      function omp_is_initial_device ()
10         logical :: omp_is_initial_device
11     end function omp_is_initial_device
12
13     function omp_get_initial_device ()
14         integer :: omp_get_initial_device
15     end function omp_get_initial_device
16
17     function omp_get_max_task_priority ()
18         integer :: omp_get_max_task_priority
19     end function omp_get_max_task_priority
20
21     subroutine omp_init_lock (svar)
22         use omp_lib_kinds
23         integer (kind=omp_lock_kind), intent(out) :: svar
24     end subroutine omp_init_lock
25
26     subroutine omp_init_lock_with_hint (svar, hint)
27         use omp_lib_kinds
28         integer (kind=omp_lock_kind), intent(out) :: svar
29         integer (kind=omp_lock_hint_kind), intent(in) :: hint
30     end subroutine omp_init_lock_with_hint
31
32     subroutine omp_destroy_lock (svar)
33         use omp_lib_kinds
34         integer (kind=omp_lock_kind), intent(inout) :: svar
35     end subroutine omp_destroy_lock
36
37     subroutine omp_set_lock (svar)
38         use omp_lib_kinds
39         integer (kind=omp_lock_kind), intent(inout) :: svar
40     end subroutine omp_set_lock
41
42     subroutine omp_unset_lock (svar)
43         use omp_lib_kinds
44         integer (kind=omp_lock_kind), intent(inout) :: svar
45     end subroutine omp_unset_lock
46
47     function omp_test_lock (svar)

```

```

1      use omp_lib_kinds
2      logical :: omp_test_lock
3      integer (kind=omp_lock_kind), intent(inout) :: svar
4  end function omp_test_lock
5
6      subroutine omp_init_nest_lock (nvar)
7          use omp_lib_kinds
8          integer (kind=omp_nest_lock_kind), intent(out) :: nvar
9  end subroutine omp_init_nest_lock
10
11     subroutine omp_init_nest_lock_with_hint (nvar, hint)
12         use omp_lib_kinds
13         integer (kind=omp_nest_lock_kind), intent(out) :: nvar
14         integer (kind=omp_lock_hint_kind), intent(in) :: hint
15     end subroutine omp_init_nest_lock_with_hint
16
17     subroutine omp_destroy_nest_lock (nvar)
18         use omp_lib_kinds
19         integer (kind=omp_nest_lock_kind), intent(inout) :: nvar
20     end subroutine omp_destroy_nest_lock
21
22     subroutine omp_set_nest_lock (nvar)
23         use omp_lib_kinds
24         integer (kind=omp_nest_lock_kind), intent(inout) :: nvar
25     end subroutine omp_set_nest_lock
26
27     subroutine omp_unset_nest_lock (nvar)
28         use omp_lib_kinds
29         integer (kind=omp_nest_lock_kind), intent(inout) :: nvar
30     end subroutine omp_unset_nest_lock
31
32     function omp_test_nest_lock (nvar)
33         use omp_lib_kinds
34         integer :: omp_test_nest_lock
35         integer (kind=omp_nest_lock_kind), intent(inout) :: nvar
36     end function omp_test_nest_lock
37
38     function omp_get_wtick ()
39         double precision :: omp_get_wtick
40     end function omp_get_wtick
41
42     function omp_get_wtime ()
43         double precision :: omp_get_wtime
44     end function omp_get_wtime
45
46     end interface
47

```

```
1          end module omp_lib
```

1 B.4 Example of a Generic Interface for a Library 2 Routine

3 Any of the OpenMP runtime library routines that take an argument may be extended with a generic
4 interface so arguments of different **KIND** type can be accommodated.

5 The **OMP_SET_NUM_THREADS** interface could be specified in the **omp_lib** module as follows:

```
interface omp_set_num_threads

    subroutine omp_set_num_threads_4(num_threads)
        use omp_lib_kinds
        integer(4), intent(in) :: num_threads
    end subroutine omp_set_num_threads_4

    subroutine omp_set_num_threads_8(num_threads)
        use omp_lib_kinds
        integer(8), intent(in) :: num_threads
    end subroutine omp_set_num_threads_8

end interface omp_set_num_threads
```

OpenMP Implementation-Defined Behaviors

This appendix summarizes the behaviors that are described as implementation defined in this API. Each behavior is cross-referenced back to its description in the main specification. An implementation is required to define and document its behavior in these cases.

- **Processor:** a hardware unit that is implementation defined (see Section 1.2.1 on page 2).
- **Device:** an implementation defined logical execution engine (see Section 1.2.1 on page 2).
- **Memory model:** the minimum size at which a memory update may also read and write back adjacent variables that are part of another variable (as array or structure elements) is implementation defined but is no larger than required by the base language (see Section 1.4.1 on page 16).
- **Memory model:** Implementations are allowed to relax the ordering imposed by implicit flush operations when the result is only visible to programs using non-sequentially consistent atomic directives (see Section 1.4.4 on page 19).
- **Internal control variables:** the initial values of *dyn-var*, *nthreads-var*, *run-sched-var*, *def-sched-var*, *bind-var*, *stacksize-var*, *wait-policy-var*, *thread-limit-var*, *max-active-levels-var*, *place-partition-var*, and *default-device-var* are implementation defined (see Section 2.3.2 on page 37).
- **Dynamic adjustment of threads:** providing the ability to dynamically adjust the number of threads is implementation defined. Implementations are allowed to deliver fewer threads (but at least one) than indicated in Algorithm 2-1 even if dynamic adjustment is disabled (see Section 2.5.1 on page 49).
- **Thread affinity:** With $T \leq P$, when T does not divide P evenly, the assignment of the remaining $P - T * S$ places into subpartitions is implementation defined. With $T > P$, when P does not divide T evenly, the assignment of the remaining $T - P * S$ threads into places is implementation defined. The determination of whether the affinity request can be fulfilled is

implementation defined. If not, the number of threads in the team and their mapping to places become implementation defined (see Section 2.5.2 on page 51).

- **Loop directive:** the integer type (or kind, for Fortran) used to compute the iteration count of a collapsed loop is implementation defined. The effect of the **schedule(runtime)** clause when the *run-sched-var* ICV is set to **auto** is implementation defined. The *simd_width* used when a **simd** schedule modifier is specified is implementation defined. See Section 2.7.1 on page 55.
- **sections construct:** the method of scheduling the structured blocks among threads in the team is implementation defined (see Section 2.7.2 on page 63).
- **single construct:** the method of choosing a thread to execute the structured block is implementation defined (see Section 2.7.3 on page 66)
- **simd construct:** the integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is implementation defined. The number of iterations that are executed concurrently at any given time is implementation defined. If the **aligned** clause is not specified, the assumed alignment is implementation defined (see Section 2.8.1 on page 71).
- **declare simd construct:** if the **simdlen** clause is not specified, the number of concurrent arguments for the function is implementation defined. If the **aligned** clause is not specified, the assumed alignment is implementation defined (see Section 2.8.2 on page 75).
- The number of loop iterations assigned to a task created from a **taskloop** construct is implementation defined, unless the **grainsize** or **num_tasks** clauses are specified (see Section 2.9.2 on page 85).
- **is_device_ptr clause:** Support for pointers created outside of the OpenMP device data management routines is implementation defined (see Section 2.10.2 on page 95).
- **teams construct:** the number of teams that are created is implementation defined but less than or equal to the value of the **num_teams** clause if specified. The maximum number of threads participating in the contention group that each team initiates is implementation defined but less than or equal to the value of the **thread_limit** clause if specified (see Section 2.10.5 on page 106).
- If no **dist_schedule** clause is specified then the schedule for the **distribute** construct is implementation defined (see Section 2.10.6 on page 109).
- **critical construct:** the effect of using a **hint** construct is implementation defined (see Section 2.13.2 on page 146).
- **atomic construct:** a compliant implementation may enforce exclusive access between **atomic** regions that update different storage locations. The circumstances under which this occurs are implementation defined. If the storage location designated by *x* is not size-aligned (that is, if the byte alignment of *x* is not a multiple of the size of *x*), then the behavior of the atomic region is implementation defined (see Section 2.13.6 on page 151).

- 1 • **omp_set_num_threads routine**: if the argument is not a positive integer the behavior is
2 implementation defined (see Section 3.2.1 on page 228).
- 3 • **omp_set_schedule routine**: for implementation specific schedule types, the values and
4 associated meanings of the second argument are implementation defined. (see Section 3.2.12 on
5 page 240).
- 6 • **omp_set_max_active_levels routine**: when called from within any explicit parallel
7 region the binding thread set (and binding region, if required) for the
8 **omp_set_max_active_levels** region is implementation defined and the behavior is
9 implementation defined. If the argument is not a non-negative integer then the behavior is
10 implementation defined (see Section 3.2.15 on page 243).
- 11 • **omp_get_max_active_levels routine**: when called from within any explicit parallel
12 region the binding thread set (and binding region, if required) for the
13 **omp_get_max_active_levels** region is implementation defined (see Section 3.2.16 on
14 page 245).
- 15 • **omp_get_place_proc_ids routine**: the meaning of the nonnegative numerical identifiers
16 returned by the **omp_get_place_proc_ids** routine is implementation defined (see
17 Section 3.2.25 on page 255).
- 18 • **omp_init_lock_with_hint** and **omp_init_nest_lock_with_hint routines**: if
19 hints are stored with a lock variable, the effect of the hints on the locks are implementation
20 defined (see Section 3.3.2 on page 270).
- 21 • **omp_target_memcpy_rect routine**: the maximum number of dimensions supported is
22 implementation defined, but must be at least three (see Section 3.5.5 on page 283).
- 23 • **OMP_SCHEDULE environment variable**: if the value of the variable does not conform to the
24 specified format then the result is implementation defined (see Section 4.1 on page 289).
- 25 • **OMP_NUM_THREADS environment variable**: if any value of the list specified in the
26 **OMP_NUM_THREADS** environment variable leads to a number of threads that is greater than the
27 implementation can support, or if any value is not a positive integer, then the result is
28 implementation defined (see Section 4.2 on page 290).
- 29 • **OMP_PROC_BIND environment variable**: if the value is not **true**, **false**, or a comma
30 separated list of **master**, **close**, or **spread**, the behavior is implementation defined. The
31 behavior is also implementation defined if an initial thread cannot be bound to the first place in
32 the OpenMP place list (see Section 4.4 on page 291).
- 33 • **OMP_DYNAMIC environment variable**: if the value is neither **true** nor **false** the behavior is
34 implementation defined (see Section 4.3 on page 291).
- 35 • **OMP_NESTED environment variable**: if the value is neither **true** nor **false** the behavior is
36 implementation defined (see Section 4.6 on page 294).

- **OMP_STACKSIZE environment variable:** if the value does not conform to the specified format or the implementation cannot provide a stack of the specified size then the behavior is implementation defined (see Section 4.7 on page 295).
- **OMP_WAIT_POLICY environment variable:** the details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined (see Section 4.8 on page 296).
- **OMP_MAX_ACTIVE_LEVELS environment variable:** if the value is not a non-negative integer or is greater than the number of parallel levels an implementation can support then the behavior is implementation defined (see Section 4.9 on page 297).
- **OMP_THREAD_LIMIT environment variable:** if the requested value is greater than the number of threads an implementation can support, or if the value is not a positive integer, the behavior of the program is implementation defined (see Section 4.10 on page 297).
- **OMP_PLACES environment variable:** the meaning of the numbers specified in the environment variable and how the numbering is done are implementation defined. The precise definitions of the abstract names are implementation defined. An implementation may add implementation-defined abstract names as appropriate for the target platform. When creating a place list of n elements by appending the number n to an abstract name, the determination of which resources to include in the place list is implementation defined. When requesting more resources than available, the length of the place list is also implementation defined. The behavior of the program is implementation defined when the execution environment cannot map a numerical value (either explicitly defined or implicitly derived from an interval) within the **OMP_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor. The behavior is also implementation defined when the **OMP_PLACES** environment variable is defined using an abstract name (see Section 4.5 on page 292).
- **Thread affinity policy:** if the affinity request for a **parallel** construct cannot be fulfilled, the behavior of the thread affinity policy is implementation defined for that **parallel** construct.

Fortran

- **threadprivate directive:** if the conditions for values of data in the threadprivate objects of threads (other than an initial thread) to persist between two consecutive active parallel regions do not all hold, the allocation status of an allocatable variable in the second region is implementation defined (see Section 2.15.2 on page 179).
- **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 2.15.3.2 on page 187).
- **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 accommodated (see Section 3.1 on page 227).

APPENDIX D

Features History

This appendix summarizes the major changes between recent versions of the OpenMP API since version 2.5.

D.1 Version 4.0 to 4.1 Differences

- Support for several features of Fortran 2003 was added (see Section 1.6 on page 21 for features that are still not supported).
- A parameter was added to the **ordered** clause of the loop construct (see Section 2.7.1 on page 55) and clauses were added to the **ordered** construct (see Section 2.13.8 on page 163) to support doacross loops and use of the **simd** construct on loops with loop-carried backward dependences.
- The **linear** clause was added to the loop construct (see Section 2.7.1 on page 55).
- The **simdlen** clause was added to the **simd** construct (see Section 2.8.1 on page 71) to support specification of the exact number of iterations desired per SIMD chunk.
- The **priority** clause was added to the **task** construct (see Section 2.9.1 on page 81) to support hints that specify the relative execution priority of explicit tasks. The **omp_get_max_task_priority** routine was added to return the maximum supported priority value (see Section 3.2.36 on page 265) and the **OMP_MAX_TASK_PRIORITY** environment variable was added to control the maximum priority value allowed (see Section 4.14 on page 300).
- Taskloop constructs (see Section 2.9.2 on page 85 and Section 2.9.3 on page 89) were added to support nestable parallel loops that create OpenMP tasks.

- To support interaction with native device implementations, the **use_device_ptr** clause was added to the **target data** construct (see Section 2.10.1 on page 93) and the **is_device_ptr** clause was added to the **target** construct (see Section 2.10.2 on page 95).
- The **nowait** and **depend** clauses were added to the **target** construct (see Section 2.10.2 on page 95) to improve support for asynchronous execution of **target** regions.
- The **private**, **firstprivate** and **defaultmap** clauses were added to the **target** construct (see Section 2.10.2 on page 95).
- The **declare target** directive was extended to allow mapping of global variables to be deferred to specific device executions and to allow an *extended-list* to be specified in C/C++ (see Section 2.10.4 on page 102).
- To support unstructured data mapping for devices, the **target enter data** (see Section 2.10.10 on page 116) and **target exit data** (see Section 2.10.11 on page 118) constructs were added and the **map** clause (see Section 2.15.5.1 on page 213) was updated.
- To support a more complete set of device construct shortcuts, the **target parallel** (see Section 2.11.5 on page 126), target parallel loop (see Section 2.11.6 on page 127), target parallel loop SIMD (see Section 2.11.7 on page 129), and **target simd** (see Section 2.11.8 on page 130), combined constructs were added.
- The **if** clause was extended to take a *directive-name-modifier* that allows it to apply to combined constructs (see Section 2.12 on page 143).
- The **hint** clause was added to the **critical** construct (see Section 2.13.2 on page 146).
- The **source** and **sink** dependence types were added to the **depend** clause (see Section 2.13.9 on page 165) to support doacross loops.
- The implicit data-sharing attribute for scalar variables in **target** regions was changed to **firstprivate** (see Section 2.15.1.1 on page 175).
- Use of some C++ reference types was allowed in some data sharing attribute clauses (see Section 2.15.3 on page 185).
- Semantics for reductions on C/C++ array sections were added and restrictions on the use of arrays and pointers in reductions were removed (see Section 2.15.3.6 on page 198).
- The **ref**, **val**, and **uval** modifiers were added to the **linear** clause (see Section 2.15.3.7 on page 204).
- Support was added to the map clauses to handle structure elements (see Section 2.15.5.1 on page 213).
- Query functions for OpenMP thread affinity were added (see Section 3.2.23 on page 253 to Section 3.2.28 on page 258).

- The lock API was extended with lock routines that support storing a hint with a lock to select a desired lock implementation for a lock's intended usage by the application code (see Section 3.3.2 on page 270).
- Device memory routines were added to allow explicit allocation, deallocation, memory transfers and memory associations (see Section 3.5 on page 279).
- C/C++ Grammar (previously Appendix B) was moved to a separate document.

D.2 Version 3.1 to 4.0 Differences

- Various changes throughout the specification were made to provide initial support of Fortran 2003 (see Section 1.6 on page 21).
- C/C++ array syntax was extended to support array sections (see Section 2.4 on page 43).
- The **proc_bind** clause (see Section 2.5.2 on page 51), the **OMP_PLACES** environment variable (see Section 4.5 on page 292), and the **omp_get_proc_bind** runtime routine (see Section 3.2.22 on page 251) were added to support thread affinity policies.
- SIMD constructs were added to support SIMD parallelism (see Section 2.8 on page 71).
- Device constructs (see Section 2.10 on page 93), the **OMP_DEFAULT_DEVICE** environment variable (see Section 4.13 on page 299), the **omp_set_default_device**, **omp_get_default_device**, **omp_get_num_devices**, **omp_get_num_teams**, **omp_get_team_num**, and **omp_is_initial_device** routines were added to support execution on devices.
- Implementation defined task scheduling points for untied tasks were removed (see Section 2.9.5 on page 92).
- The **depend** clause (see Section 2.13.9 on page 165) was added to support task dependences.
- The **taskgroup** construct (see Section 2.13.5 on page 150) was added to support more flexible deep task synchronization.
- The **reduction** clause (see Section 2.15.3.6 on page 198) was extended and the **declare reduction** construct (see Section 2.16 on page 217) was added to support user defined reductions.
- The **atomic** construct (see Section 2.13.6 on page 151) was extended to support atomic swap with the **capture** clause, to allow new atomic update and capture forms, and to support sequentially consistent atomic operations with a new **seq_cst** clause.

- The **cancel** construct (see Section 2.14.1 on page 168), the **cancellation point** construct (see Section 2.14.2 on page 172), the **omp_get_cancellation** runtime routine (see Section 3.2.9 on page 237) and the **OMP_CANCELLATION** environment variable (see Section 4.11 on page 297) were added to support the concept of cancellation.
- The **OMP_DISPLAY_ENV** environment variable (see Section 4.12 on page 298) was added to display the value of ICVs associated with the OpenMP environment variables.
- Examples (previously Appendix A) were moved to a separate document.

D.3 Version 3.0 to 3.1 Differences

- The **final** and **mergeable** clauses (see Section 2.9.1 on page 81) were added to the **task** construct to support optimization of task data environments.
- The **taskyield** construct (see Section 2.9.4 on page 90) was added to allow user-defined task scheduling points.
- The **atomic** construct (see Section 2.13.6 on page 151) was extended to include **read**, **write**, and **capture** forms, and an **update** clause was added to apply the already existing form of the **atomic** construct.
- Data environment restrictions were changed to allow **intent(in)** and **const**-qualified types for the **firstprivate** clause (see Section 2.15.3.4 on page 193).
- Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 2.15.3.4 on page 193) and **lastprivate** (see Section 2.15.3.5 on page 196).
- New reduction operators **min** and **max** were added for C and C++
- The nesting restrictions in Section 2.17 on page 224 were clarified to disallow closely-nested OpenMP regions within an **atomic** region. This allows an **atomic** region to be consistently defined with other OpenMP regions so that they include all the code in the atomic construct.
- The **omp_in_final** runtime library routine (see Section 3.2.21 on page 250) was added to support specialization of final task regions.
- The *nthreads-var* ICV has been modified to be a list of the number of threads to use at each nested parallel region level. The value of this ICV is still set with the **OMP_NUM_THREADS** environment variable (see Section 4.2 on page 290), but the algorithm for determining the number of threads used in a parallel region has been modified to handle a list (see Section 2.5.1 on page 49).

- The *bind-var* ICV has been added, which controls whether or not threads are bound to processors (see Section 2.3.1 on page 36). The value of this ICV can be set with the **OMP_PROC_BIND** environment variable (see Section 4.4 on page 291).
- Descriptions of examples (see Appendix Section A on page 301) were expanded and clarified.
- Replaced incorrect use of **omp_integer_kind** in Fortran interfaces (see Section B.3 on page 332 and Section B.4 on page 339) with **selected_int_kind(8)**.

D.4 Version 2.5 to 3.0 Differences

The concept of tasks has been added to the OpenMP execution model (see Section 1.2.5 on page 9 and Section 1.3 on page 14).

- The **task** construct (see Section 2.9 on page 81) has been added, which provides a mechanism for creating tasks explicitly.
- The **taskwait** construct (see Section 2.13.4 on page 149) has been added, which causes a task to wait for all its child tasks to complete.
- The OpenMP memory model now covers atomicity of memory accesses (see Section 1.4.1 on page 16). The description of the behavior of **volatile** in terms of **flush** was removed.
- In Version 2.5, there was a single copy of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var* internal control variables (ICVs) for the whole program. In Version 3.0, there is one copy of these ICVs per task (see Section 2.3 on page 35). As a result, the **omp_set_num_threads**, **omp_set_nested** and **omp_set_dynamic** runtime library routines now have specified effects when called from inside a **parallel** region (see Section 3.2.1 on page 228, Section 3.2.7 on page 234 and Section 3.2.10 on page 237).
- The definition of active **parallel** region has been changed: in Version 3.0 a **parallel** region is active if it is executed by a team consisting of more than one thread (see Section 1.2.2 on page 2).
- The rules for determining the number of threads used in a **parallel** region have been modified (see Section 2.5.1 on page 49).
- In Version 3.0, the assignment of iterations to threads in a loop construct with a **static** schedule kind is deterministic (see Section 2.7.1 on page 55).
- In Version 3.0, a loop construct may be associated with more than one perfectly nested loop. The number of associated loops may be controlled by the **collapse** clause (see Section 2.7.1 on page 55).

- Random access iterators, and variables of unsigned integer type, may now be used as loop iterators in loops associated with a loop construct (see Section 2.7.1 on page 55).
- The schedule kind **auto** has been added, which gives the implementation the freedom to choose any possible mapping of iterations in a loop construct to threads in the team (see Section 2.7.1 on page 55).
- Fortran assumed-size arrays now have predetermined data-sharing attributes (see Section 2.15.1.1 on page 175).
- In Fortran, **firstprivate** is now permitted as an argument to the **default** clause (see Section 2.15.3.1 on page 186).
- For list items in the **private** clause, implementations are no longer permitted to use the storage of the original list item to hold the new list item on the master thread. If no attempt is made to reference the original list item inside the **parallel** region, its value is well defined on exit from the **parallel** region (see Section 2.15.3.3 on page 189).
- In Version 3.0, Fortran allocatable arrays may appear in **private**, **firstprivate**, **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses. (see Section 2.15.2 on page 179, Section 2.15.3.3 on page 189, Section 2.15.3.4 on page 193, Section 2.15.3.5 on page 196, Section 2.15.3.6 on page 198, Section 2.15.4.1 on page 208 and Section 2.15.4.2 on page 210).
- In Version 3.0, static class members variables may appear in a **threadprivate** directive (see Section 2.15.2 on page 179).
- Version 3.0 makes clear where, and with which arguments, constructors and destructors of private and threadprivate class type variables are called (see Section 2.15.2 on page 179, Section 2.15.3.3 on page 189, Section 2.15.3.4 on page 193, Section 2.15.4.1 on page 208 and Section 2.15.4.2 on page 210).
- The runtime library routines **omp_set_schedule** and **omp_get_schedule** have been added; these routines respectively set and retrieve the value of the *run-sched-var* ICV (see Section 3.2.12 on page 240 and Section 3.2.13 on page 242).
- The *thread-limit-var* ICV has been added, which controls the maximum number of threads participating in the OpenMP program. The value of this ICV can be set with the **OMP_THREAD_LIMIT** environment variable and retrieved with the **omp_get_thread_limit** runtime library routine (see Section 2.3.1 on page 36, Section 3.2.14 on page 243 and Section 4.10 on page 297).
- The *max-active-levels-var* ICV has been added, which controls the number of nested active **parallel** regions. The value of this ICV can be set with the **OMP_MAX_ACTIVE_LEVELS** environment variable and the **omp_set_max_active_levels** runtime library routine, and it can be retrieved with the **omp_get_max_active_levels** runtime library routine (see Section 2.3.1 on page 36, Section 3.2.15 on page 243, Section 3.2.16 on page 245 and Section 4.9 on page 297).

- The *stacksize-var* ICV has been added, which controls the stack size for threads that the OpenMP implementation creates. The value of this ICV can be set with the **OMP_STACKSIZE** environment variable (see Section 2.3.1 on page 36 and Section 4.7 on page 295).
- The *wait-policy-var* ICV has been added, which controls the desired behavior of waiting threads. The value of this ICV can be set with the **OMP_WAIT_POLICY** environment variable (see Section 2.3.1 on page 36 and Section 4.8 on page 296).
- The **omp_get_level** runtime library routine has been added, which returns the number of nested **parallel** regions enclosing the task that contains the call (see Section 3.2.17 on page 246).
- The **omp_get_ancestor_thread_num** runtime library routine has been added, which returns, for a given nested level of the current thread, the thread number of the ancestor (see Section 3.2.18 on page 247).
- The **omp_get_team_size** runtime library routine has been added, which returns, for a given nested level of the current thread, the size of the thread team to which the ancestor belongs (see Section 3.2.19 on page 248).
- The **omp_get_active_level** runtime library routine has been added, which returns the number of nested, active **parallel** regions enclosing the task that contains the call (see Section 3.2.20 on page 249).
- In Version 3.0, locks are owned by tasks, not by threads (see Section 3.3 on page 267).

Index

Symbols

`_OPENMP` macro, [298](#)

`_OPENMP` macro, [33](#)

A

affinity, [51](#)

array sections, [43](#)

atomic, [151](#)

atomic construct, [341](#)

attribute clauses, [185](#)

attributes, data-mapping, [212](#)

attributes, data-sharing, [175](#)

auto, [60](#)

B

barrier, [148](#)

C

C/C++ stub routines, [302](#)

cancel, [168](#)

cancellation constructs, [168](#)

cancel, [168](#)

cancellation point, [172](#)

cancellation point, [172](#)

canonical loop form, [52](#)

clauses

 attribute data-sharing, [185](#)

collapse, [56](#), [57](#)

copyin, [208](#)

copyprivate, [210](#)

 data copying, [208](#)

 data-sharing, [185](#)

default, [186](#)

defaultmap, [217](#)

depend, [165](#)

firstprivate, [193](#)

if Clause, [143](#)

lastprivate, [196](#)

linear, [204](#)

map, [213](#)

private, [189](#)

reduction, [198](#)

schedule, [58](#)

shared, [187](#)

combined constructs, [120](#)

 parallel loop construct, [121](#)

 parallel loop SIMD construct, [125](#)

parallel sections, [122](#)

parallel workshare, [124](#)

target parallel, [126](#)

 target parallel loop, [127](#)

 target parallel loop SIMD, [129](#)

target simd, [130](#)

target teams, [132](#)

target teams distribute, [135](#)

 target teams distribute parallel loop
 construct, [139](#)

 target teams distribute parallel loop
 SIMD construct, [142](#)

target teams distribute simd,
 [137](#)

teams distribute, [133](#)

 teams distribute parallel loop construct,
 [138](#)

 teams distribute parallel loop SIMD
 construct, [141](#)

teams distribute simd, [134](#)

compilation sentinels, [33](#), [34](#)

compliance, [20](#)

conditional compilation, [33](#)

constructs

atomic, [151](#)

barrier, [148](#)

cancel, [168](#)

 cancellation constructs, [168](#)

- cancellation point**, 172
- combined constructs, 120
- critical**, 146
- declare simd**, 75
- declare target**, 102
- device constructs, 93
- distribute**, 109
- distribute parallel do**, 113
- distribute parallel do simd**, 114
- distribute parallel for**, 113
- distribute parallel for simd**, 114
- distribute parallel loop, 113
- distribute parallel loop SIMD, 114
- distribute simd**, 111
- do Fortran*, 55
- flush**, 158
- for, C/C++**, 55
- loop*, 55
- Loop SIMD, 79
- master**, 145
- ordered**, 163
- parallel**, 45
- parallel do Fortran**, 121
- parallel for C/C++**, 121
- parallel loop construct, 121
- parallel loop SIMD construct, 125
- parallel sections**, 122
- parallel workshare**, 124
- sections**, 63
- simd**, 71
- single**, 66
- target**, 95
- target data**, 93
- target enter data**, 116
- target exit data**, 118
- target parallel**, 126
- target parallel do**, 127
- target parallel do simd**, 129
- target parallel for**, 127
- target parallel for simd**, 129
- target parallel loop, 127
- target parallel loop SIMD, 129
- target simd**, 130
- target teams**, 132
- target teams distribute**, 135
- target teams distribute parallel loop
 - construct, 139
- target teams distribute parallel loop
 - SIMD construct, 142
- target teams distribute simd**, 137
- target update**, 99
- task**, 81
- taskgroup**, 150
- tasking constructs, 81
- taskloop**, 85
- taskloop simd**, 89
- taskwait**, 149
- taskyield**, 90
- teams**, 106
- teams distribute**, 133
- teams distribute parallel loop construct, 138
- teams distribute parallel loop SIMD
 - construct, 141
- teams distribute simd**, 134
- workshare**, 68
- worksharing, 55
- controlling OpenMP thread affinity, 51
- copyin**, 208
- copyprivate**, 210
- critical**, 146

D

- data copying clauses, 208
- data environment, 174
- data terminology, 10
- data-mapping rules and clauses, 212
- data-sharing attribute clauses, 185
- data-sharing attribute rules, 175
- declare reduction**, 217
- declare simd**, 75
- declare target**, 102
- default**, 186
- defaultmap**, 217

- depend**, 165
- device constructs, 93
 - declare target**, 102
 - device constructs, 93
 - distribute**, 109
 - distribute parallel loop, 113
 - distribute parallel loop SIMD, 114
 - distribute simd**, 111
 - target**, 95
 - target update**, 99
 - teams**, 106
- device data environments, 18, 116, 118
- device memory routines, 279
- directive format, 26
- directives, 25
 - declare reduction**, 217
 - declare target**, 102
 - threadprivate**, 179
- distribute**, 109
- distribute parallel loop construct, 113
- distribute parallel loop SIMD construct, 114
- distribute simd**, 111
- do**, *Fortran*, 55
- do simd**, 79
- dynamic**, 59
- dynamic thread adjustment, 340

E

- environment variables, 287
 - OMP_CANCELLATION**, 297
 - OMP_DEFAULT_DEVICE**, 299
 - OMP_DISPLAY_ENV**, 298
 - OMP_DYNAMIC**, 291
 - OMP_MAX_ACTIVE_LEVELS**, 297
 - OMP_MAX_TASK_PRIORITY**, 300
 - OMP_NESTED**, 294
 - OMP_NUM_THREADS**, 290
 - OMP_PLACES**, 292
 - OMP_PROC_BIND**, 291
 - OMP_SCHEDULE**, 289
 - OMP_STACKSIZE**, 295
 - OMP_THREAD_LIMIT**, 297
 - OMP_WAIT_POLICY**, 296
- execution environment routines, 228

- execution model, 14

F

- features history, 345
- firstprivate**, 193
- fixed source form conditional compilation
 - sentinels, 33
- fixed source form directives, 28
- flush**, 158
- flush operation, 18
- for**, *C/C++*, 55
- for simd**, 79
- free source form conditional compilation
 - sentinel, 34
- free source form directives, 29

G

- glossary, 2
- guided**, 59

H

- header files, 227, 323
- history of features, 345

I

- ICVs (internal control variables), 35
- if** Clause, 143
- implementation, 340
- implementation terminology, 13
- include files, 227, 323
- interface declarations, 323
- internal control variables, 340
- internal control variables (ICVs), 35
- introduction, 1

L

- lastprivate**, 196
- linear**, 204
- lock routines, 267
- loop**, 55
- loop SIMD Construct, 79
- loop terminology, 8

M

- map**, 213

- master**, 145
- master and synchronization constructs and clauses, 144
- memory model, 16
- modifying and retrieving ICV values, 39
- modifying ICV's, 37
- N**
- nesting of regions, 224
- normative references, 21
- O**
- omp_get_num_teams**, 261
- OMP_CANCELLATION**, 297
- OMP_DEFAULT_DEVICE**, 299
- omp_destroy_lock**, 272
- omp_destroy_nest_lock**, 272
- OMP_DISPLAY_ENV**, 298
- OMP_DYNAMIC**, 291
- omp_get_active_level**, 249
- omp_get_ancestor_thread_num**, 247
- omp_get_cancellation**, 237
- omp_get_default_device**, 260
- omp_get_dynamic**, 236
- omp_get_initial_device**, 264
- omp_get_level**, 246
- omp_get_max_active_levels**, 245
- omp_get_max_task_priority**, 265
- omp_get_max_threads**, 230
- omp_get_nested**, 239
- omp_get_num_devices**, 260
- omp_get_num_places**, 253
- omp_get_num_procs**, 233
- omp_get_num_threads**, 229
- omp_get_partition_num_places**, 257
- omp_get_partition_place_nums**, 258
- omp_get_place_num**, 256
- omp_get_place_num_procs**, 254
- omp_get_place_proc_ids**, 255
- omp_get_proc_bind**, 251
- omp_get_schedule**, 242
- omp_get_team_num**, 263
- omp_get_team_size**, 248
- omp_get_thread_limit**, 243
- omp_get_thread_num**, 232
- omp_get_wtick**, 278
- omp_get_wtime**, 276
- omp_in_final**, 250
- omp_in_parallel**, 233
- omp_init_lock**, 269, 270
- omp_init_nest_lock**, 269, 270
- omp_is_initial_device**, 264
- OMP_MAX_ACTIVE_LEVELS**, 297
- OMP_MAX_TASK_PRIORITY**, 300
- OMP_NESTED**, 294
- OMP_NUM_THREADS**, 290
- OMP_PLACES**, 292
- OMP_PROC_BIND**, 291
- OMP_SCHEDULE**, 289
- omp_set_default_device**, 259
- omp_set_dynamic**, 234
- omp_set_lock**, 273
- omp_set_max_active_levels**, 243
- omp_set_nest_lock**, 273
- omp_set_nested**, 237
- omp_set_num_threads**, 228
- omp_set_schedule**, 240
- OMP_STACKSIZE**, 295
- omp_target_alloc**, 279
- omp_target_associate_ptr**, 284
- omp_target_disassociate_ptr**, 285
- omp_target_free**, 280
- omp_target_is_present**, 281
- omp_target_memcpy**, 282
- omp_target_memcpy_rect**, 283
- omp_test_lock**, 275
- omp_test_nest_lock**, 275
- OMP_THREAD_LIMIT**, 297
- omp_unset_lock**, 274
- omp_unset_nest_lock**, 274
- OMP_WAIT_POLICY**, 296
- OpenMP compliance, 20
- ordered**, 163

P

parallel, 45
parallel loop construct, 121
parallel loop SIMD construct, 125
parallel sections, 122
parallel workshare, 124
private, 189

R

read, atomic, 151
reduction, 198
runtime library definitions, 227
runtime library routines, 226

S

scheduling, 92
sections, 63
shared, 187
simd, 71
SIMD Constructs, 71
Simple Lock Routines, 267
single, 66
stand-alone directives, 32
stub routines, 302
synchronization constructs, 144
synchronization terminology, 8

T

target, 95
target data, 93
target memory routines, 279
target parallel, 126
target parallel loop construct, 127
target parallel loop SIMD construct, 129
target simd, 130
target teams, 132
target teams distribute, 135
target teams distribute parallel loop
 construct, 139
target teams distribute parallel loop SIMD
 construct, 142
target teams distribute simd, 137
target update, 99
task, 81

task scheduling, 92
taskgroup, 150
tasking constructs, 81
tasking terminology, 9
taskloop, 85
taskloop simd, 89
taskwait, 149
taskyield, 90
teams, 106
teams distribute, 133
teams distribute parallel loop construct, 138
teams distribute parallel loop SIMD
 construct, 141
teams distribute simd, 134
thread affinity, 51
threadprivate, 179
timer, 276
timing routines, 276

U

update, atomic, 151

V

variables, environment, 287

W

wall clock timer, 276
workshare, 68
worksharing
 constructs, 55
 parallel, 121
 scheduling, 62
worksharing constructs, 55
write, atomic, 151