



OpenMP Application Program Interface

Version 4.0 - July 2013

Copyright © 1997-2013 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 is intentionally blank.

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 Synchronization Terminology	8
1.2.4 Tasking Terminology	8
1.2.5 Data Terminology	10
1.2.6 Implementation Terminology	12
1.3 Execution Model	14
1.4 Memory Model	17
1.4.1 Structure of the OpenMP Memory Model	17
1.4.2 Device Data Environments	18
1.4.3 The Flush Operation	19
1.4.4 OpenMP Memory Consistency	20
1.5 OpenMP Compliance	21
1.6 Normative References	22
1.7 Organization of this document	23
2. Directives	25
2.1 Directive Format	26
2.1.1 Fixed Source Form Directives	27
2.1.2 Free Source Form Directives	28
2.1.3 Stand-Alone Directives	31
2.2 Conditional Compilation	32
2.2.1 Fixed Source Form Conditional Compilation Sentinels	32

2.2.2	Free Source Form Conditional Compilation Sentinel	33
2.3	Internal Control Variables	34
2.3.1	ICV Descriptions	35
2.3.2	ICV Initialization	36
2.3.3	Modifying and Retrieving ICV Values	37
2.3.4	How ICVs are Scoped	39
2.3.5	ICV Override Relationships	40
2.4	Array Sections	42
2.5	<code>parallel</code> Construct	44
2.5.1	Determining the Number of Threads for a <code>parallel</code> Region	47
2.5.2	Controlling OpenMP Thread Affinity	49
2.6	Canonical Loop Form	51
2.7	Worksharing Constructs	53
2.7.1	Loop Construct	53
2.7.2	<code>sections</code> Construct	60
2.7.3	<code>single</code> Construct	63
2.7.4	<code>workshare</code> Construct	65
2.8	SIMD Constructs	68
2.8.1	<code>simd</code> construct	68
2.8.2	<code>declare simd</code> construct	72
2.8.3	Loop SIMD construct	76
2.9	Device Constructs	77
2.9.1	<code>target data</code> Construct	77
2.9.2	<code>target</code> Construct	79
2.9.3	<code>target update</code> Construct	81
2.9.4	<code>declare target</code> Directive	83
2.9.5	<code>teams</code> Construct	86
2.9.6	<code>distribute</code> Construct	88
2.9.7	<code>distribute simd</code> Construct	91

2.9.8	Distribute Parallel Loop Construct	92
2.9.9	Distribute Parallel Loop SIMD Construct	94
2.10	Combined Constructs	95
2.10.1	Parallel Loop Construct	95
2.10.2	<code>parallel sections</code> Construct	97
2.10.3	<code>parallel workshare</code> Construct	99
2.10.4	Parallel Loop SIMD Construct	100
2.10.5	<code>target teams</code> construct	101
2.10.6	<code>teams distribute</code> Construct	102
2.10.7	<code>teams distribute simd</code> Construct	104
2.10.8	<code>target teams distribute</code> Construct	105
2.10.9	<code>target teams distribute simd</code> Construct	106
2.10.10	Teams Distribute Parallel Loop Construct	107
2.10.11	Target Teams Distribute Parallel Loop Construct	109
2.10.12	Teams Distribute Parallel Loop SIMD Construct	110
2.10.13	Target Teams Distribute Parallel Loop SIMD Construct	111
2.11	Tasking Constructs	113
2.11.1	<code>task</code> Construct	113
2.11.2	<code>taskyield</code> Construct	117
2.11.3	Task Scheduling	118
2.12	Master and Synchronization Constructs	120
2.12.1	<code>master</code> Construct	120
2.12.2	<code>critical</code> Construct	122
2.12.3	<code>barrier</code> Construct	123
2.12.4	<code>taskwait</code> Construct	125
2.12.5	<code>taskgroup</code> Construct	126
2.12.6	<code>atomic</code> Construct	127
2.12.7	<code>flush</code> Construct	134
2.12.8	<code>ordered</code> Construct	138
2.13	Cancellation Constructs	140

2.13.1	<code>cancel</code> Construct	140
2.13.2	<code>cancellation point</code> Construct	143
2.14	Data Environment	146
2.14.1	Data-sharing Attribute Rules	146
2.14.2	<code>threadprivate</code> Directive	150
2.14.3	Data-Sharing Attribute Clauses	155
2.14.4	Data Copying Clauses	173
2.14.5	<code>map</code> Clause	177
2.15	<code>declare reduction</code> Directive	180
2.16	Nesting of Regions	186
3.	Runtime Library Routines	187
3.1	Runtime Library Definitions	188
3.2	Execution Environment Routines	189
3.2.1	<code>omp_set_num_threads</code>	189
3.2.2	<code>omp_get_num_threads</code>	191
3.2.3	<code>omp_get_max_threads</code>	192
3.2.4	<code>omp_get_thread_num</code>	193
3.2.5	<code>omp_get_num_procs</code>	195
3.2.6	<code>omp_in_parallel</code>	196
3.2.7	<code>omp_set_dynamic</code>	197
3.2.8	<code>omp_get_dynamic</code>	198
3.2.9	<code>omp_get_cancellation</code>	199
3.2.10	<code>omp_set_nested</code>	200
3.2.11	<code>omp_get_nested</code>	201
3.2.12	<code>omp_set_schedule</code>	203
3.2.13	<code>omp_get_schedule</code>	205
3.2.14	<code>omp_get_thread_limit</code>	206
3.2.15	<code>omp_set_max_active_levels</code>	207
3.2.16	<code>omp_get_max_active_levels</code>	209

3.2.17	omp_get_level	210
3.2.18	omp_get_ancestor_thread_num	211
3.2.19	omp_get_team_size	212
3.2.20	omp_get_active_level	214
3.2.21	omp_in_final	215
3.2.22	omp_get_proc_bind	216
3.2.23	omp_set_default_device	218
3.2.24	omp_get_default_device	219
3.2.25	omp_get_num_devices	220
3.2.26	omp_get_num_teams	221
3.2.27	omp_get_team_num	222
3.2.28	omp_is_initial_device	223
3.3	Lock Routines	224
3.3.1	omp_init_lock and omp_init_nest_lock	226
3.3.2	omp_destroy_lock and omp_destroy_nest_lock	227
3.3.3	omp_set_lock and omp_set_nest_lock	228
3.3.4	omp_unset_lock and omp_unset_nest_lock	229
3.3.5	omp_test_lock and omp_test_nest_lock	231
3.4	Timing Routines	233
3.4.1	omp_get_wtime	233
3.4.2	omp_get_wtick	234
4.	Environment Variables	237
4.1	OMP_SCHEDULE	238
4.2	OMP_NUM_THREADS	239
4.3	OMP_DYNAMIC	240
4.4	OMP_PROC_BIND	241
4.5	OMP_PLACES	241
4.6	OMP_NESTED	243
4.7	OMP_STACKSIZE	244

4.8	<code>OMP_WAIT_POLICY</code>	245
4.9	<code>OMP_MAX_ACTIVE_LEVELS</code>	245
4.10	<code>OMP_THREAD_LIMIT</code>	246
4.11	<code>OMP_CANCELLATION</code>	246
4.12	<code>OMP_DISPLAY_ENV</code>	247
4.13	<code>OMP_DEFAULT_DEVICE</code>	248
A.	Stubs for Runtime Library Routines	249
A.1	C/C++ Stub Routines	250
A.2	Fortran Stub Routines	257
B.	OpenMP C and C++ Grammar	265
B.1	Notation	265
B.2	Rules	266
C.	Interface Declarations	287
C.1	Example of the <code>omp.h</code> Header File	288
C.2	Example of an Interface Declaration <code>include</code> File	290
C.3	Example of a Fortran Interface Declaration <code>module</code>	293
C.4	Example of a Generic Interface for a Library Routine	298
D.	OpenMP Implementation-Defined Behaviors	299
E.	Features History	303
E.1	Version 3.1 to 4.0 Differences	303
E.2	Version 3.0 to 3.1 Differences	304
E.3	Version 2.5 to 3.0 Differences	305
	Index	309

2 Introduction

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

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

11 `http://www.openmp.org`

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

21 1.1 Scope

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

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

4 1.2 Glossary

5 1.2.1 Threading Concepts

6
7 **thread** An execution entity with a stack and associated static memory, called
8 *threadprivate memory*.

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

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

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

14 **device** An implementation defined logical execution engine.

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

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

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

18 1.2.2 OpenMP Language Terminology

19
20 **base language** A programming language that serves as the foundation of the OpenMP
21 specification.

22 COMMENT: See Section 1.6 on page 22 for a listing of current *base languages*
23 for the OpenMP API.

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

1	structured block	For C/C++, an executable statement, possibly compound, with a single entry at the
2		top and a single exit at the bottom, or an OpenMP <i>construct</i> .
3		For Fortran, a block of executable statements with a single entry at the top and a
4		single exit at the bottom, or an OpenMP <i>construct</i> .
5		COMMENTS:
6		For all <i>base languages</i> ,
7		• Access to the <i>structured block</i> must not be the result of a branch.
8		• The point of exit cannot be a branch out of the <i>structured block</i> .
9		
10		For C/C++:
11		• The point of entry must not be a call to <code>setjmp()</code> .
12		• <code>longjmp()</code> and <code>throw()</code> must not violate the entry/exit criteria.
13		• Calls to <code>exit()</code> are allowed in a <i>structured block</i> .
14		• An expression statement, iteration statement, selection statement,
15		or try block is considered to be a <i>structured block</i> if the
16		corresponding compound statement obtained by enclosing it in {
17		and } would be a <i>structured block</i> .
18		
19		For Fortran:
20		• STOP statements are allowed in a <i>structured block</i> .
21		
22	enclosing context	In C/C++, the innermost scope enclosing an OpenMP <i>directive</i> .
23		In Fortran, the innermost scoping unit enclosing an OpenMP <i>directive</i> .
24	directive	In C/C++, a #pragma , and in Fortran, a comment, that specifies <i>OpenMP</i>
25		<i>program</i> behavior.
26		COMMENT: See Section 2.1 on page 26 for a description of OpenMP <i>directive</i>
27		syntax.
28	white space	A non-empty sequence of space and/or horizontal tab characters.
29	OpenMP program	A program that consists of a <i>base program</i> , annotated with OpenMP <i>directives</i> and
30		runtime library routines.
31	conforming program	An <i>OpenMP program</i> that follows all the rules and restrictions of the OpenMP
32		specification.

1	declarative directive	An OpenMP <i>directive</i> that may only be placed in a declarative context. A
2		<i>declarative directive</i> results in one or more declarations only; it is not associated
3		with the immediate execution of any user code.
4	executable directive	An OpenMP <i>directive</i> that is not declarative. That is, it may be placed in an
5		executable context.
6	stand-alone directive	An OpenMP <i>executable directive</i> that has no associated executable user code.
7	loop directive	An OpenMP <i>executable directive</i> whose associated user code must be a loop nest
8		that is a <i>structured block</i> .
9	associated loop(s)	The loop(s) controlled by a <i>loop directive</i> .
10		COMMENT: If the <i>loop directive</i> contains a collapse clause then there may be
11		more than one <i>associated loop</i> .
12	construct	An OpenMP <i>executable directive</i> (and for Fortran, the paired end directive , if
13		any) and the associated statement, loop or <i>structured block</i> , if any, not including
14		the code in any called routines. That is, in the lexical extent of an <i>executable</i>
15		<i>directive</i> .
16	region	All code encountered during a specific instance of the execution of a given
17		<i>construct</i> or of an OpenMP library routine. A <i>region</i> includes any code in called
18		routines as well as any implicit code introduced by the OpenMP implementation.
19		The generation of a <i>task</i> at the point where a task directive is encountered is a
20		part of the <i>region</i> of the <i>encountering thread</i> , but the <i>explicit task region</i>
21		associated with the task directive is not. The point where a target or teams
22		directive is encountered is a part of the <i>region</i> of the <i>encountering thread</i> , but the
23		<i>region</i> associated with the target or teams directive is not.
24		COMMENTS:
25		<i>A region</i> may also be thought of as the dynamic or runtime extent of a
26		<i>construct</i> or of an OpenMP library routine.
27		During the execution of an <i>OpenMP program</i> , a <i>construct</i> may give
28		rise to many <i>regions</i> .
29		
30	active parallel region	A parallel region that is executed by a <i>team</i> consisting of more than one
31		<i>thread</i> .
32	inactive parallel region	A parallel region that is executed by a <i>team</i> of only one <i>thread</i> .

1	sequential part	All code encountered during the execution of an <i>initial task region</i> that is not part
2		of a parallel region corresponding to a parallel construct or a task
3		<i>region</i> corresponding to a task construct .
4		COMMENTS:
5		<i>A sequential part is enclosed by an implicit parallel region.</i>
6		Executable statements in called routines may be in both a <i>sequential</i>
7		<i>part</i> and any number of explicit parallel regions at different points
8		in the program execution.
9	master thread	The <i>thread</i> that encounters a parallel construct , creates a <i>team</i> , generates a set
10		of <i>implicit tasks</i> , then executes one of those <i>tasks</i> as <i>thread</i> number 0.
11	parent thread	The <i>thread</i> that encountered the parallel construct and generated a
12		parallel region is the <i>parent thread</i> of each of the <i>threads</i> in the <i>team</i> of that
13		parallel region . The <i>master thread</i> of a parallel region is the same <i>thread</i>
14		as its <i>parent thread</i> with respect to any resources associated with an <i>OpenMP</i>
15		<i>thread</i> .
16	child thread	When a <i>thread</i> encounters a parallel construct , each of the <i>threads</i> in the
17		generated parallel region's <i>team</i> are <i>child threads</i> of the encountering <i>thread</i> .
18		The target or teams region's <i>initial thread</i> is not a <i>child thread</i> of the <i>thread</i>
19		that encountered the target or teams construct.
20	ancestor thread	For a given <i>thread</i> , its <i>parent thread</i> or one of its <i>parent thread's</i> <i>ancestor threads</i> .
21	descendent thread	For a given <i>thread</i> , one of its <i>child threads</i> or one of its <i>child threads' descendent</i>
22		<i>threads</i> .
23	team	A set of one or more <i>threads</i> participating in the execution of a parallel
24		<i>region</i> .
25		COMMENTS:
26		For an <i>active parallel region</i> , the <i>team</i> comprises the <i>master thread</i>
27		and at least one additional <i>thread</i> .
28		For an <i>inactive parallel region</i> , the <i>team</i> comprises only the <i>master</i>
29		<i>thread</i> .
30	league	The set of <i>thread teams</i> created by a target construct or a teams construct.
31	contention group	An <i>initial thread</i> and its <i>descendent threads</i> .
32	implicit parallel region	An <i>inactive parallel region</i> that generates an <i>initial task region</i> . <i>Implicit parallel</i>
33		<i>regions</i> surround the whole <i>OpenMP</i> program, all target regions, and all
34		teams regions.

1	initial thread	A <i>thread</i> that executes an <i>implicit parallel region</i> .
2	nested construct	A <i>construct</i> (lexically) enclosed by another <i>construct</i> .
3	closely nested construct	A <i>construct</i> nested inside another <i>construct</i> with no other <i>construct</i> nested between them.
4		
5	nested region	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> encountered during the execution of another <i>region</i> .
6		
7		COMMENT: Some nestings are <i>conforming</i> and some are not. See Section 2.16 on
8		page 186 for the restrictions on nesting.
9	closely nested region	A <i>region nested</i> inside another <i>region</i> with no parallel <i>region nested</i> between them.
10		
11	all threads	All OpenMP <i>threads</i> participating in the <i>OpenMP program</i> .
12	current team	All <i>threads</i> in the <i>team</i> executing the innermost enclosing parallel <i>region</i> .
13	encountering thread	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
14	all tasks	All <i>tasks</i> participating in the <i>OpenMP program</i> .
15	current team tasks	All <i>tasks</i> encountered by the corresponding <i>team</i> . Note that the <i>implicit tasks</i> constituting the parallel <i>region</i> and any <i>descendent tasks</i> encountered during the execution of these <i>implicit tasks</i> are included in this set of tasks.
16		
17		
18	generating task	For a given <i>region</i> , the <i>task</i> whose execution by a <i>thread</i> generated the <i>region</i> .
19	binding thread set	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a <i>region</i> .
20		
21		The <i>binding thread set</i> for a given <i>region</i> can be <i>all threads</i> on a <i>device</i> , <i>all threads</i> in a <i>contention group</i> , the <i>current team</i> , or the <i>encountering thread</i> .
22		
23		COMMENT: The <i>binding thread set</i> for a particular <i>region</i> is described in its corresponding subsection of this specification.
24		
25	binding task set	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a <i>region</i> .
26		
27		The <i>binding task set</i> for a given <i>region</i> can be <i>all tasks</i> , the <i>current team tasks</i> , or the <i>generating task</i> .
28		
29		COMMENT: The <i>binding task set</i> for a particular <i>region</i> (if applicable) is described in its corresponding subsection of this specification.
30		

1	binding region	The enclosing <i>region</i> that determines the execution context and limits the scope of the effects of the bound <i>region</i> is called the <i>binding region</i> .
2		
3		<i>Binding region</i> is not defined for <i>regions</i> whose <i>binding thread set</i> is <i>all threads</i>
4		or the <i>encountering thread</i> , nor is it defined for <i>regions</i> whose <i>binding task set</i> is
5		<i>all 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		For all other <i>regions</i> for which the <i>binding thread set</i> is the <i>current</i>
12		<i>team</i> or the <i>binding task set</i> is the <i>current team tasks</i> , the <i>binding</i>
13		<i>region</i> is the innermost enclosing parallel <i>region</i> .
14		For <i>regions</i> for which the <i>binding task set</i> is the <i>generating task</i> , the
15		<i>binding region</i> is the <i>region</i> of the <i>generating task</i> .
16		A parallel <i>region</i> need not be <i>active</i> nor explicit to be a <i>binding</i>
17		<i>region</i> .
18		A <i>task region</i> need not be explicit to be a <i>binding region</i> .
19		A <i>region</i> never binds to any <i>region</i> outside of the innermost enclosing
20		parallel <i>region</i> .
21	orphaned construct	A <i>construct</i> that gives rise to a <i>region</i> whose <i>binding thread set</i> is the <i>current</i>
22		<i>team</i> , but is not nested within another <i>construct</i> giving rise to the <i>binding region</i> .
23	worksharing construct	A <i>construct</i> that defines units of work, each of which is executed exactly once by
24		one of the <i>threads</i> in the <i>team</i> executing the <i>construct</i> .
25		For C/C++, <i>worksharing constructs</i> are for , sections , and single .
26		For Fortran, <i>worksharing constructs</i> are do , sections , single and
27		workshare .
28	sequential loop	A loop that is not associated with any OpenMP <i>loop directive</i> .
29	place	Unordered set of <i>processors</i> that is treated by the execution environment as a
30		location unit when dealing with OpenMP thread affinity.
31	place list	The ordered list that describes all OpenMP <i>places</i> available to the execution
32		environment.

1	place partition	An ordered list that corresponds to a contiguous interval in the OpenMP <i>place list</i> .
2		It describes the <i>places</i> currently available to the execution environment for a given
3		parallel region.
4	SIMD instruction	A single machine instruction that can can operate on multiple data elements.
5	SIMD lane	A software or hardware mechanism capable of processing one data element from a
6		<i>SIMD instruction</i> .
7	SIMD chunk	A set of iterations executed concurrently, each by a <i>SIMD lane</i> , by a single <i>thread</i>
8		by means of <i>SIMD instructions</i> .
9	SIMD loop	A loop that includes at least one <i>SIMD chunk</i> .

10 1.2.3 Synchronization Terminology

11	barrier	A point in the execution of a program encountered by a <i>team of threads</i> , beyond
12		which no <i>thread</i> in the team may execute until all <i>threads</i> in the <i>team</i> have
13		reached the barrier and all <i>explicit tasks</i> generated by the <i>team</i> have executed to
14		completion. If <i>cancellation</i> has been requested, threads may proceed to the end of
15		the canceled <i>region</i> even if some threads in the team have not reached the <i>barrier</i> .
16	cancellation	An action that cancels (that is, aborts) an OpenMP <i>region</i> and causes executing
17		<i>implicit</i> or <i>explicit</i> tasks to proceed to the end of the canceled <i>region</i> .
18	cancellation point	A point at which implicit and explicit tasks check if cancellation has been
19		requested. If cancellation has been observed, they perform the <i>cancellation</i> .
20		COMMENT: For a list of cancellation points, see Section 2.13.1 on page 140.

21 1.2.4 Tasking Terminology

22	task	A specific instance of executable code and its <i>data environment</i> , generated when a
23		<i>thread</i> encounters a task construct or a parallel construct .
24	task region	A <i>region</i> consisting of all code encountered during the execution of a <i>task</i> .
25		COMMENT: A parallel region consists of one or more implicit <i>task regions</i> .
26	explicit task	A <i>task</i> generated when a task construct is encountered during execution.
27	implicit task	A <i>task</i> generated by an <i>implicit parallel region</i> or generated when a parallel
28		<i>construct</i> is encountered during execution.
29	initial task	An <i>implicit task</i> associated with an <i>implicit parallel region</i> .
30	current task	For a given <i>thread</i> , the <i>task</i> corresponding to the <i>task region</i> in which it is
31		executing.

1	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
2		of its generating <i>task region</i> .
3	sibling tasks	<i>Tasks</i> that are <i>child tasks</i> of the same <i>task region</i> .
4	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</i>
5		<i>regions</i> .
6	task completion	<i>Task completion</i> occurs when the end of the <i>structured block</i> associated with the
7		<i>construct</i> that generated the <i>task</i> is reached.
8		COMMENT: Completion of the <i>initial task</i> occurs at program exit.
9	task scheduling point	A point during the execution of the current <i>task region</i> at which it can be
10		suspended to be resumed later; or the point of <i>task completion</i> , after which the
11		executing thread may switch to a different <i>task region</i> .
12		COMMENT: For a list of task scheduling points, see Section 2.11.3 on page 118.
13	task switching	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
14	tied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same
15		<i>thread</i> that suspended it. That is, the <i>task</i> is tied to that <i>thread</i> .
16	untied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in
17		the team. That is, the <i>task</i> is not tied to any <i>thread</i> .
18	undeferrred task	A <i>task</i> for which execution is not deferred with respect to its generating <i>task</i>
19		<i>region</i> . That is, its generating <i>task region</i> is suspended until execution of the
20		<i>undeferrred task</i> is completed.
21	included task	A <i>task</i> for which execution is sequentially included in the generating <i>task region</i> .
22		That is, an <i>included task</i> is <i>undeferrred</i> and executed immediately by the
23		<i>encountering thread</i> .
24	merged task	A <i>task</i> whose <i>data environment</i> , inclusive of ICVs, is the same as that of its
25		generating <i>task region</i> .
26	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> .
27	task dependence	An ordering relation between two <i>sibling tasks</i> : the <i>dependent task</i> and a
28		previously generated <i>predecessor task</i> . The <i>task dependence</i> is fulfilled when the
29		<i>predecessor task</i> has completed.
30	dependent task	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor</i>
31		<i>tasks</i> have completed.
32	predecessor task	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
33	task synchronization construct	A taskwait , taskgroup , or a barrier <i>construct</i> .

1 1.2.5 Data Terminology

2 **variable** A named data storage block, whose value can be defined and redefined during the
3 execution of a program.

4 **Note** – An array or structure element is a variable that is part of another variable.

5 **array section** A designated subset of the elements of an array.

6 **array item** An array, an array section or an array element.

7 **private variable** With respect to a given set of *task regions* or *SIMD lanes* that bind to the same
8 **parallel region**, a *variable* whose name provides access to a different block of
9 storage for each *task region* or *SIMD lane*.

10 A *variable* that is part of another variable (as an array or structure element) cannot
11 be made private independently of other components.

12 **shared variable** With respect to a given set of *task regions* that bind to the same **parallel**
13 *region*, a *variable* whose name provides access to the same block of storage for
14 each *task region*.

15 A *variable* that is part of another variable (as an array or structure element) cannot
16 be *shared* independently of the other components, except for static data members
17 of C++ classes.

18 **threadprivate variable** A *variable* that is replicated, one instance per *thread*, by the OpenMP
19 implementation. Its name then provides access to a different block of storage for
20 each *thread*.

21 A *variable* that is part of another variable (as an array or structure element) cannot
22 be made *threadprivate* independently of the other components, except for static
23 data members of C++ classes.

24 **threadprivate memory** The set of *threadprivate variables* associated with each *thread*.

25 **data environment** The *variables* associated with the execution of a given *region*.

26 **device data environment** A *data environment* defined by a **target data** or **target** construct.

27

1	mapped variable	An original <i>variable</i> in a <i>data environment</i> with a corresponding <i>variable</i> in a device <i>data environment</i> .
2		
3		COMMENT: The original and corresponding <i>variables</i> may share
4		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
7		not mappable then the type is not mappable.
8		COMMENT: Pointer types are <i>mappable</i> but the memory block to
9		which the pointer refers is not <i>mapped</i> .
10		For C:
11		The type must be a complete type.
12		For C++:
13		The type must be a complete type.
14		In addition, for class types:
15		• All member functions accessed in any target region must appear in a
16		declare target directive.
17		• All data members must be non-static.
18		• A <i>mappable type</i> cannot contain virtual members.
19		
20		For Fortran:
21		The type must be definable.
22	defined	For <i>variables</i> , the property of having a valid value.
23		For C:
24		For the contents of <i>variables</i> , the property of having a valid value.
25		For C++:
26		For the contents of <i>variables</i> of POD (plain old data) type, the property of having
27		a valid value.
28		For <i>variables</i> of non-POD class type, the property of having been constructed but
29		not subsequently destructed.
30		For Fortran:
31		For the contents of <i>variables</i> , the property of having a valid value. For the
32		allocation or association status of <i>variables</i> , the property of having a valid status.
33		COMMENT: Programs that rely upon <i>variables</i> that are not <i>defined</i> are
34		<i>non-conforming programs</i> .
35	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 1.2.6 **Implementation Terminology**

- 4 **supporting n levels of**
 parallelism Implies allowing an *active parallel region* to be enclosed by $n-1$ *active parallel regions*.

- 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 *threads* or *tasks* in an *OpenMP program*.

10 COMMENT: The acronym ICV is used interchangeably with the term *internal control variable* in the remainder of this specification.

- 12 **compliant**
 implementation An implementation of the OpenMP specification that compiles and executes any *conforming program* as defined by the specification.

14 COMMENT: A *compliant implementation* may exhibit *unspecified behavior* when compiling or executing a *non-conforming program*.

- 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 *OpenMP program*.

18 Such *unspecified behavior* may result from:

- 19 • Issues documented by the OpenMP specification as having *unspecified behavior*.
- 20 • A *non-conforming program*.
- 21 • A *conforming program* exhibiting an *implementation defined* behavior.

implementation

1 **defined** Behavior that must be documented by the implementation, and is allowed to vary
2 among different *compliant implementations*. An implementation is allowed to
3 define this behavior as *unspecified*.

4 COMMENT: All features that have *implementation defined* behavior
5 are documented in Appendix D.

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 a target device does not exist, or the target device is not supported by the implementation, or the target device cannot execute the **target** construct then the **target** region is executed by the host device.

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.

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

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

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

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

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

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

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

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

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

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

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

28

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 2.14.3.3 on page 159 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 2.14 on page 146.

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.

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

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

14 1.4.2 Device Data Environments

15 When an OpenMP program begins, each device has an initial device data environment.
16 The initial device data environment for the host device is the data environment
17 associated with the initial task region. Directives that accept data-mapping attribute
18 clauses determine how an original variable is mapped to a corresponding variable in a
19 device data environment. The original variable is the variable with the same name that
20 exists in the data environment of the task that encounters the directive.

21 If a corresponding variable is present in the enclosing device data environment, the new
22 device data environment inherits the corresponding variable from the enclosing device
23 data environment. If a corresponding variable is not present in the enclosing device data
24 environment, a new corresponding variable (of the same type and size) is created in the
25 new device data environment. In the latter case, the initial value of the new
26 corresponding variable is determined from the clauses and the data environment of the
27 encountering thread.

28 The corresponding variable in the device data environment may share storage with the
29 original variable. Writes to the corresponding variable may alter the value of the original
30 variable. The impact of this on memory consistency is discussed in Section 1.4.4 on
31 page 20. When a task executes in the context of a device data environment, references to
32 the original variable refer to the corresponding variable in the device data environment.

33 The relationship between the value of the original variable and the initial or final value
34 of the corresponding variable depends on the *map-type*. Details of this issue, as well as
35 other issues with mapping a variable, are provided in Section 2.14.5 on page 177.

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

1 1.4.3 The Flush Operation

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

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

13 If a thread has performed a write to its temporary view of a shared variable since its last
14 flush of that variable, then when it executes another flush of the variable, the flush does
15 not complete until the value of the variable has been written to the variable in memory.
16 If a thread performs multiple writes to the same variable between two flushes of that
17 variable, the flush ensures that the value of the last write is written to the variable in
18 memory. A flush of a variable executed by a thread also causes its temporary view of the
19 variable to be discarded, so that if its next memory operation for that variable is a read,
20 then the thread will read from memory when it may again capture the value in the
21 temporary view. When a thread executes a flush, no later memory operation by that
22 thread for a variable involved in that flush is allowed to start until the flush completes.
23 The completion of a flush of a set of variables executed by a thread is defined as the
24 point at which all writes to those variables performed by the thread before the flush are
25 visible in memory to all other threads and that thread's temporary view of all variables
26 involved is discarded.

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

- 33 1. The value is written to the variable by the first thread.
- 34 2. The variable is flushed by the first thread.
- 35 3. The variable is flushed by the second thread.
- 36 4. The value is read from the variable by the second thread.

1 **Note** – OpenMP synchronization operations, described in Section 2.12 on page 120 and
2 in Section 3.3 on page 224, are recommended for enforcing this order. Synchronization
3 through variables is possible but is not recommended because the proper timing of
4 flushes is difficult.

5 1.4.4 OpenMP Memory Consistency

6 The restrictions in Section 1.4.3 on page 19 on reordering with respect to flush
7 operations guarantee the following:

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

16 The flush operation can be specified using the **flush** directive, and is also implied at
17 various locations in an OpenMP program: see Section 2.12.7 on page 134 for details.

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

21 OpenMP programs that:

- 22 • do not use non-sequentially consistent atomic directives,
- 23 • do not rely on the accuracy of a *false* result from **omp_test_lock** and
24 **omp_test_nest_lock**, and
- 25 • correctly avoid data races as required in Section 1.4.1 on page 17

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

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

4 1.5 OpenMP Compliance

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

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

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

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

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

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
- Allocatable enhancement
- Parameterized derived types
- Finalization
- Procedures bound by name to a type
- The **PASS** attribute

- 1 • Procedures bound to a type as operators
- 2 • Type extension
- 3 • Overriding a type-bound procedure
- 4 • Polymorphic entities
- 5 • **SELECT TYPE** construct
- 6 • Deferred bindings and abstract types
- 7 • Controlling IEEE underflow
- 8 • Another IEEE class value

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



11 1.7 Organization of this document

12 The remainder of this document is structured as follows:

- 13 • Chapter 2 “Directives”
- 14 • Chapter 3 “Runtime Library Routines”
- 15 • Chapter 4 “Environment Variables”
- 16 • Appendix A “Stubs for Runtime Library Routines”
- 17 • Appendix B “OpenMP C and C++ Grammar”
- 18 • Appendix C “Interface Declarations”
- 19 • Appendix D “OpenMP Implementation-Defined Behaviors”
- 20 • Appendix E “Features History”

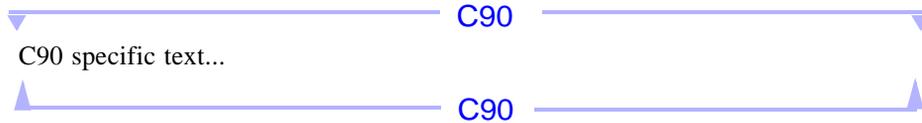
21 Some sections of this document only apply to programs written in a certain base
 22 language. Text that applies only to programs whose base language is C or C++ is shown
 23 as follows:



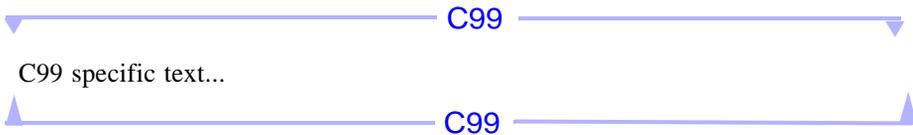
25 Text that applies only to programs whose base language is C only is shown as follows:



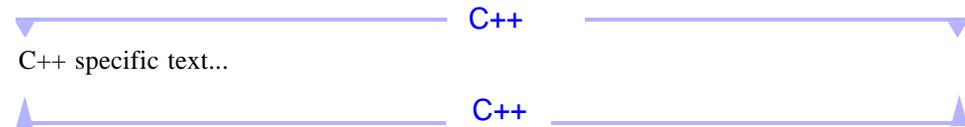
1 Text that applies only to programs whose base language is C90 only is shown as
2 follows:



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



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



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



12 Where an entire page consists of, for example, Fortran specific text, a marker is shown
13 at the top of the page like this:



14 Some text is for information only, and is not part of the normative specification. Such
15 text is designated as a note, like this:



2

Directives

3 This chapter describes the syntax and behavior of OpenMP directives, and is divided
4 into the 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 32)
- 7 • How to specify and to use array sections for all base languages (Section 2.4 on page
8 42)
- 9 • Control of OpenMP API ICVs (Section 2.3 on page 34)
- 10 • Details of each OpenMP directive (Section 2.5 on page 44 to Section 2.16 on page
11 186)

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

 C/C++ 

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

 Fortran 

17 Compilers can therefore ignore OpenMP directives and conditionally compiled code if
18 support of the OpenMP API is not provided or enabled. A compliant implementation
19 must provide an option or interface that ensures that underlying support of all OpenMP
20 directives and OpenMP conditional compilation mechanisms is enabled. In the
21 remainder of this document, the phrase *OpenMP compilation* is used to mean a
22 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 formally specified by the grammar in Appendix B, and informally 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.

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 27 and Section 2.1.2 on page 28.

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

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

Fortran

5 Only one *directive-name* can be specified per directive (note that this includes combined
6 directives, see Section 2.10 on page 95). The order in which clauses appear on directives
7 is not significant. Clauses on directives may be repeated as needed, subject to the
8 restrictions listed in the description of each clause.

9 Some data-sharing attribute clauses (Section 2.14.3 on page 155), data copying clauses
10 (Section 2.14.4 on page 173), the **threadprivate** directive (Section 2.14.2 on page
11 150) and the **flush** directive (Section 2.12.7 on page 134) accept a *list*. A *list* consists
12 of a comma-separated collection of one or more *list items*.

C/C++

13 A *list item* is a variable or array section, subject to the restrictions specified in
14 Section 2.4 on page 42 and in each of the sections describing clauses and directives for
15 which a *list* appears.

C/C++

Fortran

16 A *list item* is a variable, array section or common block name (enclosed in slashes),
17 subject to the restrictions specified in Section 2.4 on page 42 and in each of the sections
18 describing clauses and directives for which a *list* appears.

Fortran

Fortran

20 2.1.1 Fixed Source Form Directives

21 The following sentinels are recognized in fixed form source files:

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

22 Sentinels must start in column 1 and appear as a single word with no intervening
23 characters. Fortran fixed form line length, white space, continuation, and column rules
24 apply to the directive line. Initial directive lines must have a space or zero in column 6,
25 and continuation directive lines must have a character other than a space or a zero in
26 column 6.

Fortran (cont.)

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

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

```
8           c23456789  
9           !$omp parallel do shared(a,b,c)  
  
10           c$omp parallel do  
11           c$omp+shared(a,b,c)  
  
12           c$omp paralleldoshared(a,b,c)  
13  
14           c$omp paralleldoshared(a,b,c)
```

2.1.2 Free Source Form Directives

16 The following sentinel is recognized in free form source files:

```
!$omp
```

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

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

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

```

4     declare reduction
5     declare simd
6     declare target
7     distribute parallel do
8     distribute parallel do simd
9     distribute simd
10    do simd
11    end atomic
12    end critical
13    end distribute
14    end distribute parallel do
15    end distribute parallel do simd
16    end distribute simd
17    end do
18    end do simd
19    end master
20    end ordered
21    end parallel
22    end parallel do
23    end parallel do simd
24    end parallel sections
25    end parallel workshare
26    end sections
27    end simd
  
```

```

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

```

```
1      target update
2      teams distribute
3      teams distribute parallel do
4      teams distribute parallel do simd
5      teams distribute simd
```



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

```
8      !23456789
9          !$omp parallel do &
10             !$omp shared(a,b,c)
11
12             !$omp parallel &
13             !$omp&do shared(a,b,c)
14
15             !$omp paralleldo shared(a,b,c)
```



16



Fortran

17 2.1.3 Stand-Alone Directives

18 Summary

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

20 Description

21 Stand-alone directives do not have any associated executable user code. Instead, they
22 represent executable statements that typically do not have succinct equivalent statements
23 in the base languages. There are some restrictions on the placement of a stand-alone
24 directive within a program. A stand-alone directive may be placed only at a point where
25 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**. See Appendix B for the formal grammar.

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

2.2 Conditional Compilation

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

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

Fortran

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

2.2.1 Fixed Source Form Conditional Compilation Sentinels

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

!\$ *\$ c\$

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

- The sentinel must start in column 1 and appear as a single word with no intervening white space.

Fortran (cont.)

- 1 • After the sentinel is replaced with two spaces, initial lines must have a space or zero
- 2 in column 6 and only white space and numbers in columns 1 through 5.
- 3 • After the sentinel is replaced with two spaces, continuation lines must have a
- 4 character other than a space or zero in column 6 and only white space in columns 1
- 5 through 5.
- 6 If these criteria are met, the sentinel is replaced by two spaces. If these criteria are not
- 7 met, the line 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:

```
!$
```

- 23 To enable conditional compilation, a line with a conditional compilation sentinel must
- 24 satisfy the following criteria:
- 25 • The sentinel can appear in any column but must be preceded only by white space.
- 26 • The sentinel must appear as a single word with no intervening white space.
- 27 • Initial lines must have a space after the sentinel.

- 1 • Continued lines must have an ampersand as the last non-blank character on the line,
2 prior to any comment appearing on the conditionally compiled line. Continued lines
3 can have an ampersand after the sentinel, with optional white space before and after
4 the ampersand.

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

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

```
10 c23456789  
11 !$ iam = omp_get_thread_num() + &  
12 !$& index  
13  
14 #ifdef _OPENMP  
15     iam = omp_get_thread_num() + &  
16     index  
17 #endif
```

▲
18
▲————— Fortran —————▲

19 2.3 Internal Control Variables

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

28 For purposes of exposition, this document refers to the ICVs by certain names, but an
29 implementation is not required to use these names or to offer any way to access the
30 variables other than through the ways shown in Section 2.3.2 on page 36.

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

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

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

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

- 33 • *stacksize-var* - controls the stack size for threads that the OpenMP implementation
34 creates. There is one copy of this ICV per device.
- 35 • *wait-policy-var* - controls the desired behavior of waiting threads. There is one copy
36 of this ICV per device.
- 37 • *cancel-var* - controls the desired behavior of the **cancel** construct and cancellation
38 points. There is one copy of the ICV for the whole program (the scope is global).

- *default-device-var* - controls the default target device. There is one copy of this ICV per data environment.

2.3.2 ICV Initialization

The following table 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

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 1.2.6 on page 12 for further details.

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

6 **Cross References:**

- 7 • **OMP_SCHEDULE** environment variable, see Section 4.1 on page 238.
- 8 • **OMP_NUM_THREADS** environment variable, see Section 4.2 on page 239.
- 9 • **OMP_DYNAMIC** environment variable, see Section 4.3 on page 240.
- 10 • **OMP_PROC_BIND** environment variable, see Section 4.4 on page 241.
- 11 • **OMP_PLACES** environment variable, see Section 4.5 on page 241.
- 12 • **OMP_NESTED** environment variable, see Section 4.6 on page 243.
- 13 • **OMP_STACKSIZE** environment variable, see Section 4.7 on page 244.
- 14 • **OMP_WAIT_POLICY** environment variable, see Section 4.8 on page 245.
- 15 • **OMP_MAX_ACTIVE_LEVELS** environment variable, see Section 4.9 on page 245.
- 16 • **OMP_THREAD_LIMIT** environment variable, see Section 4.10 on page 246.
- 17 • **OMP_CANCELLATION** environment variable, see Section 4.11 on page 246.
- 18 • **OMP_DEFAULT_DEVICE** environment variable, see Section 4.13 on page 248.

19 **2.3.3 Modifying and Retrieving ICV Values**

20 The following table shows the method for modifying and retrieving the values of ICVs
 21 through OpenMP API routines:

ICV	Ways to modify value	Way 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>

ICV	Ways to modify value	Way to retrieve value
<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_levels()</code>
<i>levels-var</i>	(none)	<code>omp_get_level()</code>
<i>place-partition-var</i>	(none)	(none)
<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>

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.

Cross References:

- `thread_limit` clause of the `teams` construct, see Section 2.9.5 on page 86.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 189.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 192.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 197.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 198.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 199.
- `omp_set_nested` routine, see Section 3.2.10 on page 200.
- `omp_get_nested` routine, see Section 3.2.11 on page 201.
- `omp_set_schedule` routine, see Section 3.2.12 on page 203.
- `omp_get_schedule` routine, see Section 3.2.13 on page 205.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 206.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 207.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 209.
- `omp_get_level` routine, see Section 3.2.17 on page 210.
- `omp_get_active_level` routine, see Section 3.2.20 on page 214.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 216
- `omp_set_default_device` routine, see Section 3.2.23 on page 218.
- `omp_get_default_device` routine, see Section 3.2.24 on page 219.

1 2.3.4 How ICVs are Scoped

2 The following table 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>	device
<i>default-device-var</i>	data environment

3 **Comments:**

- 4 • There is one copy per device of each ICV with device scope
- 5 • Each data environment has its own copies of ICVs with data environment scope
- 6 • Each implicit task has its own copy of ICVs with implicit task scope

7 Calls to OpenMP API routines retrieve or modify data environment scoped ICVs in the
8 data environment of their binding tasks.

9 2.3.4.1 How the Per-Data Environment ICVs Work

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

13 When a **task** construct is encountered, the generated task inherits the value of
14 *nthreads-var* from the generating task's *nthreads-var* value. When a **parallel**
15 construct is encountered, and the generating task's *nthreads-var* list contains a single

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

6 When a device construct is encountered, the new device data environment inherits the
 7 values of the data environment scoped ICVs from the enclosing device data environment
 8 of the device that will execute the region. If a **teams** construct with a **thread_limit**
 9 clause is encountered, the *thread-limit-var* ICV of the new device data environment is
 10 not inherited but instead is set to a value that is less than or equal to the value specified
 11 in the clause.

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

15 2.3.5 ICV Override Relationships

16 The override relationships among construct clauses and ICVs are shown in the following
 17 table:

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)
<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)

1
2
3
4
5

6
7
8
9
10
11

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 elements of the *bind-var* ICV; otherwise, the **proc_bind** clause has no effect.

Cross References:

- **parallel** construct, see Section 2.5 on page 44.
- **proc_bind** clause, Section 2.5 on page 44.
- **num_threads** clause, see Section 2.5.1 on page 47.
- Loop construct, see Section 2.7.1 on page 53.
- **schedule** clause, see Section 2.7.1.1 on page 59.

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.

C/C++

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.

C/C++

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++

1 2.5 parallel Construct

2 Summary

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

5 Syntax

6  The syntax of the `parallel` construct is as follows:

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

7 where *clause* is one of the following:

```
    if (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)
```

8 

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

```

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

2 where *clause* is one of the following:

```

    if (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)
    
```

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

4 Binding

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

7 Description

8 When a thread encounters a **parallel** construct, a team of threads is created to
9 execute the **parallel** region (see Section 2.5.1 on page 47 for more information about
10 how the number of threads in the team is determined, including the evaluation of the **if**
11 and **num_threads** clauses). The thread that encountered the **parallel** construct
12 becomes the master thread of the new team, with a thread number of zero for the
13 duration of the new **parallel** region. All threads in the new team, including the
14 master thread, execute the region. Once the team is created, the number of threads in the
15 team remains constant for the duration of that **parallel** region.

1 The optional **proc_bind** clause, described in Section 2.5.2 on page 49, specifies the
2 mapping of OpenMP threads to places within the current place partition, that is, within
3 the places listed in the *place-partition-var* ICV for the implicit task of the encountering
4 thread.

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

9 A set of implicit tasks, equal in number to the number of threads in the team, is
10 generated by the encountering thread. The structured block of the parallel construct
11 determines the code that will be executed in each implicit task. Each task is assigned to
12 a different thread in the team and becomes tied. The task region of the task being
13 executed by the encountering thread is suspended and each thread in the team executes
14 its implicit task. Each thread can execute a path of statements that is different from that
15 of the other threads.

16 The implementation may cause any thread to suspend execution of its implicit task at a
17 task scheduling point, and switch to execute any explicit task generated by any of the
18 threads in the team, before eventually resuming execution of the implicit task (for more
19 details see Section 2.11 on page 113).

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

23 If a thread in a team executing a **parallel** region encounters another **parallel**
24 directive, it creates a new team, according to the rules in Section 2.5.1 on page 47, and
25 it becomes the master of that new team.

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

31 Restrictions

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

- 33 • A program that branches into or out of a **parallel** region is non-conforming.
- 34 • A program must not depend on any ordering of the evaluations of the clauses of the
35 **parallel** directive, or on any side effects of the evaluations of the clauses.
- 36 • At most one **if** clause can appear on the directive.
- 37 • At most one **proc_bind** clause can appear on the directive.

1 • At most one `num_threads` clause can appear on the directive. The `num_threads`
2 expression must evaluate to a positive integer value.

C/C++

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

C/C++

Fortran

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

Fortran

Cross References

- 8
- 9 • `default`, `shared`, `private`, `firstprivate`, and `reduction` clauses, see
10 Section 2.14.3 on page 155.
 - 11 • `copyin` clause, see Section 2.14.4 on page 173.
 - 12 • `omp_get_thread_num` routine, see Section 3.2.4 on page 193.

2.5.1 Determining the Number of Threads for a parallel Region

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

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

1
2

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;
else let ThreadsRequested = value of the first element of nthreads-var;
let ThreadsAvailable = (thread-limit-var - ThreadsBusy + 1);
if (IfClauseValue = false)
then number of threads = 1;
else if (ActiveParRegions >= 1) and (nest-var = false)
then number of threads = 1;
else if (ActiveParRegions = max-active-levels-var)
then number of threads = 1;
else if (dyn-var = true) and (ThreadsRequested <= ThreadsAvailable)
then number of threads = [ 1 : ThreadsRequested ];
else if (dyn-var = true) and (ThreadsRequested > ThreadsAvailable)
then number of threads = [ 1 : ThreadsAvailable ];
else if (dyn-var = false) and (ThreadsRequested <= ThreadsAvailable)
then number of threads = ThreadsRequested;
else if (dyn-var = false) and (ThreadsRequested > ThreadsAvailable)
then behavior is implementation defined;
```

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

4 **Cross References**

- 5 • *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs, see
6 Section 2.3 on page 34.

7 **2.5.2 Controlling OpenMP Thread Affinity**

8 When creating a team for a parallel region, the **proc_bind** clause specifies a policy
9 for assigning OpenMP threads to places within the current place partition, that is, the
10 places listed in the *place-partition-var* ICV for the implicit task of the encountering
11 thread. Once a thread is assigned to a place, the OpenMP implementation should not
12 move it to another place.

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

17 The **close** thread affinity policy instructs the execution environment to assign the
18 threads to places close to the place of the parent thread. The master thread executes on
19 the parent's place and the remaining threads in the team execute on places from the
20 place list consecutive from the parent's position in the list, with wrap around with
21 respect to the place partition of the parent thread's implicit task. The place partition is
22 not changed by this policy, and each implicit task inherits the *place-partition-var* ICV of
23 the parent implicit task.

24 The purpose of the **spread** thread affinity policy is to create a sparse distribution for a
25 team of T threads among the P places of the parent's place partition. It accomplishes this
26 by first subdividing the parent partition into T subpartitions if T is less than or equal to
27 P , or P subpartitions if T is greater than P . Then it assigns one thread ($T \leq P$) or a set of
28 threads ($T > P$) to each subpartition. The *place-partition-var* ICV of each implicit task is
29 set to its subpartition. The subpartitioning is not only a mechanism for achieving a
30 sparse distribution, it also defines a subset of places for a thread to use when creating a
31 nested parallel region.

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- $T \leq P$. The parent's partition is split into T subpartitions, where each subpartition contains at least $S = \text{floor}(P/T)$ 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. For the other threads, assignment is to the first place in the corresponding subpartition. When T does not divide P evenly, the assignment of the remaining $P - T * S$ places into subpartitions is implementation defined.
 - $T > P$. The parent's partition is split into P unit-sized subpartitions. Each place is assigned $S = \text{floor}(T/P)$ threads. When P does not divide T evenly, the assignment of the remaining $T - P * S$ threads into places is implementation defined.
- 8
- 9
- 10

11 For the **close** and **spread** thread affinity policies, the threads with the smallest thread
12 numbers execute on the place of the master thread, then the threads with the next
13 smaller thread numbers execute on the next place in the partition; and so on, with wrap
14 around with respect to the encountering thread's place partition.

15 The determination of whether the affinity request can be fulfilled is implementation
16 defined. If the affinity request cannot be fulfilled, then the number of threads in the team
17 and their mapping to places are implementation defined.

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*

<i>init-expr</i>	One of the following: <i>var</i> = <i>lb</i> <i>integer-type var</i> = <i>lb</i> <i>random-access-iterator-type var</i> = <i>lb</i> <i>pointer-type var</i> = <i>lb</i>
<i>test-expr</i>	One of the following: <i>var relational-op b</i> <i>b relational-op var</i>
<i>incr-expr</i>	One of the following: ++ <i>var</i> <i>var</i> ++ -- <i>var</i> <i>var</i> -- <i>var</i> += <i>incr</i> <i>var</i> -= <i>incr</i> <i>var</i> = <i>var</i> + <i>incr</i> <i>var</i> = <i>incr</i> + <i>var</i> <i>var</i> = <i>var</i> - <i>incr</i>
<i>var</i>	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 <i>for-loop</i> other than in <i>incr-expr</i> . Unless the variable is specified lastprivate on the loop construct, its value after the loop is unspecified.
<i>relational-op</i>	One of the following: < <= > >=
<i>lb</i> and <i>b</i>	Loop invariant expressions of a type compatible with the type of <i>var</i> .
<i>incr</i>	A loop invariant integer expression.

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

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

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

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

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

18 The following restrictions also apply:

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

29 C/C++

2.7 Worksharing Constructs

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

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

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

- **loop** construct
- **sections** construct
- **single** construct
- **workshare** construct

Restrictions

The following restrictions apply to worksharing constructs:

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

2.7.1 Loop Construct

Summary

The loop construct specifies that the iterations of one or more associated loops will be executed in parallel by threads in the team in the context of their implicit tasks. The iterations are distributed across threads that already exist in the team executing the **parallel** region to which the loop 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)
reduction (reduction-identifier: list)
schedule (kind[, chunk_size])
collapse (n)
ordered
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 51).

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:

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

`schedule(kind[, chunk_size])`

`collapse(n)`

`ordered`

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

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

7 If any of the loop iteration variables would otherwise be shared, they are implicitly
8 made private on the loop construct. Unless the loop iteration variables are specified
9 **lastprivate** on the loop construct, their values after the loop are unspecified.



Fortran

10 Binding

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

15 Description

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

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

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

24 If more than one loop is associated with the loop construct, then the iterations of all
25 associated loops are collapsed into one larger iteration space that is then divided
26 according to the **schedule** clause. The sequential execution of the iterations in all
27 associated loops determines the order of the iterations in the collapsed iteration space.

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

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

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

17 Different loop regions with the same schedule and iteration count, even if they occur in
18 the same parallel region, can distribute iterations among threads differently. The only
19 exception is for the **static** schedule as specified in Table 2-1. Programs that depend
20 on which thread executes a particular iteration under any other circumstances are
21 non-conforming.

22 See Section 2.7.1.1 on page 59 for details of how the schedule for a worksharing loop is
23 determined.

24 The schedule *kind* can be one of those specified in Table 2-1.

TABLE 2-1 `schedule` clause *kind* values

static	<p>When <code>schedule(static, chunk_size)</code> is specified, iterations are divided into chunks of size <code>chunk_size</code>, and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.</p> <p>When no <code>chunk_size</code> 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.</p> <p>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 <code>chunk_size</code> specified, or both loop regions have no <code>chunk_size</code> 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.</p>
dynamic	<p>When <code>schedule(dynamic, chunk_size)</code> 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.</p> <p>Each chunk contains <code>chunk_size</code> iterations, except for the last chunk to be distributed, which may have fewer iterations.</p> <p>When no <code>chunk_size</code> is specified, it defaults to 1.</p>
guided	<p>When <code>schedule(guided, chunk_size)</code> 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.</p> <p>For a <code>chunk_size</code> 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 <code>chunk_size</code> with value <i>k</i> (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 <i>k</i> iterations (except for the last chunk to be assigned, which may have fewer than <i>k</i> iterations).</p> <p>When no <code>chunk_size</code> is specified, it defaults to 1.</p>
auto	<p>When <code>schedule(auto)</code> 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.</p>

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 `run-sched-var` ICV. If the ICV is set to **auto**, the schedule is implementation defined.

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 $2*p*k$.

Restrictions

Restrictions to the loop construct are as follows:

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

C/C++

- 1 • The associated *for-loops* must be structured blocks.
- 2 • Only an iteration of the innermost associated loop may be curtailed by a **continue**
- 3 statement.
- 4 • No statement can branch to any associated **for** statement.
- 5 • Only one **nowait** clause can appear on a **for** directive.
- 6 • A throw executed inside a loop region must cause execution to resume within the
- 7 same iteration of the loop region, and the same thread that threw the exception must
- 8 catch it.

C/C++

Fortran

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

Fortran

Cross References

- 16 • **private**, **firstprivate**, **lastprivate**, and **reduction** clauses, see
- 17 Section 2.14.3 on page 155.
- 18 • **OMP_SCHEDULE** environment variable, see Section 4.1 on page 238.
- 19 • **ordered** construct, see Section 2.12.8 on page 138.
- 20

2.7.1.1 Determining the Schedule of a Worksharing Loop

21 When execution encounters a loop directive, the **schedule** clause (if any) on the

22 directive, and the *run-sched-var* and *def-sched-var* ICVs are used to determine how loop

23 iterations are assigned to threads. See Section 2.3 on page 34 for details of how the

24 values of the ICVs are determined. If the loop directive does not have a **schedule**

25 clause then the current value of the *def-sched-var* ICV determines the schedule. If the

26 loop directive has a **schedule** clause that specifies the **runtime** schedule kind then

27 the current value of the *run-sched-var* ICV determines the schedule. Otherwise, the

28 value of the **schedule** clause determines the schedule. Figure 2-1 describes how the

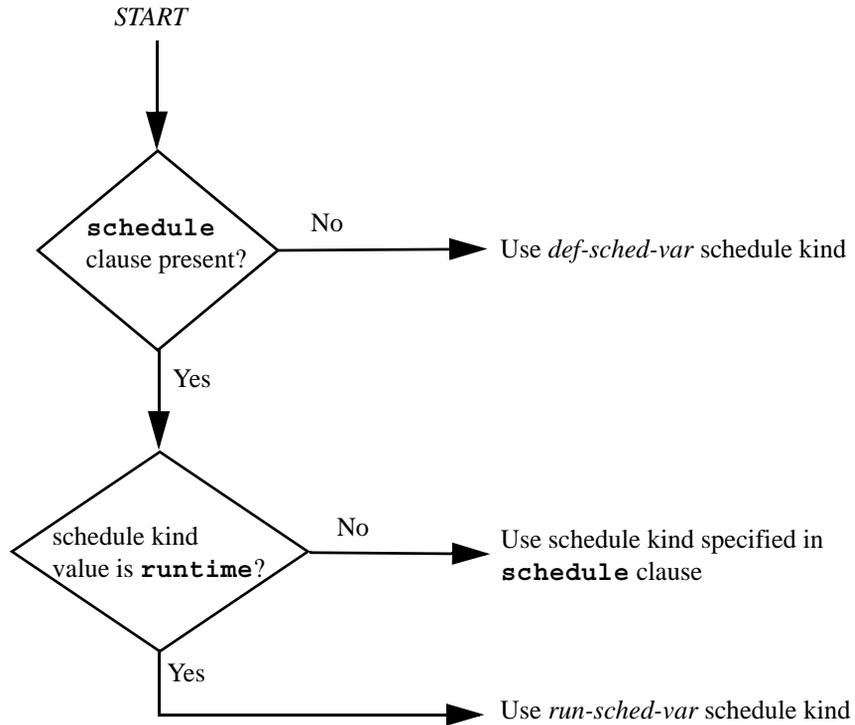
29 schedule for a worksharing loop is determined.

30

1
2
3

Cross References

- ICVs, see Section 2.3 on page 34.



4 **FIGURE 2-1** Determining the schedule for a worksharing loop.

5 **2.7.2 sections Construct**

6 **Summary**

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

1

Syntax

C/C++

2

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 ]
    ...
  }

```

3

where *clause* is one of the following:

```

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

```

4

C/C++

Fortran

5

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]

```

6

where *clause* is one of the following:

```

  private (list)

```

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

1



2

Binding

3

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.

4

5

6

7

8

Description

9

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.

10

11

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

12

13

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

14

15

Restrictions

16

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

17

- 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.

18

19

20

- The code enclosed in a **sections** construct must be a structured block.

21

- Only a single **nowait** clause can appear on a **sections** directive.

1
2
3

- 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.

4
5
6

Cross References

- **private**, **firstprivate**, **lastprivate**, and **reduction** clauses, see Section 2.14.3 on page 155.

7 2.7.3 single Construct

8

Summary

9
10
11
12

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.

13

Syntax

14

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

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

15

where *clause* is one of the following:

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

16

Fortran

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

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

2 where *clause* is one of the following:

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

3 and *end_clause* is one of the following:

```
copyprivate (list)
nowait
```

Fortran

5 Binding

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

11 Description

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

15 Restrictions

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

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

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

Cross References

- 3 • **private** and **firstprivate** clauses, see Section 2.14.3 on page 155.
4 • **copyprivate** clause, see Section 2.14.4.2 on page 175.
5

6 2.7.4 workshare Construct

7 Summary

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

11 Syntax

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

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

13 The enclosed structured block must consist of only the following:

- 14 • array assignments
15 • scalar assignments
16 • **FORALL** statements
17 • **FORALL** constructs
18 • **WHERE** statements
19 • **WHERE** constructs
20 • **atomic** constructs

- 1 • **critical** constructs
- 2 • **parallel** constructs
- 3 Statements contained in any enclosed **critical** construct are also subject to these
- 4 restrictions. Statements in any enclosed **parallel** construct are not restricted.

5 **Binding**

6 The binding thread set for a **workshare** region is the current team. A **workshare**

7 region binds to the innermost enclosing **parallel** region. Only the threads of the team

8 executing the binding **parallel** region participate in the execution of the units of

9 work and the implied barrier of the **workshare** region if the barrier is not eliminated

10 by a **nowait** clause.

11 **Description**

12 There is an implicit barrier at the end of a **workshare** construct unless a **nowait**

13 clause is specified.

14 An implementation of the **workshare** construct must insert any synchronization that is

15 required to maintain standard Fortran semantics. For example, the effects of one

16 statement within the structured block must appear to occur before the execution of

17 succeeding statements, and the evaluation of the right hand side of an assignment must

18 appear to complete prior to the effects of assigning to the left hand side.

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

- 20 • For array expressions within each statement, including transformational array
- 21 intrinsic functions that compute scalar values from arrays:
 - 22 • Evaluation of each element of the array expression, including any references to
 - 23 **ELEMENTAL** functions, is a unit of work.
 - 24 • Evaluation of transformational array intrinsic functions may be freely subdivided
 - 25 into any number of units of work.
- 26 • For an array assignment statement, the assignment of each element is a unit of work.
- 27 • For a scalar assignment statement, the assignment operation is a unit of work.
- 28 • For a **WHERE** statement or construct, the evaluation of the mask expression and the
- 29 masked assignments are each a unit of work.
- 30 • For a **FORALL** statement or construct, the evaluation of the mask expression,
- 31 expressions occurring in the specification of the iteration space, and the masked
- 32 assignments are each a unit of work.

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
5 into a SIMD loop (that is, multiple iterations of the loop can be executed concurrently
6 using SIMD instructions).

7 Syntax

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

9  C/C++

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

10 where *clause* is one of the following:

```
safelen(length)  
linear(list[:linear-step])  
aligned(list[:alignment])  
private(list)  
lastprivate(list)  
reduction(reduction-identifier:list)  
collapse(n)
```

11 The `simd` directive places restrictions on the structure of the associated *for-loops*.
12 Specifically, all associated *for-loops* must have *canonical loop form* (Section 2.6 on
13 page 51).

 C/C++

1

Fortran

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

2

where *clause* is one of the following:

```
safelen(length)
linear(list[:linear-step])
aligned(list[:alignment])
private(list)
lastprivate(list)
reduction(reduction-identifier:list)
collapse(n)
```

3

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

4

5

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.

6

7

8

Fortran

9

Binding

10

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

11

12

Description

13

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

14

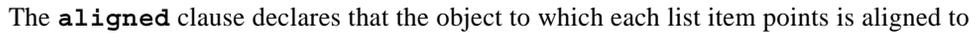
1 The **collapse** clause may be used to specify how many loops are associated with the
2 construct. The parameter of the **collapse** clause must be a constant positive integer
3 expression. If no **collapse** clause is present, the only loop that is associated with the
4 loop construct is the one that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

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

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

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

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

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

 C/C++ 

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

 Fortran 

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

30 Restrictions

- 31 • All loops associated with the construct must be perfectly nested; that is, there must be
32 no intervening code nor any OpenMP directive between any two loops.

- 1 • The associated loops must be structured blocks.
- 2 • A program that branches into or out of a **simd** region is non-conforming.
- 3 • Only one **collapse** clause can appear on a **simd** directive.
- 4 • A *list-item* cannot appear in more than one **aligned** clause.
- 5 • Only one **safelen** clause can appear on a **simd** directive.
- 6 • No OpenMP construct can appear in the **simd** region.

C/C++

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

C/C++

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

C

- 9 • The type of list items appearing in the **aligned** clause must be array, pointer,
10 reference to array, or reference to pointer.
- 11 • No exception can be raised in the **simd** region.

C++

C++

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

Fortran

Cross References

- 16 • **private**, **lastprivate**, **linear** and **reduction** clauses, see Section 2.14.3
17 on page 155.
- 18

1 2.8.2 declare simd construct

2 Summary

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

8 Syntax

9 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
```

10 where *clause* is one of the following:

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

▲ C/C++ ▲

▼ Fortran ▼

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

12

1 where *clause* is one of the following::

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

2  Fortran 

3 **Description**

4  C/C++ 

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

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

10  C/C++ 

11  Fortran 

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

15  Fortran 

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

18 If a SIMD version is created, the number of concurrent arguments for the function is
19 determined by the **simdlen** clause. If the **simdlen** clause is used its value
20 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.

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++

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

C/C++

Fortran

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

Fortran

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

8 The **inbranch** clause specifies that the function will always be called from inside a
9 conditional statement of a SIMD loop. The **notinbranch** clause specifies that the
10 function will never be called from inside a conditional statement of a SIMD loop. If
11 neither clause is specified, then the function may or may not be called from inside a
12 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.
- 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

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

C

C++

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

C++

Fortran

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

Fortran

Cross References

- 20 • **reduction** clause, see Section 2.14.3.6 on page 167.
- 21 • **linear** clause, see Section 2.14.3.7 on page 172.

1 2.8.3 Loop SIMD construct

2 Summary

3 The loop SIMD construct specifies a loop that can be executed concurrently using SIMD
4 instructions and that those iterations will also be executed in parallel by threads in the
5 team.

6 Syntax

▼ C/C++ ▼

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

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

▲ C/C++ ▲

▼ Fortran ▼

9

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

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

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

▲ Fortran ▲

14 Description

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

1
2
3
4

5
6
7
8

9

10

11
12

13

14

15
16
17
18

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.

Cross References

- loop construct, see Section 2.7.1 on page 53.
- `simd` construct, see Section 2.8.1 on page 68.
- Data attribute clauses, see Section 2.14.3 on page 155.



2.9 Device Constructs

2.9.1 target data Construct

Summary

Create 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:

- `device (integer-expression)`
- `map ([map-type :] list)`
- `if (scalar-expression)`

C/C++

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

```

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

```

2 where *clause* is one of the following:

3 **device** (*scalar-integer-expression*)

4 **map** ([*map-type* :] *list*)

5 **if** (*scalar-logical-expression*)

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

7 Binding

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

10 Description

11 When a **target data** construct is encountered, a new device data environment is
12 created, and the encountering task executes the **target data** region. If there is no
13 **device** clause, the default device is determined by the *default-device-var* ICV. The
14 new device data environment is constructed from the enclosing device data environment,
15 the data environment of the encountering task and any data-mapping clauses on the
16 construct. When an **if** clause is present and the **if** clause expression evaluates to *false*,
17 the device is the host.

18 Restrictions

- 19 • A program must not depend on any ordering of the evaluations of the clauses of the
20 **target data** directive, or on any side effects of the evaluations of the clauses.
- 21 • At most one **device** clause can appear on the directive. The **device** expression
22 must evaluate to a non-negative integer value.
- 23 • At most one **if** clause can appear on the directive.

Cross References

- `map` clause, see Section 2.14.5 on page 177.
- `default-device-var`, see Section 2.3 on page 34.

2.9.2 target Construct

Summary

Create a device data environment and execute the construct on the same 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:

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

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

```
if ( scalar-expression )
```

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:

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

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

```
if ( scalar-logical-expression )
```

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

Fortran

2 Binding

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

5 Description

6 The **target** construct provides a superset of the functionality and restrictions provided
7 by the **target data** directive. The functionality added to the **target** directive is the
8 inclusion of an executable region to be executed by a device. That is, the **target**
9 directive is an executable directive. The encountering task waits for the device to
10 complete the target region. When an **if** clause is present and the **if** clause expression
11 evaluates to *false*, the target region is executed by the host device.

12 Restrictions

- 13 • If a **target**, **target update**, or **target data** construct appears within a target
14 region then the behavior is unspecified.
- 15 • The result of an **omp_set_default_device**, **omp_get_default_device**,
16 or **omp_get_num_devices** routine called within a target region is unspecified.
- 17 • The effect of an access to a **threadprivate** variable in a target region is
18 unspecified.
- 19 • A variable referenced in a **target construct** that is not declared in the construct
20 is implicitly treated as if it had appeared in a **map** clause with a *map-type* of
21 **tofrom**.
- 22 • A variable referenced in a target region but not the target construct that is not
23 declared in the target region must appear in a **declare target** directive.

C++

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

C++

26 Cross References

- 27 • **target data** construct, see Section 2.9.1 on page 77.

- 1 • *default-device-var*, see Section 2.3 on page 34.
- 2 • **map** clause, see Section 2.14.5 on page 177.

3 2.9.3 target update Construct

4 Summary

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

8 Syntax

C/C++

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

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

10 where *motion-clause* is one of the following:

11 **to** (*list*)

12 **from** (*list*)

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

14 **device** (*integer-expression*)

15 **if** (*scalar-expression*)

C/C++

Fortran

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

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

17 where *motion-clause* is one of the following:

18 **to** (*list*)

19 **from** (*list*)

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

2 **device** (*scalar-integer-expression*)

3 **if** (*scalar-logical-expression*)

Fortran

4 **Binding**

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

7 **Description**

8 For each list item in a **to** or **from** clause there is a corresponding list item and an
9 original list item. If the corresponding list item is not present in the device data
10 environment, the behavior is unspecified. Otherwise, each corresponding list item in the
11 device data environment has an original list item in the current task's data environment.

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

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

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

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

20 **Restrictions**

- 21 • A program must not depend on any ordering of the evaluations of the clauses of the
- 22 **target update** directive, or on any side effects of the evaluations of the clauses.
- 23 • At least one *motion-clause* must be specified.
- 24 • If a list item is an array section it must specify contiguous storage.
- 25 • A variable that is part of another variable (such as a field of a structure) but is not an
- 26 array element or an array section cannot appear as a list item in a clause of a
- 27 **target update** construct.
- 28 • A list item can only appear in a **to** or **from** clause, but not both.
- 29 • A list item in a **to** or **from** clause must have a mappable type.

- 1 • At most one **device** clause can appear on the directive. The **device** expression
2 must evaluate to a non-negative integer value.
- 3 • At most one **if** clause can appear on the directive.

4 **Cross References**

- 5 • *default-device-var*, see Section 2.3 on page 34.
- 6 • **target data**, see Section 2.9.1 on page 77.
- 7 • Array sections, Section 2.4 on page 42

8 **2.9.4 declare target Directive**

9 **Summary**

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

13 **Syntax**

14  C/C++ 
The syntax of the **declare target** directive is as follows:

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

15  C/C++ 

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

17 For variables, functions and subroutines:

```
!$omp declare target( list )
```

1 where *list* is a comma-separated list of named variables, procedure names and named
2 common blocks. Common block names must appear between slashes.

3 For functions and subroutines:

```
!$omp declare target
```

4 Fortran

5 Description

6 C/C++

7 Variable and routine declarations that appear between the **declare target** and **end**
8 **declare target** directives form an implicit list where each list item is the variable
or function name.

9 C/C++

10 Fortran

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

12 Fortran

13 If a list item 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.

14 If a list item is a variable then the original variable is mapped to a corresponding
15 variable in the initial device data environment for all devices. If the original variable is
16 initialized, the corresponding variable in the device data environment is initialized with
17 the same value.

18 Restrictions

- 19 • A threadprivate variable cannot appear in a **declare target** directive.
- 20 • A variable declared in a **declare target** directive must have a mappable type.

21 C/C++

- 22 • A variable declared in a **declare target** directive must be at file or namespace
scope.
- 23 • A function declared in a **declare target** directive must be at file, namespace, or
24 class scope.

- 1 • All declarations and definitions for a function must have a **declare target**
2 directive if one is specified for any of them. Otherwise, the result is unspecified.

▲ C/C++ ▲

▼ Fortran ▼

- 3 • If a list item is a procedure name, it must not be a generic name, procedure pointer or
4 entry name.
- 5 • Any **declare target** directive with a list can only appear in a specification part
6 of a subroutine subprogram, function subprogram, program or module.
- 7 • Any **declare target** directive without a list can only appear in a specification
8 part of a subroutine subprogram, function subprogram or interface body to which it
9 applies.
- 10 • If a **declare target** directive is specified in an interface block for a procedure, it
11 must match a **declare target** directive in the definition of the procedure.
- 12 • If any procedure is declared via a procedure declaration statement, any **declare**
13 **target** directive with the procedure name must appear in the same specification
14 part.
- 15 • A variable that is part of another variable (as an array or structure element) cannot
16 appear in a **declare target** directive.
- 17 • The **declare target** directive must appear in the declaration section of a scoping
18 unit in which the common block or variable is declared. Although variables in
19 common blocks can be accessed by use association or host association, common
20 block names cannot. This means that a common block name specified in a **declare**
21 **target** directive must be declared to be a common block in the same scoping unit
22 in which the **declare target** directive appears.
- 23 • If a **declare target** directive specifying a common block name appears in one
24 program unit, then such a directive must also appear in every other program unit that
25 contains a **COMMON** statement specifying the same name. It must appear after the last
26 such **COMMON** statement in the program unit.
- 27 • If a **declare target** variable or a **declare target** common block is declared
28 with the **BIND** attribute, the corresponding C entities must also be specified in a
29 **declare target** directive in the C program.
- 30 • A blank common block cannot appear in a **declare target** directive.
- 31 • A variable can only appear in a **declare target** directive in the scope in which it
32 is declared. It must not be an element of a common block or appear in an
33 **EQUIVALENCE** statement.
- 34 • A variable that appears in a **declare target** directive must be declared in the
35 Fortran scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

▲ Fortran ▲

1 2.9.5 teams Construct

2 Summary

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

5 Syntax

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

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

7 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 )
```

8  

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

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

16 where *clause* is one of the following:

```
num_teams ( scalar-integer-expression )  
thread_limit ( scalar-integer-expression )
```

```
1      default ( shared | firstprivate | private | none )
2      private ( list )
3      firstprivate ( list )
4      shared ( list )
5      reduction ( reduction-identifier : list )
```

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

Fortran

7 **Binding**

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

9 **Description**

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

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

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

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

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

23 The threads other than the master thread do not begin execution until the master thread
24 encounters a **parallel** region.

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

27 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 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:

- *num_teams_var*, see Section 2.3.5 on page 40.
- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section 2.14.3 on page 155.
- **omp_get_num_teams** routine, see Section 3.2.26 on page 221.
- **omp_get_team_num** routine, see Section 3.2.27 on page 222.

2.9.6 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 51

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-loop*.

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

Fortran

Binding

5
6 The binding thread set for a **distribute** region is the set of master threads created by
7 a **teams** construct. A **distribute** region binds to the innermost enclosing **teams**
8 region. Only the threads executing the binding **teams** region participate in the
9 execution of the loop iterations.

Description

10
11 The **distribute** construct is associated with a loop nest consisting of one or more
12 loops that follow the directive.

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

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

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

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

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

Restrictions

30
31 Restrictions to the **distribute** construct are as follows:

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

3 **Cross References:**

- 4 • loop construct, see Section 2.7.1 on page 53.
- 5 • **teams** construct, see Section 2.9.5 on page 86.

6 **2.9.7 distribute simd Construct**

7 **Summary**

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

10 **Syntax**

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

▼ C/C++ ▼

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

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

▲ C/C++ ▲

▼ Fortran ▼

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

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

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

Fortran

3 Description

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

10 Restrictions

11 The restrictions for the **distribute** and **simd** constructs apply.

12 Cross References

- 13 • **simd** construct, see Section 2.8.1 on page 68.
- 14 • **distribute** construct, see Section 2.9.6 on page 88.
- 15 • Data attribute clauses, see Section 2.14.3 on page 155.

16 2.9.8 Distribute Parallel Loop Construct

17 Summary

18 The distribute parallel loop construct specifies a loop that can be executed in parallel by
19 multiple threads that are members of multiple teams.

20 Syntax

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

C/C++

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

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

▲ C/C++ ▲

3

▼ Fortran ▼

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

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

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

▲ Fortran ▲

8 Description

9 The distribute parallel loop construct will first distribute the iterations of the associated
10 loop(s) according to the semantics of the **distribute** construct and any clauses that
11 apply to the **distribute** construct. The resulting loops will then be distributed across
12 the threads contained within the **teams** region to which the **distribute** construct
13 binds in a manner consistent with any clauses that apply to the parallel loop construct.
14 The effect of any clause that applies to both the **distribute** and parallel loop
15 constructs is as if it were applied to both constructs separately.

16 Restrictions

17 The restrictions for the **distribute** and parallel loop constructs apply.

18 Cross References

- 19 • **distribute** construct, see Section 2.9.6 on page 88.
- 20 • Parallel loop construct, see Section 2.10.1 on page 95.
- 21 • Data attribute clauses, see Section Section 2.14.3 on page 155.

1 2.9.9 Distribute Parallel Loop SIMD Construct

2 Summary

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

6 Syntax

7  C/C++ 
The syntax of the distribute parallel loop SIMD construct is as follows:

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

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

 C/C++ 

10  Fortran 
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 ]
```

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

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

 Fortran 

15 Description

16 The distribute parallel loop SIMD construct will first distribute the iterations of the
17 associated loop(s) according to the semantics of the **distribute** construct and any
18 clauses that apply to the **distribute** construct. The resulting loops will then be
19 distributed across the threads contained within the **teams** region to which the

1 **distribute** construct binds in a manner consistent with any clauses that apply to the
2 parallel loop construct. The resulting chunks of iterations will then be converted to a
3 SIMD loop in a manner consistent with any clauses that apply to the **simd** construct.
4 The effect of any clause that applies to both the **distribute** and parallel loop SIMD
5 constructs is as if it were applied to both constructs separately.

6 **Restrictions**

7 The restrictions for the **distribute** and parallel loop SIMD constructs apply.

8 **Cross References**

- 9 • **distribute** construct, see Section 2.9.6 on page 88.
- 10 • Parallel loop SIMD construct, see Section 2.10.4 on page 100.
- 11 • Data attribute clauses, see Section Section 2.14.3 on page 155.



12 **2.10 Combined Constructs**

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

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

22 **2.10.1 Parallel Loop Construct**

23 **Summary**

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

1

Syntax

2

▼ C/C++ ▼

The syntax of the parallel loop construct is as follows:

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

3

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

4

▲ C/C++ ▲

5

▼ Fortran ▼

The syntax of the parallel loop construct is as follows:

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

6

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

7

8

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

9

10

▲ Fortran ▲

11

Description

12

▼ C/C++ ▼

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

13

▲ C/C++ ▲

14

▼ Fortran ▼

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

15

16

▲ Fortran ▲

1
2
3
4
5
6
7
8
9
10
11
12
13
14

Restrictions

The restrictions for the `parallel` construct and the loop construct apply.

Cross References

- `parallel` construct, see Section 2.5 on page 44.
- loop construct, see Section 2.7.1 on page 53.
- Data attribute clauses, see Section 2.14.3 on page 155.

2.10.2 `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

1 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
```

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

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

Fortran

6 Description

C/C++

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

C/C++

Fortran

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

Fortran

12 Restrictions

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

14 Cross References:

- 15 • **parallel** construct, see Section 2.5 on page 44.
- 16 • **sections** construct, see Section 2.7.2 on page 60.
- 17 • Data attribute clauses, see Section 2.14.3 on page 155.

1 2.10.3 parallel workshare Construct

2 Summary

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

5 Syntax

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

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

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

10 Description

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

14 Restrictions

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

16 Cross References

- 17 • **parallel** construct, see Section 2.5 on page 44.
- 18 • **workshare** construct, see Section 2.7.4 on page 65.
- 19 • Data attribute clauses, see Section 2.14.3 on page 155.

1 2.10.4 Parallel Loop SIMD Construct

2 Summary

3 The parallel loop SIMD construct is a shortcut for specifying a **parallel** construct
4 containing one loop SIMD construct and no other statement.

5 Syntax

C/C++

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

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

C/C++

8

Fortran

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

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

11 If an **end parallel do simd** directive is not specified, an **end parallel do**
12 **simd** directive is assumed at the end of the *do-loop*. **nowait** may not be specified on
13 an **end parallel do simd** directive.

Fortran

14 Description

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

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

Restrictions

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

Cross References

- `parallel` construct, see Section 2.5 on page 44.
- loop SIMD construct, see Section 2.8.3 on page 76.
- Data attribute clauses, see Section 2.14.3 on page 155.

2.10.5 target teams construct

Summary

The `target teams` construct is a shortcut for specifying a `target` construct containing a `teams` construct.

Syntax

The syntax of the target teams construct is as follows:

C/C++

```
#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

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

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

Fortran

3 Description

C/C++

4 The semantics are identical to explicitly specifying a **target** directive immediately
5 followed by a **teams** directive.

C/C++

Fortran

6 The semantics are identical to explicitly specifying a **target** directive immediately
7 followed by a **teams** directive, and an **end teams** directive immediately followed by
8 an **end target** directive.

Fortran

9 Restrictions

10 The restrictions for the **target** and **teams** constructs apply.

11 Cross References

- 12 • **target** construct, see Section 2.9.2 on page 79.
- 13 • **teams** construct, see Section 2.9.5 on page 86.
- 14 • Data attribute clauses, see Section 2.14.3 on page 155.

15 2.10.6 teams distribute Construct

16 Summary

17 The **teams distribute** construct is a shortcut for specifying a **teams** construct
18 containing a **distribute** construct.

Syntax

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

C/C++

```
#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

```
!$omp teams distribute [clause[[,] clause]...]
    do-loops
[ !$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 2.9.5 on page 86.
- **distribute** construct, see Section 2.9.6 on page 88.
- Data attribute clauses, see Section 2.14.3 on page 155.

1 2.10.7 teams distribute simd Construct

2 Summary

3 The **teams distribute simd** construct is a shortcut for specifying a **teams**
4 construct containing a **distribute simd** construct.

5 Syntax

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

▼ C/C++ ▼

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

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

▲ C/C++ ▲

9

▼ Fortran ▼

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

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

12 If an **end teams distribute** directive is not specified, an **end teams**
13 **distribute** 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
16 followed by a **distribute simd** directive. Some clauses are permitted on both
17 constructs.

Restrictions

The restrictions for the `teams` and `distribute simd` constructs apply.

Cross References

- `teams` construct, see Section 2.9.5 on page 86.
- `distribute simd` construct, see Section 2.9.7 on page 91.
- Data attribute clauses, see Section 2.14.3 on page 155.

2.10.8 target teams distribute Construct

Summary

The `target teams distribute` construct is a shortcut for specifying a `target` construct containing a `teams distribute` construct.

Syntax

The syntax of the `target teams distribute` construct is as follows:

C/C++

```
#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

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

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

▲ Fortran ▲

3 Description

4 The semantics are identical to explicitly specifying a **target** directive immediately
5 followed by a **teams distribute** directive.

6 Restrictions

7 The restrictions for the **target** and **teams distribute** constructs apply.

8 Cross References

- 9 • **target** construct, see Section 2.9.1 on page 77.
- 10 • **teams distribute** construct, see Section 2.10.6 on page 102.
- 11 • Data attribute clauses, see Section 2.14.3 on page 155.

12 2.10.9 target teams distribute simd Construct

13 Summary

14 The **target teams distribute simd** construct is a shortcut for specifying a
15 **target** construct containing a **teams distribute simd** construct.

16 Syntax

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

▼ C/C++ ▼

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

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

▲ C/C++ ▲

1

Fortran

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

2

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

3

4

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

5

Fortran

6

Description

7

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

8

9

Restrictions

10

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

11

Cross References

12

- **target** construct, see Section 2.9.1 on page 77

13

- **teams distribute simd** construct, see Section 2.10.7 on page 104.

14

- Data attribute clauses, see Section 2.14.3 on page 155.

15 2.10.10 Teams Distribute Parallel Loop Construct

16

Summary

17

The **teams distribute parallel loop** construct is a shortcut for specifying a **teams** construct containing a **distribute parallel loop** construct.

18

Syntax

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

C/C++

```
#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

```
!$omp teams distribute parallel do [clause[[,] clause]...]
      do-loops
[ !$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 2.9.5 on page 86.
- Distribute parallel loop construct, see Section 2.9.8 on page 92.
- Data attribute clauses, see Section 2.14.3 on page 155.

1 2.10.11 Target Teams Distribute Parallel Loop 2 Construct

3 Summary

4 The target teams distribute parallel loop construct is a shortcut for specifying a **target**
5 construct containing a teams distribute parallel loop construct.

6 Syntax

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

▶ C/C++ ◀

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

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

▶ C/C++ ◀

10

▶ Fortran ◀

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

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

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

▶ Fortran ◀

16 Description

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

Restrictions

The restrictions for the **target** and **teams distribute parallel loop** constructs apply.

Cross References

- **target** construct, see Section 2.9.2 on page 79.
- Distribute parallel loop construct, see Section 2.10.10 on page 107.
- Data attribute clauses, see Section 2.14.3 on page 155.

2.10.12 Teams Distribute Parallel Loop SIMD Construct

Summary

The **teams distribute parallel loop SIMD** construct is a shortcut for specifying a **teams** construct containing a distribute parallel loop SIMD construct.

Syntax

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

C/C++

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

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

C/C++

Fortran

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

1 where *clause* can be any of the clauses accepted by the **teams** or **distribute**
2 **parallel do simd** directives with identical meanings and restrictions.
3 If an **end teams distribute parallel do simd** directive is not specified, an
4 **end teams distribute parallel do simd** directive is assumed at the end of the
5 *do-loops*.

Description

The semantics are identical to explicitly specifying a **teams** directive immediately followed by a distribute parallel loop SIMD 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 SIMD constructs apply.

Cross References

- **teams** construct, see Section 2.9.5 on page 86.
- Distribute parallel loop SIMD construct, see Section 2.9.9 on page 94.
- Data attribute clauses, see Section 2.14.3 on page 155.

2.10.13 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.

Syntax

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

C/C++

```
#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

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

Cross References

- `target` construct, see Section 2.9.2 on page 79.
- Teams distribute parallel loop SIMD construct, see Section 2.10.12 on page 110.
- Data attribute clauses, see Section 2.14.3 on page 155.



2.11 Tasking Constructs

2.11.1 `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 (scalar-expression)
final (scalar-expression)
untied
default(shared | none)
mergeable
private(list)
firstprivate(list)
shared(list)
depend(dependence-type : list)
```

1



2



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

```

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

```

3

where *clause* is one of the following:

- if** (*scalar-logical-expression*)
- final** (*scalar-logical-expression*)
- untied**
- default**(**private** | **firstprivate** | **shared** | **none**)
- mergeable**
- private** (*list*)
- firstprivate** (*list*)
- shared** (*list*)
- depend** (*dependence-type* : *list*)

4



5

Binding

6

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

7

8

Description

9

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.

10

11

12

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

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

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

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

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

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

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

32 **Note** – When storage is shared by an explicit **task** region, it is the programmer's
33 responsibility to ensure, by adding proper synchronization, that the storage does not
34 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.

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++

Fortran

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

Fortran

2.11.1.1 depend Clause

Summary

The **depend** clause enforces additional constraints on the scheduling of tasks. These constraints establish dependences only between sibling tasks. The clause consists of a *dependence-type* with one or more list items.

Syntax

The syntax of the **depend** clause is as follows:

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

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:

1 The **in** *dependence-type*. The generated task will be a dependent task of all previously
2 generated sibling tasks that reference at least one of the list items in an **out** or **inout**
3 *dependence-type* list.

4 The **out** and **inout** *dependence-types*. The generated task will be a dependent task of
5 all previously generated sibling tasks that reference at least one of the list items in an
6 **in**, **out**, or **inout** *dependence-type* list.

7 The list items that appear in the **depend** clause may include array sections.

8 **Note** – The enforced task dependence establishes a synchronization of memory
9 accesses performed by a dependent task with respect to accesses performed by the
10 predecessor tasks. However, it is the responsibility of the programmer to synchronize
11 properly with respect to other concurrent accesses that occur outside of those tasks.

12 **Restrictions**

13 Restrictions to the **depend** clause are as follows:

- 14 • List items used in **depend** clauses of the same task or sibling tasks must indicate
15 identical storage or disjoint storage.
- 16 • List items used in **depend** clauses cannot be zero-length array sections.
- 17 • A variable that is part of another variable (such as a field of a structure) but is not an
18 array element or an array section cannot appear in a **depend** clause.

19 **Cross References**

- 20 • Array sections, Section 2.4 on page 42.
- 21 • Task scheduling constraints, Section 2.11.3 on page 118.

22 **2.11.2 taskyield Construct**

23 **Summary**

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

Syntax

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

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

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 2.11.3 on page 118.

2.11.3 Task Scheduling

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

- the point immediately following the generation of an explicit task

- 1 • after the point of completion of a **task** region
- 2 • in a **taskyield** region
- 3 • in a **taskwait** region
- 4 • at the end of a **taskgroup** region
- 5 • in an implicit and explicit **barrier** region
- 6 • the point immediately following the generation of a **target** region
- 7 • at the beginning and end of a **target data** region
- 8 • in a **target update** region

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

- 11 • begin execution of a tied task bound to the current team
- 12 • resume any suspended task region, bound to the current team, to which it is tied
- 13 • begin execution of an untied task bound to the current team
- 14 • resume any suspended untied task region bound to the current team.

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

17 *Task Scheduling Constraints* are as follows:

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

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

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

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

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

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

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

13 2.12 Master and Synchronization Constructs

14 OpenMP provides the following synchronization constructs:

- 15 • the **master** construct.
- 16 • the **critical** construct.
- 17 • the **barrier** construct.
- 18 • the **taskwait** construct.
- 19 • the **taskgroup** construct.
- 20 • the **atomic** construct.
- 21 • the **flush** construct.
- 22 • the **ordered** construct.

23 2.12.1 master Construct

24 Summary

25 The **master** construct specifies a structured block that is executed by the master thread
26 of the team.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16

Syntax

C/C++

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

```
#pragma omp master new-line
    structured-block
```

C/C++

Fortran

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

```
!$omp master
    structured-block
!$omp end master
```

Fortran

Binding

The binding thread set for a **master** region is the current team. A **master** region binds to the innermost enclosing **parallel** region. Only the master thread of the team executing the binding **parallel** region participates in the execution of the structured block of the **master** region.

Description

Other threads in the team do not execute the associated structured block. There is no implied barrier either on entry to, or exit from, the **master** construct.

Restrictions

C++

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

C++

1 2.12.2 critical Construct

2 Summary

3 The **critical** construct restricts execution of the associated structured block to a
4 single thread at a time.

5 Syntax

6  C/C++
The syntax of the **critical** construct is as follows:

```
#pragma omp critical [(name)] new-line  
    structured-block
```

7  C/C++

8  Fortran
The syntax of the **critical** construct is as follows:

```
!$omp critical [(name)]  
    structured-block  
!$omp end critical [(name)]
```

9  Fortran

10 Binding

11 The binding thread set for a **critical** region is all threads in the contention group.
12 Region execution is restricted to a single thread at a time among all threads in the
13 contention group, without regard to the team(s) to which the threads belong.

14 Description

15 An optional *name* may be used to identify the **critical** construct. All **critical**
16 constructs without a name are considered to have the same unspecified name. A thread
17 waits at the beginning of a **critical** region until no thread in the contention group is

1 executing a **critical** region with the same name. The **critical** construct enforces
2 exclusive access with respect to all **critical** constructs with the same name in all
3 threads in the contention group, not just those threads in the current team.

4  C/C++ 
5 Identifiers used to identify a **critical** construct have external linkage and are in a
6 name space that is separate from the name spaces used by labels, tags, members, and
ordinary identifiers.

7  C/C++ 
8  Fortran 
9 The names of **critical** constructs are global entities of the program. If a name
conflicts with any other entity, the behavior of the program is unspecified.

10 **Restrictions**

11  C++ 
12 • A throw executed inside a **critical** region must cause execution to resume within
the same **critical** region, and the same thread that threw the exception must catch
it.

13  C++ 
14  Fortran 
15 The following restrictions apply to the **critical** construct:
16 • If a *name* is specified on a **critical** directive, the same *name* must also be
17 specified on the **end critical** directive.
• If no *name* appears on the **critical** directive, no *name* can appear on the **end
critical** directive.

18  Fortran 

18 **2.12.3 barrier Construct**

19 **Summary**

20 The **barrier** construct specifies an explicit barrier at the point at which the construct
21 appears. The **barrier** construct is a stand-alone directive.

Syntax

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

```
#pragma omp barrier new-line
```

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

```
!$omp barrier
```

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.

Restrictions

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

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

1 2.12.4 `taskwait` Construct

2 Summary

3 The `taskwait` construct specifies a wait on the completion of child tasks of the
4 current task. The `taskwait` construct is a stand-alone directive.

5 Syntax

C/C++

6 The syntax of the `taskwait` construct is as follows:

```
#pragma omp taskwait newline
```

7

C/C++

Fortran

8 The syntax of the `taskwait` construct is as follows:

```
!$omp taskwait
```

9

Fortran

10 Binding

11 A `taskwait` region binds to the current task region. The binding thread set of the
12 `taskwait` region is the current team.

13 Description

14 The `taskwait` region includes an implicit task scheduling point in the current task
15 region. The current task region is suspended at the task scheduling point until all child
16 tasks that it generated before the `taskwait` region complete execution.

1 2.12.5 taskgroup Construct

2 Summary

3 The **taskgroup** construct specifies a wait on completion of child tasks of the current
4 task and their descendent tasks.

5 Syntax

6  C/C++ 
The syntax of the **taskgroup** construct is as follows:

```
#pragma omp taskgroup new-line  
structured-block
```

7  C/C++ 

8  Fortran 
The syntax of the **taskgroup** construct is as follows:

```
!$omp taskgroup  
structured-block  
!$omp end taskgroup
```

9  Fortran 

10 Binding

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

13 Description

14 When a thread encounters a **taskgroup** construct, it starts executing the region. There
15 is an implicit task scheduling point at the end of the **taskgroup** region. The current
16 task is suspended at the task scheduling point until all child tasks that it generated in the
17 **taskgroup** region and all of their descendent tasks complete execution.

Cross References

- Task scheduling, see Section 2.11.3 on page 118

2.12.6 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

C/C++

The syntax of the **atomic** construct takes either of the following forms:

```
#pragma omp atomic [read | write | update |  
capture] [seq_cst] new-line  
    expression-stmt
```

or:

```
#pragma omp atomic capture [seq_cst] new-line  
    structured-block
```

where *expression-stmt* is an expression statement with one of the following forms:

- If clause is **read**:
`v = x;`
- If clause is **write**:
`x = expr;`
- If clause is **update** or not present:
`x++;`
`x--;`
`++x;`
`--x;`
`x binop= expr;`
`x = x binop expr;`
`x = expr binop x;`

- 1 • If clause is **capture**:
- 2 `v = x++;`
- 3 `v = x--;`
- 4 `v = ++x;`
- 5 `v = --x;`
- 6 `v = x binop= expr;`
- 7 `v = x = x binop expr;`
- 8 `v = x = expr binop x;`

9 and where *structured-block* is a structured block with one of the following forms:

- 10 `{v = x; x binop= expr;}`
- 11 `{x binop= expr; v = x;}`
- 12 `{v = x; x = x binop expr;}`
- 13 `{v = x; x = expr binop x;}`
- 14 `{x = x binop expr; v = x;}`
- 15 `{x = expr binop x; v = x;}`
- 16 `{v = x; x = expr;}`
- 17 `{v = x; x++;}`
- 18 `{v = x; ++x;}`
- 19 `{++x; v = x;}`
- 20 `{x++; v = x;}`
- 21 `{v = x; x--;}`
- 22 `{v = x; --x;}`
- 23 `{--x; v = x;}`
- 24 `{x--; v = x;}`

25 In the preceding expressions:

- 26 • *x* and *v* (as applicable) are both *l-value* expressions with scalar type.
- 27 • During the execution of an atomic region, multiple syntactic occurrences of *x* must
- 28 designate the same storage location.
- 29 • Neither of *v* and *expr* (as applicable) may access the storage location designated by *x*.
- 30 • Neither of *x* and *expr* (as applicable) may access the storage location designated by *v*.
- 31 • *expr* is an expression with scalar type.
- 32 • *binop* is one of `+`, `*`, `-`, `/`, `&`, `^`, `|`, `<<`, or `>>`.
- 33 • *binop*, *binop=*, `++`, and `--` are not overloaded operators.
- 34 • The expression *x binop expr* must be numerically equivalent to *x binop (expr)*. This
- 35 requirement is satisfied if the operators in *expr* have precedence greater than *binop*,
- 36 or by using parentheses around *expr* or subexpressions of *expr*.

- 1 • The expression *expr binop x* must be numerically equivalent to *(expr) binop x*. This
2 requirement is satisfied if the operators in *expr* have precedence equal to or greater
3 than *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
4 • For forms that allow multiple occurrences of *x*, the number of times that *x* is
5 evaluated is unspecified.

▶ C/C++ ◀

▼ Fortran ▶

6 The syntax of the **atomic** construct takes any of the following forms:

```
!$omp atomic read [seq_cst]
    capture-statement
[!$omp end atomic]
```

7 or

```
!$omp atomic write [seq_cst]
    write-statement
[!$omp end atomic]
```

8 or

```
!$omp atomic [update] [seq_cst]
    update-statement
[!$omp end atomic]
```

9 or

```
!$omp atomic capture [seq_cst]
    update-statement
    capture-statement
!$omp end atomic
```

10 or

```
!$omp atomic capture [seq_cst]
    capture-statement
    update-statement
!$omp end atomic
```

1 or

```

!$omp atomic capture [seq_cst]
    capture-statement
    write-statement
!$omp end atomic
    
```

2 where *write-statement* has the following form (if clause is **write**):

3 $x = expr$

4 where *capture-statement* has the following form (if clause is **capture** or **read**):

5 $v = x$

6 and where *update-statement* has one of the following forms (if clause is **update**,
7 **capture**, or not present):

8 $x = x \text{ operator } expr$

9 $x = expr \text{ operator } x$

10 $x = \textit{intrinsic_procedure_name} (x, \textit{expr_list})$

11 $x = \textit{intrinsic_procedure_name} (\textit{expr_list}, x)$

12 In the preceding statements:

- 13 • x and v (as applicable) are both scalar variables of intrinsic type.
- 14 • x must not be an allocatable variable.
- 15 • During the execution of an atomic region, multiple syntactic occurrences of x must
16 designate the same storage location.
- 17 • None of v , $expr$ and $expr_list$ (as applicable) may access the same storage location as
18 x .
- 19 • None of x , $expr$ and $expr_list$ (as applicable) may access the same storage location as
20 v .
- 21 • $expr$ is a scalar expression.
- 22 • $expr_list$ is a comma-separated, non-empty list of scalar expressions. If
23 $\textit{intrinsic_procedure_name}$ refers to **IAND**, **IOR**, or **IEOR**, exactly one expression
24 must appear in $expr_list$.
- 25 • $\textit{intrinsic_procedure_name}$ is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- 26 • $operator$ is one of **+**, *****, **-**, **/**, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**

- 1 • The expression *x operator expr* must be numerically equivalent to *x operator (expr)*.
2 This requirement is satisfied if the operators in *expr* have precedence greater than
3 *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- 4 • The expression *expr operator x* must be mathematically equivalent to *(expr) operator*
5 *x*. This requirement is satisfied if the operators in *expr* have precedence equal to or
6 greater than *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- 7 • *intrinsic_procedure_name* must refer to the intrinsic procedure name and not to other
8 program entities.
- 9 • *operator* must refer to the intrinsic operator and not to a user-defined operator.
- 10 • All assignments must be intrinsic assignments.
- 11 • For forms that allow multiple occurrences of *x*, the number of times that *x* is
12 evaluated is unspecified.

Fortran

- 13 • In all **atomic** construct forms, the *seq_cst* clause and the clause that denotes the
14 type of the atomic construct can appear in any order. In addition, an optional comma
15 may be used to separate the clauses.

16 Binding

17 The binding thread set for an atomic region is all threads in the contention group.
18 **atomic** regions enforce exclusive access with respect to other **atomic** regions that
19 access the same storage location *x* among all threads in the contention group without
20 regard to the teams to which the threads belong.

21 Description

22 The **atomic** construct with the **read** clause forces an atomic read of the location
23 designated by *x* regardless of the native machine word size.

24 The **atomic** construct with the **write** clause forces an atomic write of the location
25 designated by *x* regardless of the native machine word size.

26 The **atomic** construct with the **update** clause forces an atomic update of the location
27 designated by *x* using the designated operator or intrinsic. Note that when no clause is
28 present, the semantics are equivalent to atomic update. Only the read and write of the
29 location designated by *x* are performed mutually atomically. The evaluation of *expr* or
30 *expr_list* need not be atomic with respect to the read or write of the location designated
31 by *x*. No task scheduling points are allowed between the read and the write of the
32 location designated by *x*.

1 The **atomic** construct with the **capture** clause forces an atomic update of the
2 location designated by *x* using the designated operator or intrinsic while also capturing
3 the original or final value of the location designated by *x* with respect to the atomic
4 update. The original or final value of the location designated by *x* is written in the
5 location designated by *v* depending on the form of the **atomic** construct structured
6 block or statements following the usual language semantics. Only the read and write of
7 the location designated by *x* are performed mutually atomically. Neither the evaluation
8 of *expr* or *expr_list*, nor the write to the location designated by *v* need be atomic with
9 respect to the read or write of the location designated by *x*. No task scheduling points
10 are allowed between the read and the write of the location designated by *x*.

11 Any **atomic** construct with a **seq_cst** clause forces the atomically performed
12 operation to include an implicit flush operation without a list.

13 **Note** – As with other implicit flush regions, Section 1.4.4 on page 20 reduces the
14 ordering that must be enforced. The intent is that, when the analogous operation exists
15 in C++11 or C11, a sequentially consistent **atomic** construct has the same semantics as
16 a **memory_order_seq_cst** atomic operation in C++11/C11. Similarly, a
17 non-sequentially consistent **atomic** construct has the same semantics as a
18 **memory_order_relaxed** atomic operation in C++11/C11.

19
20 Unlike non-sequentially consistent **atomic** constructs, sequentially consistent **atomic**
21 constructs preserve the interleaving (sequentially consistent) behavior of correct,
22 data-race-free programs. However, they are not designed to replace the **flush** directive
23 as a mechanism to enforce ordering for non-sequentially consistent **atomic** constructs,
24 and attempts to do so require extreme caution. For example, a sequentially consistent
25 **atomic write** construct may appear to be reordered with a subsequent
26 non-sequentially consistent **atomic write** construct, since such reordering would not
27 be observable by a correct program if the second write were outside an **atomic**
28 directive.

29 For all forms of the **atomic** construct, any combination of two or more of these
30 **atomic** constructs enforces mutually exclusive access to the locations designated by *x*.
31 To avoid race conditions, all accesses of the locations designated by *x* that could
32 potentially occur in parallel must be protected with an **atomic** construct.

33 **atomic** regions do not guarantee exclusive access with respect to any accesses outside
34 of **atomic** regions to the same storage location *x* even if those accesses occur during a
35 **critical** or **ordered** region, while an OpenMP lock is owned by the executing
36 task, or during the execution of a **reduction** clause.

37 However, other OpenMP synchronization can ensure the desired exclusive access. For
38 example, a barrier following a series of atomic updates to *x* guarantees that subsequent
39 accesses do not form a race with the atomic accesses.

1 A compliant implementation may enforce exclusive access between **atomic** regions
2 that update different storage locations. The circumstances under which this occurs are
3 implementation defined.

4 If the storage location designated by *x* is not size-aligned (that is, if the byte alignment
5 of *x* is not a multiple of the size of *x*), then the behavior of the **atomic** region is
6 implementation defined.

7 **Restrictions**

C/C++

8 The following restriction applies to the **atomic** construct:

- 9 • All atomic accesses to the storage locations designated by *x* throughout the program
10 are required to have a compatible type.

C/C++

Fortran

11 The following restriction applies to the **atomic** construct:

- 12 • All atomic accesses to the storage locations designated by *x* throughout the program
13 are required to have the same type and type parameters.

Fortran

14 **Cross References**

- 15 • **critical** construct, see Section 2.12.2 on page 122.
16 • **barrier** construct, see Section 2.12.3 on page 123.
17 • **flush** construct, see Section 2.12.7 on page 134.
18 • **ordered** construct, see Section 2.12.8 on page 138.
19 • **reduction** clause, see Section 2.14.3.6 on page 167.
20 • lock routines, see Section 3.3 on page 224.

1 2.12.7 flush Construct

2 Summary

3 The **flush** construct executes the OpenMP flush operation. This operation makes a
4 thread's temporary view of memory consistent with memory, and enforces an order on
5 the memory operations of the variables explicitly specified or implied. See the memory
6 model description in Section 1.4 on page 17 for more details. The **flush** construct is a
7 stand-alone directive.

8 Syntax

9  C/C++ 
The syntax of the **flush** construct is as follows:

```
#pragma omp flush [(list)] new-line
```

10  C/C++ 

11  Fortran 
The syntax of the **flush** construct is as follows:

```
!$omp flush [(list)]
```

12  Fortran 

13 Binding

14 The binding thread set for a **flush** region is the encountering thread. Execution of a
15 **flush** region affects the memory and the temporary view of memory of only the thread
16 that executes the region. It does not affect the temporary view of other threads. Other
17 threads must themselves execute a flush operation in order to be guaranteed to observe
18 the effects of the encountering thread's flush operation.

1
2
3
4
5
6
7

8
9

10
11
12
13
14
15
16

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

1
2
3
4
5
6
7
8
9
10

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

```
atomic(b = 1)
flush(b)
flush(a)
atomic(tmp = a)
if (tmp == 0) then
    protected section
end if
```

thread 2

```
atomic(a = 1)
flush(a)
flush(b)
atomic(tmp = b)
if (tmp == 0) then
    protected section
end if
```

11
12
13
14
15
16

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.

1 The following pseudocode example correctly ensures that the protected section is executed
2 by not more than one of the two threads at any one time. Notice that execution of the
3 protected section by neither thread is considered correct in this example. This occurs if
4 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
```

5 The compiler is prohibited from moving the flush at all for either thread, ensuring that the
6 respective assignment is complete and the data is flushed before the **if** statement is
7 executed.



8 A **flush** region without a list is implied at the following locations:

- 9 • During a barrier region.
- 10 • At entry to and exit from **parallel**, **critical**, and **ordered** regions.
- 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)
13 performed in a 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**
16 and **omp_test_nest_lock** regions, if the region causes the lock to be set or
17 unset.
- 18 • Immediately before and immediately after every task scheduling point.

19 A **flush** region with a list is implied at the following locations:

- 20 • At entry to and exit from the **atomic** operation (read, write, update, or capture)
21 performed in a non-sequentially consistent **atomic** region, where the list contains
22 only the storage location designated as *x* according to the description of the syntax of
23 the **atomic** construct in Section 2.12.6 on page 127.

24

1
2
3

Note – A **flush** region is not implied at the following locations:

- At entry to worksharing regions.
- At entry to or exit from a **master** region.

4 2.12.8 ordered Construct

5

Summary

6
7
8

The **ordered** construct specifies a structured block in a loop region that will be executed in the order of the loop iterations. This sequentializes and orders the code within an **ordered** region while allowing code outside the region to run in parallel.

9

Syntax

10

C/C++
The syntax of the **ordered** construct is as follows:

```
#pragma omp ordered new-line  
                  structured-block
```

11

C/C++
Fortran

12

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

```
!$omp ordered  
      structured-block  
!$omp end ordered
```

13

Fortran

1
2
3
4

5
6
7
8
9
10
11

12
13
14
15
16
17
18

19
20
21

22
23
24

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.

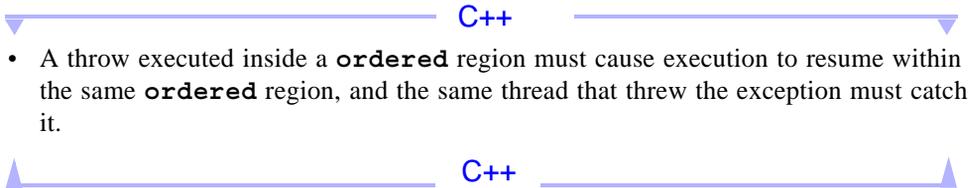
Description

The threads in the team executing the loop region execute **ordered** regions sequentially in the order of the loop iterations. 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 the **ordered** regions belonging to all previous iterations have completed.

Restrictions

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

- The loop region to which an **ordered** region binds must have an **ordered** clause specified on the corresponding loop (or parallel loop) 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 binds to the same loop region.



Cross References

- loop construct, see Section 2.7.1 on page 53.
- parallel loop construct, see Section 2.10.1 on page 95.

1 2.13 Cancellation Constructs

2 2.13.1 `cancel` Construct

3 Summary

4 The `cancel` construct activates cancellation of the innermost enclosing region of the
5 type specified. The `cancel` construct is a stand-alone directive.

6 Syntax

7  C/C++ 
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

9 `parallel`

10 `sections`

11 `for`

12 `taskgroup`

13 and *if-clause* is

14 `if (scalar-expression)`

 C/C++ 

 Fortran 

15 The syntax of the `cancel` construct is as follows:

```
!$omp cancel construct-type-clause[[,] if-clause] new-line
```

16 where *construct-type-clause* is one of the following

17 `parallel`

1 **sections**
2 **do**
3 **taskgroup**
4 and *if-clause* is
5 **if** (*scalar-logical-expression*)

Fortran

6 **Binding**

7 The binding thread set of the **cancel** region is the current team. The **cancel** region
8 binds to the innermost enclosing construct of the type corresponding to the *type-clause*
9 specified in the directive (that is, the innermost **parallel**, **sections**, **do**, or
10 **taskgroup** construct).

11 **Description**

12 The **cancel** construct activates cancellation of the binding construct only if *cancel-var*
13 is **true**, in which case the construct causes the encountering task to continue execution
14 at the end of the canceled construct. If *cancel-var* is **false**, the **cancel** construct is
15 ignored.

16 Threads check for active cancellation only at cancellation points. Cancellation points are
17 implied at the following locations:

- 18 • implicit barriers
- 19 • **barrier** regions
- 20 • **cancel** regions
- 21 • **cancellation point** regions

22 When a thread reaches one of the above cancellation points and if *cancel-var* is **true**,
23 the thread immediately checks for active cancellation (that is, if cancellation has been
24 activated by a **cancel** construct). If cancellation is active, the encountering thread
25 continues execution at the end of the canceled construct.

26 **Note** – If one thread activates cancellation and another thread encounters a cancellation
27 point, the absolute order of execution between the two threads is non-deterministic.
28 Whether the thread that encounters a cancellation point detects the activated cancellation
29 depends on the underlying hardware and operating system.

1 When cancellation of tasks is activated through the **cancel taskgroup** construct, the
2 innermost enclosing **taskgroup** will be canceled. The task that encountered the
3 **cancel taskgroup** construct continues execution at the end of its **task** region,
4 which implies completion of that task. Any task that belongs to the innermost enclosing
5 **taskgroup** and has already begun execution must run to completion or until a
6 cancellation point is reached. Upon reaching a cancellation point and if cancellation is
7 active, the task continues execution at the end of its **taskgroup** region, which implies
8 its completion. Any task that belongs to the innermost enclosing **taskgroup** and that
9 has not begun execution may be discarded, which implies its completion.

10 When cancellation is active for a **parallel**, **sections**, **for**, or **do** region, each
11 thread of the binding thread set resumes execution at the end of the canceled region if a
12 cancellation point is encountered. If the canceled region is a **parallel** region, any
13 tasks that have been created by a **task** construct and their descendent tasks are
14 canceled according to the above **taskgroup** cancellation semantics. If the canceled
15 region is a **sections**, **for**, or **do** region, no task cancellation occurs.

16 
The usual C++ rules for object destruction are followed when cancellation is performed.



17 
All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside
18 the canceled construct are deallocated.



19 **Note** – The user is responsible for releasing locks and similar data structures that might
20 cause a deadlock when a **cancel** construct is encountered and blocked threads cannot
21 be canceled.

22 If the canceled construct contains a **reduction** or **lastprivate** clause, the final
23 value of the **reduction** or **lastprivate** variable is undefined.

24 When an **if** clause is present on a **cancel** construct and the **if** expression evaluates
25 to *false*, the **cancel** construct does not activate cancellation. The cancellation point
26 associated with the **cancel** construct is always encountered regardless of the value of
27 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. 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** construct is not 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 35
- **cancellation point** construct, see Section 2.13.2 on page 143
- **omp_get_cancellation** routine, see Section 3.2.9 on page 199

2.13.2 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

C/C++

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

A **cancellation point** region binds to the current task region.

1
2
3
4
5
6
7
8
9

10
11
12
13
14
15
16
17

18
19
20
21

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** the task immediately checks whether cancellation of the region specified in the clause has been activated. If so, the encountering task continues execution at the end of the canceled construct.

Restrictions

- A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be closely nested inside a **task** construct. A **cancellation point** construct for which *construct-type-clause* is not **taskgroup** 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 35.
- **cancel** construct, see Section 2.13.1 on page 140.
- **omp_get_cancellation** routine, see Section 3.2.9 on page 199.

2.14 Data Environment

This section presents a directive and several clauses for controlling the data environment during the execution of **parallel**, **task**, **simd**, and worksharing regions.

- Section 2.14.1 on page 146 describes how the data-sharing attributes of variables referenced in **parallel**, **task**, **simd**, and worksharing regions are determined.
- The **threadprivate** directive, which is provided to create threadprivate memory, is described in Section 2.14.2 on page 150.
- Clauses that may be specified on directives to control the data-sharing attributes of variables referenced in **parallel**, **task**, **simd** or worksharing constructs are described in Section 2.14.3 on page 155.
- 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 2.14.4 on page 173.
- Clauses that may be specified on directives to map variables to devices are described in Section 2.14.5 on page 177.

2.14.1 Data-sharing Attribute Rules

This section describes how the data-sharing attributes of variables referenced in **parallel**, **task**, **simd**, and worksharing regions are determined. The following two cases are described separately:

- Section 2.14.1.1 on page 146 describes the data-sharing attribute rules for variables referenced in a construct.
- Section 2.14.1.2 on page 149 describes the data-sharing attribute rules for variables referenced in a region, but outside any construct.

2.14.1.1 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

1 construct may cause an implicit reference to the variable in the enclosing construct.
2 Such implicit references are also subject to the data-sharing attribute rules outlined in
3 this section.

4 Certain variables and objects have predetermined data-sharing attributes as follows:

C/C++

- 5 • Variables appearing in **threadprivate** directives are threadprivate.
- 6 • Variables with automatic storage duration that are declared in a scope inside the
7 construct are private.
- 8 • Objects with dynamic storage duration are shared.
- 9 • Static data members are shared.
- 10 • The loop iteration variable(s) in the associated *for-loop(s)* of a **for** or **parallel**
11 **for** construct is (are) private.
- 12 • The loop iteration variable in the associated *for-loop* of a **simd** construct with just
13 one associated *for-loop* is linear with a *constant-linear-step* that is the increment of
14 the associated *for-loop*.
- 15 • The loop iteration variables in the associated *for-loops* of a **simd** construct with
16 multiple associated *for-loops* are lastprivate.
- 17 • Variables with static storage duration that are declared in a scope inside the construct
18 are shared.

C/C++

Fortran

- 19 • Variables and common blocks appearing in **threadprivate** directives are
20 threadprivate.
- 21 • The loop iteration variable(s) in the associated *do-loop(s)* of a **do** or **parallel do**
22 construct is (are) private.
- 23 • The loop iteration variable in the associated *do-loop* of a **simd** construct with just
24 one associated *do-loop* is linear with a *constant-linear-step* that is the increment of
25 the associated *do-loop*.
- 26 • The loop iteration variables in the associated *do-loops* of a **simd** construct with
27 multiple associated *do-loops* are lastprivate.
- 28 • A loop iteration variable for a sequential loop in a **parallel** or **task** construct is
29 private in the innermost such construct that encloses the loop.
- 30 • Implied-do indices and **forall** indices are private.
- 31 • Cray pointees inherit the data-sharing attribute of the storage with which their Cray
32 pointers are associated.
- 33 • 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** or **parallel for** 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** or **parallel do** 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** or **task** 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 2.14.3 on page 155.

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.

1 Variables with *implicitly determined* data-sharing attributes are those that are referenced
2 in a given construct, do not have predetermined data-sharing attributes, and are not
3 listed in a data-sharing attribute clause on the construct.

4 Rules for variables with *implicitly determined* data-sharing attributes are as follows:

- 5 • In a **parallel** or **task** construct, the data-sharing attributes of these variables are
6 determined by the **default** clause, if present (see Section 2.14.3.1 on page 156).
- 7 • In a **parallel** construct, if no **default** clause is present, these variables are
8 shared.
- 9 • For constructs other than **task**, if no **default** clause is present, these variables
10 inherit their data-sharing attributes from the enclosing context.
- 11 • In a **task** construct, if no **default** clause is present, a variable that in the
12 enclosing context is determined to be shared by all implicit tasks bound to the current
13 team is shared.

Fortran

- 14 • In an orphaned **task** construct, if no **default** clause is present, dummy arguments
15 are **firstprivate**.

Fortran

- 16 • In a **task** construct, if no **default** clause is present, a variable whose data-sharing
17 attribute is not determined by the rules above is **firstprivate**.

18 Additional restrictions on the variables for which data-sharing attributes cannot be
19 implicitly determined in a **task** construct are described in Section 2.14.3.4 on page
20 162.

21 2.14.1.2 Data-sharing Attribute Rules for Variables Referenced in a 22 Region but not in a Construct

23 The data-sharing attributes of variables that are referenced in a region, but not in a
24 construct, are determined as follows:

C/C++

- 25 • Variables with static storage duration that are declared in called routines in the region
26 are shared.
- 27 • Variables with **const**-qualified type having no mutable member, and that are
28 declared in called routines, are shared.
- 29 • File-scope or namespace-scope variables referenced in called routines in the region
30 are shared unless they appear in a **threadprivate** directive.
- 31 • Objects with dynamic storage duration are shared.
- 32 • Static data members are shared unless they appear in a **threadprivate** directive.

- 1 • Formal arguments of called routines in the region that are passed by reference inherit
- 2 the data-sharing attributes of the associated actual argument.
- 3 • Other variables declared in called routines in the region are private.

▲ C/C++ ▼

▼ Fortran ▲

- 4 • Local variables declared in called routines in the region and that have the **save**
- 5 attribute, or that are data initialized, are shared unless they appear in a
- 6 **threadprivate** directive.
- 7 • Variables belonging to common blocks, or declared in modules, and referenced in
- 8 called routines in the region are shared unless they appear in a **threadprivate**
- 9 directive.
- 10 • Dummy arguments of called routines in the region that are passed by reference inherit
- 11 the data-sharing attributes of the associated actual argument.
- 12 • Cray pointees inherit the data-sharing attribute of the storage with which their Cray
- 13 pointers are associated.
- 14 • Implied-do indices, **forall** indices, and other local variables declared in called
- 15 routines in the region are private.

▲ Fortran ▼

16 2.14.2 threadprivate Directive

17 Summary

18 The **threadprivate** directive specifies that variables are replicated, with each thread

19 having its own copy. The **threadprivate** directive is a declarative directive.

20 Syntax

▼ C/C++ ▼

21 The syntax of the **threadprivate** directive is as follows:

```
#pragma omp threadprivate(list) new-line
```

22 where *list* is a comma-separated list of file-scope, namespace-scope, or static

23 block-scope variables that do not have incomplete types.

▲ C/C++ ▼

1 The syntax of the **threadprivate** directive is as follows:

```
!$omp threadprivate(list)
```

2 where *list* is a comma-separated list of named variables and named common blocks.
3 Common block names must appear between slashes.

4 Description

5 Each copy of a threadprivate variable is initialized once, in the manner specified by the
6 program, but at an unspecified point in the program prior to the first reference to that
7 copy. The storage of all copies of a threadprivate variable is freed according to how
8 static variables are handled in the base language, but at an unspecified point in the
9 program.

10 A program in which a thread references another thread's copy of a threadprivate variable
11 is non-conforming.

12 The content of a threadprivate variable can change across a task scheduling point if the
13 executing thread switches to another task that modifies the variable. For more details on
14 task scheduling, see Section 1.3 on page 14 and Section 2.11 on page 113.

15 In **parallel** regions, references by the master thread will be to the copy of the
16 variable in the thread that encountered the **parallel** region.

17 During a sequential part references will be to the initial thread's copy of the variable.
18 The values of data in the initial thread's copy of a threadprivate variable are guaranteed
19 to persist between any two consecutive references to the variable in the program.

20 The values of data in the threadprivate variables of non-initial threads are guaranteed to
21 persist between two consecutive active **parallel** regions only if all the following
22 conditions hold:

- 23 • Neither **parallel** region is nested inside another explicit **parallel** region.
- 24 • The number of threads used to execute both **parallel** regions is the same.
- 25 • The thread affinity policies used to execute both **parallel** regions are the same.
- 26 • The value of the *dyn-var* internal control variable in the enclosing task region is *false*
27 at entry to both **parallel** regions.

28 If these conditions all hold, and if a threadprivate variable is referenced in both regions,
29 then threads with the same thread number in their respective regions will reference the
30 same copy of that variable.

C/C++

1 If the above conditions hold, the storage duration, lifetime, and value of a thread's copy
2 of a `threadprivate` variable that does not appear in any `copyin` clause on the second
3 region will be retained. Otherwise, the storage duration, lifetime, and value of a thread's
4 copy of the variable in the second region is unspecified.

5 If the value of a variable referenced in an explicit initializer of a `threadprivate` variable
6 is modified prior to the first reference to any instance of the `threadprivate` variable, then
7 the behavior is unspecified.

C/C++

C++

8 The order in which any constructors for different `threadprivate` variables of class type
9 are called is unspecified. The order in which any destructors for different `threadprivate`
10 variables of class type are called is unspecified.

C++

Fortran

11 A variable is affected by a `copyin` clause if the variable appears in the `copyin` clause
12 or it is in a common block that appears in the `copyin` clause.

13 If the above conditions hold, the definition, association, or allocation status of a thread's
14 copy of a `threadprivate` variable or a variable in a `threadprivate` common
15 block, that is not affected by any `copyin` clause that appears on the second region, will
16 be retained. Otherwise, the definition and association status of a thread's copy of the
17 variable in the second region is undefined, and the allocation status of an allocatable
18 variable will be implementation defined.

19 If a `threadprivate` variable or a variable in a `threadprivate` common block is
20 not affected by any `copyin` clause that appears on the first `parallel` region in which
21 it is referenced, the variable or any subobject of the variable is initially defined or
22 undefined according to the following rules:

- 23 • If it has the **ALLOCATABLE** attribute, each copy created will have an initial
24 allocation status of not currently allocated.
- 25 • If it has the **POINTER** attribute:
 - 26 • if it has an initial association status of disassociated, either through explicit
27 initialization or default initialization, each copy created will have an association
28 status of disassociated;
 - 29 • otherwise, each copy created will have an association status of undefined.
- 30 • If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
 - 31 • if it is initially defined, either through explicit initialization or default
32 initialization, each copy created is so defined;

- 1
- otherwise, each copy created is undefined.

Fortran

Restrictions

2 The restrictions to the **threadprivate** directive are as follows:

- 3
- 4 • A **threadprivate** variable must not appear in any clause except the **copyin**,
 - 5 **copyprivate**, **schedule**, **num_threads**, **thread_limit**, and **if** clauses.
 - 6 • A program in which an untied task accesses **threadprivate** storage is non-conforming.

C/C++

- 7
- 8 • A variable that is part of another variable (as an array or structure element) cannot
 - 9 appear in a **threadprivate** clause unless it is a static data member of a C++
 - 10 class.
 - 11 • A **threadprivate** directive for file-scope variables must appear outside any
 - 12 definition or declaration, and must lexically precede all references to any of the
 - 13 variables in its list.
 - 14 • A **threadprivate** directive for namespace-scope variables must appear outside
 - 15 any definition or declaration other than the namespace definition itself, and must
 - 16 lexically precede all references to any of the variables in its list.
 - 17 • Each variable in the list of a **threadprivate** directive at file, namespace, or class
 - 18 scope must refer to a variable declaration at file, namespace, or class scope that
 - 19 lexically precedes the directive.
 - 20 • A **threadprivate** directive for static block-scope variables must appear in the
 - 21 scope of the variable and not in a nested scope. The directive must lexically precede
 - 22 all references to any of the variables in its list.
 - 23 • Each variable in the list of a **threadprivate** directive in block scope must refer to
 - 24 a variable declaration in the same scope that lexically precedes the directive. The
 - 25 variable declaration must use the static storage-class specifier.
 - 26 • If a variable is specified in a **threadprivate** directive in one translation unit, it
 - 27 must be specified in a **threadprivate** directive in every translation unit in which
 - 28 it is declared.
 - The address of a **threadprivate** variable is not an address constant.

C/C++

C++

- 29
- 30 • A **threadprivate** directive for static class member variables must appear in the
 - 31 class definition, in the same scope in which the member variables are declared, and
 - 32 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/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 ▲

Cross References:

- *dyn-var* ICV, see Section 2.3 on page 34.
- number of threads used to execute a `parallel` region, see Section 2.5.1 on page 47.
- `copyin` clause, see Section 2.14.4.1 on page 173.

1 2.14.3 Data-Sharing Attribute Clauses

2 Several constructs accept clauses that allow a user to control the data-sharing attributes
3 of variables referenced in the construct. Data-sharing attribute clauses apply only to
4 variables for which the names are visible in the construct on which the clause appears.

5 Not all of the clauses listed in this section are valid on all directives. The set of clauses
6 that is valid on a particular directive is described with the directive.

7 Most of the clauses accept a comma-separated list of list items (see Section 2.1 on page
8 26). All list items appearing in a clause must be visible, according to the scoping rules
9 of the base language. With the exception of the **default** clause, clauses may be
10 repeated as needed. A list item that specifies a given variable may not appear in more
11 than one clause on the same directive, except that a variable may be specified in both
12 **firstprivate** and **lastprivate** clauses.

▼ C++ ▼

13 If a variable referenced in a data-sharing attribute clause has a type derived from a
14 template, and there are no other references to that variable in the program, then any
15 behavior related to that variable is unspecified.

▲ C++ ▲

▼ Fortran ▼

16 A named common block may be specified in a list by enclosing the name in slashes.
17 When a named common block appears in a list, it has the same meaning as if every
18 explicit member of the common block appeared in the list. An explicit member of a
19 common block is a variable that is named in a **COMMON** statement that specifies the
20 common block name and is declared in the same scoping unit in which the clause
21 appears.

22 Although variables in common blocks can be accessed by use association or host
23 association, common block names cannot. As a result, a common block name specified
24 in a data-sharing attribute clause must be declared to be a common block in the same
25 scoping unit in which the data-sharing attribute clause appears.

26 When a named common block appears in a **private**, **firstprivate**,
27 **lastprivate**, or **shared** clause of a directive, none of its members may be declared
28 in another data-sharing attribute clause in that directive. When individual members of a
29 common block appear in a **private**, **firstprivate**, **lastprivate**, or
30 **reduction** clause of a directive, the storage of the specified variables is no longer
31 associated with the storage of the common block itself.

▲ Fortran ▲

1 2.14.3.1 default clause

2 Summary

3 The **default** clause explicitly determines the data-sharing attributes of variables that
4 are referenced in a **parallel**, **task** or **teams** construct and would otherwise be
5 implicitly determined (see Section 2.14.1.1 on page 146).

6 Syntax

7  C/C++
The syntax of the **default** clause is as follows:

```
default(shared | none)
```

8  C/C++

9  Fortran

10 The syntax of the **default** clause is as follows:

```
default(private | firstprivate | shared | none)
```

11  Fortran

12 Description

13 The **default(shared)** clause causes all variables referenced in the construct that
14 have implicitly determined data-sharing attributes to be shared.

15  Fortran

16 The **default(firstprivate)** clause causes all variables in the construct that have
implicitly determined data-sharing attributes to be firstprivate.

17 The **default(private)** clause causes all variables referenced in the construct that
18 have implicitly determined data-sharing attributes to be private.

 Fortran

1 The **default (none)** clause requires that each variable that is referenced in the
2 construct, and that does not have a predetermined data-sharing attribute, must have its
3 data-sharing attribute explicitly determined by being listed in a data-sharing attribute
4 clause.

5 **Restrictions**

6 The restrictions to the **default** clause are as follows:

- 7 • Only a single default clause may be specified on a **parallel**, **task**, or **teams**
8 directive.

9 **2.14.3.2 shared clause**

10 **Summary**

11 The **shared** clause declares one or more list items to be shared by tasks generated by
12 a **parallel**, **task** or **teams** construct.

13 **Syntax**

14 The syntax of the **shared** clause is as follows:

```
shared (list)
```

15 **Description**

16 All references to a list item within a task refer to the storage area of the original variable
17 at the point the directive was encountered.

18 It is the programmer's responsibility to ensure, by adding proper synchronization, that
19 storage shared by an explicit **task** region does not reach the end of its lifetime before
20 the explicit **task** region completes its execution.

Fortran

21 The association status of a shared pointer becomes undefined upon entry to and on exit
22 from the **parallel**, **task** or **teams** construct if it is associated with a target or a
23 subobject of a target that is in a **private**, **firstprivate**, **lastprivate**, or
24 **reduction** clause inside the construct.

1 Under certain conditions, passing a shared variable to a non-intrinsic procedure may
2 result in the value of the shared variable being copied into temporary storage before the
3 procedure reference, and back out of the temporary storage into the actual argument
4 storage after the procedure reference. It is implementation defined when this situation
5 occurs.

6 **Note** – Use of intervening temporary storage may occur when the following three
7 conditions hold regarding an actual argument in a reference to a non-intrinsic procedure:

8 a. The actual argument is one of the following:

- 9 • A shared variable.
- 10 • A subobject of a shared variable.
- 11 • An object associated with a shared variable.
- 12 • An object associated with a subobject of a shared variable.

13 b. The actual argument is also one of the following:

- 14 • An array section.
- 15 • An array section with a vector subscript.
- 16 • An assumed-shape array.
- 17 • A pointer array.

18 c. The associated dummy argument for this actual argument is an explicit-shape array
19 or an assumed-size array.

20 These conditions effectively result in references to, and definitions of, the temporary
21 storage during the procedure reference. Any references to (or definitions of) the shared
22 storage that is associated with the dummy argument by any other task must be
23 synchronized with the procedure reference to avoid possible race conditions.

24 **Fortran**

25 Restrictions

26 The restrictions for the **shared** clause are as follows:

- 27 **C/C++**
- 28 • A variable that is part of another variable (as an array or structure element) cannot appear in a **shared** clause unless it is a static data member of a C++ class.

C/C++

- 1 • A variable that is part of another variable (as an array or structure element) cannot
2 appear in a **shared** clause.

3 2.14.3.3 **private** clause

4 **Summary**

5 The **private** clause declares one or more list items to be private to a task or to a
6 SIMD lane.

7 **Syntax**

8 The syntax of the **private** clause is as follows:

<code>private (<i>list</i>)</code>

9 **Description**

10 Each task that references a list item that appears in a **private** clause in any statement
11 in the construct receives a new list item. Each SIMD lane used in a **simd** construct that
12 references a list item that appears in a private clause in any statement in the construct
13 receives a new list item. Language-specific attributes for new list items are derived from
14 the corresponding original list item. Inside the construct, all references to the original
15 list item are replaced by references to the new list item. In the rest of the region, it is
16 unspecified whether references are to the new list item or the original list item.
17 Therefore, if an attempt is made to reference the original item, its value after the region
18 is also unspecified. If a SIMD construct or a task does not reference a list item that
19 appears in a **private** clause, it is unspecified whether SIMD lanes or the task receive
20 a new list item.

21 The value and/or allocation status of the original list item will change only:

- 22 • if accessed and modified via pointer,
23 • if possibly accessed in the region but outside of the construct,
24 • 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++

A new list item of the same type, with automatic storage duration, is allocated for the construct. The storage and thus lifetime of these list items lasts until the block in which they are created exits. The size and alignment of the new list item are determined by the type of the variable. This allocation occurs once for each task generated by the construct and/or once for each SIMD lane used by the construct.

The new list item is initialized, or has an undefined initial value, as if it had been locally declared without an initializer.

C/C++

C++

The order in which any default constructors for different private variables of class type are called is unspecified. The order in which any destructors for different private variables of class type are called is unspecified.

C++

Fortran

If any statement of the construct references a list item, a new list item of the same type and type parameters is allocated: once for each implicit task in the **parallel** construct; once for each task generated by a **task** construct; and once for each SIMD lane used by a **simd** construct. The initial value of the new list item is undefined. Within a **parallel**, **worksharing**, **task**, **teams**, or **simd** region, the initial status of a private pointer is undefined.

For a list item or the subobject of a list item with the **ALLOCATABLE** attribute:

- if the allocation status is "not currently allocated", the new list item or the subobject of the new list item will have an initial allocation status of "not currently allocated";

- 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

2.14.3.4 **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 2.14.3.3 on page 159, 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 **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.

1 To avoid race conditions, concurrent updates of the original list item must be
2 synchronized with the read of the original list item that occurs as a result of the
3 **firstprivate** clause.

4 If a list item appears in both **firstprivate** and **lastprivate** clauses, the update
5 required for **lastprivate** occurs after all the initializations for **firstprivate**.

C/C++

6 For variables of non-array type, the initialization occurs by copy assignment. For an
7 array of elements of non-array type, each element is initialized as if by assignment from
8 an element of the original array to the corresponding element of the new array.

C/C++

C++

9 For variables of class type, a copy constructor is invoked to perform the initialization.
10 The order in which copy constructors for different variables of class type are called is
11 unspecified.

C++

Fortran

12 If the original list item does not have the **POINTER** attribute, initialization of the new
13 list items occurs as if by intrinsic assignment, unless the original list item has the
14 allocation status of not currently allocated, in which case the new list items will have the
15 same status.

16
17 If the original list item has the **POINTER** attribute, the new list items receive the same
18 association status of the original list item as if by pointer assignment.

Fortran

19 Restrictions

20 The restrictions to the **firstprivate** clause are as follows:

- 21 • A variable that is part of another variable (as an array or structure element) cannot
22 appear in a **firstprivate** clause.
- 23 • A list item that is private within a **parallel** region must not appear in a
24 **firstprivate** clause on a worksharing construct if any of the worksharing
25 regions arising from the worksharing construct ever bind to any of the **parallel**
26 regions arising from the **parallel** construct.

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- A list item that is private within a **teams** region must not appear in a **firstprivate** clause on a **distribute** construct if any of the **distribute** regions arising from the **distribute** construct ever bind to any of the **teams** regions arising from the **teams** construct.
 - A list item that appears in a **reduction** clause of a **parallel** construct must not appear in a **firstprivate** clause on a worksharing or **task** construct if any of the worksharing or **task** regions arising from the worksharing or **task** construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
 - A list item that appears in a **reduction** clause of a **teams** construct must not appear in a **firstprivate** clause on a **distribute** construct if any of the **distribute** regions arising from the **distribute** construct ever bind to any of the **teams** regions arising from the **teams** construct.
 - A list item that appears in a **reduction** clause in a worksharing construct must not appear in a **firstprivate** clause in a task construct encountered during execution of any of the worksharing regions arising from the worksharing construct.

▼ C++ ▼

- 16
- 17
- A variable of class type (or array thereof) that appears in a **firstprivate** clause requires an accessible, unambiguous copy constructor for the class type.

▲ C++ ▲

▼ C/C++ ▼

- 18
- 19
- A variable that appears in a **firstprivate** clause must not have an incomplete type or a reference type.

▲ C/C++ ▲

▼ Fortran ▼

- 20
- 21
- 22
- Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions, may not appear in a **firstprivate** clause.

▲ Fortran ▲

23 **2.14.3.5 lastprivate clause**

24 **Summary**

25 The **lastprivate** clause declares one or more list items to be private to an implicit
26 task or to a SIMD lane, and causes the corresponding original list item to be updated
27 after the end of the region.

Syntax

The syntax of the `lastprivate` clause is as follows:

```
lastprivate (list)
```

Description

The `lastprivate` clause provides a superset of the functionality provided by the `private` clause.

A list item that appears in a `lastprivate` clause is subject to the `private` clause semantics described in Section 2.14.3.3 on page 159. In addition, when a `lastprivate` clause appears on the directive that identifies a worksharing construct or a SIMD construct, the value of each new list item from the sequentially last iteration of the associated loops, or the lexically last `section` construct, is assigned to the original list item.

C/C++
For an array of elements of non-array type, each element is assigned to the corresponding element of the original array.

C/C++
Fortran

If the original list item does not have the `POINTER` attribute, its update occurs as if by intrinsic assignment.

If the original list item has the `POINTER` attribute, its update occurs as if by pointer assignment.

Fortran

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.

1 If a list item appears in both **firstprivate** and **lastprivate** clauses, the update
2 required for **lastprivate** occurs after all initializations for **firstprivate**.

3 **Restrictions**

4 The restrictions to the **lastprivate** clause are as follows:

- 5 • A variable that is part of another variable (as an array or structure element) cannot
6 appear in a **lastprivate** clause.
- 7 • A list item that is private within a **parallel** region, or that appears in the
8 **reduction** clause of a **parallel** construct, must not appear in a **lastprivate**
9 clause on a worksharing construct if any of the corresponding worksharing regions
10 ever binds to any of the corresponding **parallel** regions.

C++

- 11 • A variable of class type (or array thereof) that appears in a **lastprivate** clause
12 requires an accessible, unambiguous default constructor for the class type, unless the
13 list item is also specified in a **firstprivate** clause.
- 14 • A variable of class type (or array thereof) that appears in a **lastprivate** clause
15 requires an accessible, unambiguous copy assignment operator for the class type. The
16 order in which copy assignment operators for different variables of class type are
17 called is unspecified.

C/C++

C/C++

- 18 • A variable that appears in a **lastprivate** clause must not have a **const**-qualified
19 type unless it is of class type with a **mutable** member.
- 20 • A variable that appears in a **lastprivate** clause must not have an incomplete type
21 or a reference type.

C/C++

Fortran

- 22 • A variable that appears in a **lastprivate** clause must be definable.
- 23 • An original list item with the **ALLOCATABLE** attribute in the sequentially last
24 iteration or lexically last section must have an allocation status of allocated upon exit
25 from that iteration or section.
- 26 • Variables that appear in namelist statements, in variable format expressions, and in
27 expressions for statement function definitions, may not appear in a **lastprivate**
28 clause.

Fortran

1 2.14.3.6 reduction clause

2 Summary

3 The **reduction** clause specifies a *reduction-identifier* and one or more list items. For
4 each list item, a private copy is created in each implicit task or SIMD lane, and is
5 initialized with the initializer value of the *reduction-identifier*. After the end of the
6 region, the original list item is updated with the values of the private copies using the
7 combiner associated with the *reduction-identifier*.

8 Syntax

9 C/C++
The syntax of the **reduction** clause is as follows:

`reduction (reduction-identifier : list)`

10 where:

11 C
reduction-identifier is either an *identifier* or one of the following operators: +, -, *,
12 &, |, ^, && and ||

13 C++
reduction-identifier is either an *id-expression* or one of the following operators: +, -,
14 *, &, |, ^, && and ||

15 C++

16 The following table lists each *reduction-identifier* that is implicitly declared at every
17 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
+	<code>omp_priv = 0</code>	<code>omp_out += omp_in</code>
*	<code>omp_priv = 1</code>	<code>omp_out *= omp_in</code>
-	<code>omp_priv = 0</code>	<code>omp_out += omp_in</code>
&	<code>omp_priv = ~0</code>	<code>omp_out &= omp_in</code>

	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
min	omp_priv = <i>Largest representable number in the reduction list item type</i>	omp_out = omp_in < omp_out ? omp_in : omp_out

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`.

The following table 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
<code>+</code>	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
<code>*</code>	<code>omp_priv = 1</code>	<code>omp_out = omp_in * omp_out</code>
<code>-</code>	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
<code>.and.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .and. omp_out</code>
<code>.or.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .or. omp_out</code>
<code>.eqv.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .eqv. omp_out</code>
<code>.neqv.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .neqv. omp_out</code>

max	<code>omp_priv = <i>Least representable number in the reduction list item type</i></code>	<code>omp_out = max(omp_in, omp_out)</code>
min	<code>omp_priv = <i>Largest representable number in the reduction list item type</i></code>	<code>omp_out = min(omp_in, omp_out)</code>
iand	<code>omp_priv = All bits on</code>	<code>omp_out = iand(omp_in, omp_out)</code>
ior	<code>omp_priv = 0</code>	<code>omp_out = ior(omp_in, omp_out)</code>
ieor	<code>omp_priv = 0</code>	<code>omp_out = ieor(omp_in, omp_out)</code>

1

Fortran

2

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.

3

4

5

6

Description

7

The reduction clause can be used to perform some forms of recurrence calculations (involving mathematically associative and commutative operators) in parallel.

8

9

10

11

12

13

14

15

16

17

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*.

18

19

20

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++

21

22

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 purposes of this clause.

23

24

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.

1 If the *reduction-identifier* is not an *id-expression* then it is implicitly converted to one by
2 prepending the keyword operator (for example, + becomes *operator+*).

3 If the *reduction-identifier* is qualified then a qualified name lookup is used to find the
4 declaration.

5 If the *reduction-identifier* is unqualified then an *argument-dependent name lookup* must
6 be performed using the type of each list item.

▲ C++ ▲

7 If **nowait** is not used, the reduction computation will be complete at the end of the
8 construct; however, if the reduction clause is used on a construct to which **nowait** is
9 also applied, accesses to the original list item will create a race and, thus, have
10 unspecified effect unless synchronization ensures that they occur after all threads have
11 executed all of their iterations or **section** constructs, and the reduction computation
12 has completed and stored the computed value of that list item. This can most simply be
13 ensured through a barrier synchronization.

14 The location in the OpenMP program at which the values are combined and the order in
15 which the values are combined are unspecified. Therefore, when comparing sequential
16 and parallel runs, or when comparing one parallel run to another (even if the number of
17 threads used is the same), there is no guarantee that bit-identical results will be obtained
18 or that side effects (such as floating-point exceptions) will be identical or take place at
19 the same location in the OpenMP program.

20 To avoid race conditions, concurrent reads or updates of the original list item must be
21 synchronized with the update of the original list item that occurs as a result of the
22 **reduction** computation.

23 Restrictions

24 The restrictions to the **reduction** clause are as follows:

- 25 • A list item that appears in a **reduction** clause of a worksharing construct must be
26 shared in the **parallel** regions to which any of the worksharing regions arising
27 from the worksharing construct bind.
- 28 • A list item that appears in a **reduction** clause of the innermost enclosing
29 worksharing or **parallel** construct may not be accessed in an explicit task.
- 30 • Any number of **reduction** clauses can be specified on the directive, but a list item
31 can appear only once in the **reduction** clauses for that directive.
- 32 • For a *reduction-identifier* declared with the **declare reduction** construct, the
33 directive must appear before its use in a **reduction** clause.

C/C++

- 1 • The type of a list item that appears in a **reduction** clause must be valid for the
2 *reduction-identifier*. For a **max** or **min** reduction in C, the type of the list item must
3 be an allowed arithmetic data type: **char**, **int**, **float**, **double**, or **_Bool**,
4 possibly modified with **long**, **short**, **signed**, or **unsigned**. For a **max** or **min**
5 reduction in C++, the type of the list item must be an allowed arithmetic data type:
6 **char**, **wchar_t**, **int**, **float**, **double**, or **bool**, possibly modified with **long**,
7 **short**, **signed**, or **unsigned**.
- 8 • Arrays may not appear in a **reduction** clause.
- 9 • A list item that appears in a **reduction** clause must not be **const**-qualified.
- 10 • If a list item is a reference type then it must bind to the same object for all threads of
11 the team.
- 12 • The *reduction-identifier* for any list item must be unambiguous and accessible.

C/C++

Fortran

- 13 • The type of a list item that appears in a **reduction** clause must be valid for the
14 reduction operator or intrinsic.
- 15 • A list item that appears in a **reduction** clause must be definable.
- 16 • A procedure pointer may not appear in a **reduction** clause.
- 17 • A pointer with the **INTENT (IN)** attribute may not appear in the **reduction**
18 clause.
- 19 • A pointer must be associated upon entry and exit to the region.
- 20 • A pointer must not have its association status changed within the region.
- 21 • An original list item with the **POINTER** attribute must be associated at entry to the
22 construct containing the **reduction** clause. Additionally, the list item must not be
23 deallocated, allocated, or pointer assigned within the region.
- 24 • An original list item with the **ALLOCATABLE** attribute must be in the allocated state
25 at entry to the construct containing the **reduction** clause. Additionally, the list
26 item must not be deallocated and/or allocated within the region.
- 27 • If the *reduction-identifier* is defined in a **declare reduction** directive, the
28 **declare reduction** directive must be in the same subprogram, or accessible by
29 host or use association.
- 30 • If the *reduction-identifier* is a user-defined operator, the same explicit interface for
31 that operator must be accessible as at the **declare reduction** directive.

- If the *reduction-identifier* is defined in a **declare reduction** directive, any subroutine or function referenced in the initializer clause or combiner expression must be an intrinsic function, or must have an explicit interface where the same explicit interface is accessible as at the **declare reduction** directive.

2.14.3.7 **linear** clause

Summary

The **linear** clause declares one or more list items to be private to a SIMD lane and to have a linear relationship with respect to the iteration space of a loop.

Syntax

The syntax of the **linear** clause is as follows:

```
linear ( list[:linear-step] )
```

Description

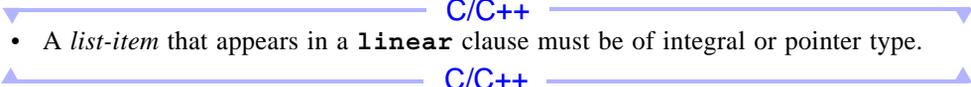
The **linear** clause provides a superset of the functionality provided by the **private** clause.

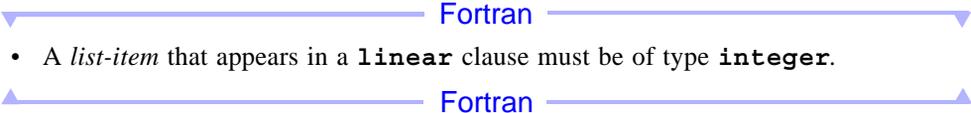
A list item that appears in a **linear** clause is subject to the **private** clause semantics described in Section 2.14.3.3 on page 159 except as noted. In addition, the value of the new list item on each iteration of the associated loop(s) corresponds to the value of the original list item before entering the construct plus the logical number of the iteration times *linear-step*. If *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.

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.

1 • A *list-item* that appears in a **linear** clause cannot appear in any other data-sharing
2 attribute clause.

3  C/C++
• A *list-item* that appears in a **linear** clause must be of integral or pointer type.

4  Fortran
• A *list-item* that appears in a **linear** clause must be of type **integer**.

5 2.14.4 Data Copying Clauses

6 This section describes the **copyin** clause (allowed on the **parallel** directive and
7 combined parallel worksharing directives) and the **copyprivate** clause (allowed on
8 the **single** directive).

9 These clauses support the copying of data values from private or threadprivate variables
10 on one implicit task or thread to the corresponding variables on other implicit tasks or
11 threads in the team.

12 The clauses accept a comma-separated list of list items (see Section 2.1 on page 26). All
13 list items appearing in a clause must be visible, according to the scoping rules of the
14 base language. Clauses may be repeated as needed, but a list item that specifies a given
15 variable may not appear in more than one clause on the same directive.

16 2.14.4.1 **copyin** clause

17 **Summary**

18 The **copyin** clause provides a mechanism to copy the value of the master thread's
19 threadprivate variable to the threadprivate variable of each other member of the team
20 executing the **parallel** region.

21 **Syntax**

22 The syntax of the **copyin** clause is as follows:

```
copyin (list)
```

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

- 1 • A list item that appears in a **copyin** clause must be threadprivate. Named variables
2 appearing in a threadprivate common block may be specified: it is not necessary to
3 specify the whole common block.
- 4 • A common block name that appears in a **copyin** clause must be declared to be a
5 common block in the same scoping unit in which the **copyin** clause appears.

Fortran

6 2.14.4.2 **copyprivate** clause

7 **Summary**

8 The **copyprivate** clause provides a mechanism to use a private variable to broadcast
9 a value from the data environment of one implicit task to the data environments of the
10 other implicit tasks belonging to the **parallel** region.

11 To avoid race conditions, concurrent reads or updates of the list item must be
12 synchronized with the update of the list item that occurs as a result of the
13 **copyprivate** clause.

14 **Syntax**

15 The syntax of the **copyprivate** clause is as follows:

```
copyprivate (list)
```

16 **Description**

17 The effect of the **copyprivate** clause on the specified list items occurs after the
18 execution of the structured block associated with the **single** construct (see
19 Section 2.7.3 on page 63), and before any of the threads in the team have left the barrier
20 at the end of the construct.

C/C++

21 In all other implicit tasks belonging to the **parallel** region, each specified list item
22 becomes defined with the value of the corresponding list item in the implicit task whose
23 thread executed the structured block. For variables of non-array type, the definition
24 occurs by copy assignment. For an array of elements of non-array type, each element is

1 copied by copy assignment from an element of the array in the data environment of the
2 implicit task associated with the thread that executed the structured block to the
3 corresponding element of the array in the data environment of the other implicit tasks.

▲ C/C++ ▲

▼ C++ ▼

4 For class types, a copy assignment operator is invoked. The order in which copy
5 assignment operators for different variables of class type are called is unspecified.

▲ C++ ▲

▼ Fortran ▼

6 If a list item does not have the **POINTER** attribute, then in all other implicit tasks
7 belonging to the **parallel** region, the list item becomes defined as if by intrinsic
8 assignment with the value of the corresponding list item in the implicit task associated
9 with the thread that executed the structured block.

10 If the list item has the **POINTER** attribute, then, in all other implicit tasks belonging to
11 the **parallel** region, the list item receives, as if by pointer assignment, the same
12 association status of the corresponding list item in the implicit task associated with the
13 thread that executed the structured block.

▲ Fortran ▲

▼ ▼

14 **Note** – The **copyprivate** clause is an alternative to using a shared variable for the
15 value when providing such a shared variable would be difficult (for example, in a
16 recursion requiring a different variable at each level).

▲ ▲

17 Restrictions

18 The restrictions to the **copyprivate** clause are as follows:

- 19 • All list items that appear in the **copyprivate** clause must be either **threadprivate**
20 or **private** in the enclosing context.
- 21 • A list item that appears in a **copyprivate** clause may not appear in a **private** or
22 **firstprivate** clause on the **single** construct.

▼ C++ ▼

- 23 • A variable of class type (or array thereof) that appears in a **copyprivate** clause
24 requires an accessible unambiguous copy assignment operator for the class type.

▲ C++ ▲

- 1 • A common block that appears in a **copyprivate** clause must be threadprivate.
 2 • Pointers with the **INTENT (IN)** attribute may not appear in the **copyprivate**
 3 clause.

4 2.14.5 map Clause

5 Summary

6 The **map** clause maps a variable from the current task's data environment to the device
 7 data environment associated with the construct.

8 Syntax

9 The syntax of the **map** clause is as follows:

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

10 Description

11 The list items that appear in a **map** clause may include array sections.

12 For list items that appear in a **map** clause, corresponding new list items are created in
 13 the device data environment associated with the construct.

14 The original and corresponding list items may share storage such that writes to either
 15 item by one task followed by a read or write of the other item by another task without
 16 intervening synchronization can result in data races.

17 If a corresponding list item of the original list item is in the enclosing device data
 18 environment, the new device data environment uses the corresponding list item from the
 19 enclosing device data environment. No additional storage is allocated in the new device
 20 data environment and neither initialization nor assignment is performed, regardless of
 21 the *map-type* that is specified.

1 If a corresponding list item is not in the enclosing device data environment, a new list
2 item with language-specific attributes is derived from the original list item and created
3 in the new device data environment. This new list item becomes the corresponding list
4 item to the original list item in the new device data environment. Initialization and
5 assignment are performed if specified by the *map-type*.

C/C++

6 If a new list item is created then a new list item of the same type, with automatic storage
7 duration, is allocated for the construct. The storage and thus lifetime of this list item
8 lasts until the block in which it is created exits. The size and alignment of the new list
9 item are determined by the type of the variable. This allocation occurs if the region
10 references the list item in any statement.

11 If the type of the variable appearing in an array section is pointer, reference to array, or
12 reference to pointer then the variable is implicitly treated as if it had appeared in a **map**
13 clause with a *map-type* of **alloc**. The corresponding variable is assigned the address of
14 the storage location of the corresponding array section in the new device data
15 environment. If the variable appears in a **to** or **from** clause in a **target update**
16 region enclosed by the new device data environment but not as part of the specification
17 of an array section, the behavior is unspecified.

C/C++

Fortran

18 If a new list item is created then a new list item of the same type, type parameter, and
19 rank is allocated.

Fortran

20 The *map-type* determines how the new list item is initialized.

21 The **alloc** *map-type* declares that on entry to the region each new corresponding list
22 item has an undefined initial value.

23 The **to** *map-type* declares that on entry to the region each new corresponding list item
24 is initialized with the original list item's value.

25 The **from** *map-type* declares that on exit from the region the corresponding list item's
26 value is assigned to each original list item.

27 The **tofrom** *map-type* declares that on entry to the region each new corresponding list
28 item is initialized with the original list item's value and that on exit from the region the
29 corresponding list item's value is assigned to each original list item.

30 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 enclosing device data environment, all of the original storage must have corresponding storage in the enclosing device data environment.
- A variable that is part of another variable (such as a field of a structure) but is not an array element or an array section cannot appear in a **map** clause.
- 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 enclosing 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 enclosing device data environment already contains that variable.

C/C++

Fortran

- The value of the new list item becomes that of the original list item in the **map** initialization and assignment.

Fortran

2.15 declare reduction Directive

Summary

The following section describes the directive for declaring user-defined reductions. The **declare reduction** directive declares a *reduction-identifier* that can be used in a **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 a list of type names
- *combiner* is an expression

- 1 • *initializer-clause* is **initializer** (*initializer-expr*) where *initializer-expr* is
2 **omp_priv** *initializer* or *function-name* (*argument-list*)

▲ **C++** ▲

3

▼ **Fortran** ▼

```
!$omp declare reduction( reduction-identifier : type-list : combiner )  
[initializer-clause]
```

4 where:

- 5 • *reduction-identifier* is either a base language identifier, or a user-defined operator, or
6 one of the following operators: +, -, *, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of
7 the following intrinsic procedure names: **max**, **min**, **iand**, **ior**, **ieor**.
- 8 • *type-list* is a list of type specifiers
- 9 • *combiner* is either an assignment statement or a subroutine name followed by an
10 argument list
- 11 • *initializer-clause* is **initializer** (*initializer-expr*), where *initializer-expr* is
12 **omp_priv** = *expression* or *subroutine-name* (*argument-list*)

▲ **Fortran** ▲

13 Description

14 Custom reductions can be defined using the **declare reduction** directive; the
15 *reduction-identifier* and the type identify the **declare reduction** directive. The
16 *reduction-identifier* can later be used in a **reduction** clause using variables of the
17 type or types specified in the **declare reduction** directive. If the directive applies
18 to several types then it is considered as if there were multiple **declare reduction**
19 directives, one for each type.

▼ **Fortran** ▼

20 If a type with deferred or assumed length type parameter is specified in a **declare**
21 **reduction** directive, the *reduction-identifier* of that directive can be used in a
22 **reduction** clause with any variable of the same type, regardless of the length type
23 parameters with which the variable is declared.

▲ **Fortran** ▲

1 The visibility and accessibility of this declaration are the same as those of a variable
2 declared at the same point in the program. The enclosing context of the *combiner* and of
3 the *initializer-expr* will be that of the **declare reduction** directive. The *combiner*
4 and the *initializer-expr* must be correct in the base language as if they were the body of
5 a function defined at the same point in the program.

C++

6 The **declare reduction** directive can also appear at points in the program at which
7 a static data member could be declared. In this case, the visibility and accessibility of
8 the declaration are the same as those of a static data member declared at the same point
9 in the program.

C++

10 The *combiner* specifies how partial results can be combined into a single value. The
11 *combiner* can use the special variable identifiers **omp_in** and **omp_out** that are of the
12 type of the variables being reduced with this *reduction-identifier*. Each of them will
13 denote one of the values to be combined before executing the *combiner*. It is assumed
14 that the special **omp_out** identifier will refer to the storage that holds the resulting
15 combined value after executing the *combiner*.

16 The number of times the *combiner* is executed, and the order of these executions, for
17 any **reduction** clause is unspecified.

Fortran

18 If the *combiner* is a subroutine name with an argument list, the *combiner* is evaluated by
19 calling the subroutine with the specified argument list.

20 If the *combiner* is an assignment statement, the *combiner* is evaluated by executing the
21 assignment statement.

Fortran

22 As the *initializer-expr* value of a user-defined reduction is not known *a priori* the
23 *initializer-clause* can be used to specify one. Then the contents of the *initializer-clause*
24 will be used as the initializer for private copies of reduction list items where the
25 **omp_priv** identifier will refer to the storage to be initialized. The special identifier
26 **omp_orig** can also appear in the *initializer-clause* and it will refer to the storage of the
27 original variable to be reduced.

28 The number of times that the *initializer-expr* is evaluated, and the order of these
29 evaluations, is unspecified.

C/C++

1 If the *initializer-expr* is a function name with an argument list, the *initializer-expr* is
2 evaluated by calling the function with the specified argument list. Otherwise, the
3 *initializer-expr* specifies how `omp_priv` is declared and initialized.

C/C++

C

4 If no *initializer-clause* is specified, the private variables will be initialized following the
5 rules for initialization of objects with static storage duration.

C

C++

6 If no *initializer-expr* is specified, the private variables will be initialized following the
7 rules for *default-initialization*.

C++

Fortran

8 If the *initializer-expr* is a subroutine name with an argument list, the *initializer-expr* is
9 evaluated by calling the subroutine with the specified argument list.

10 If the *initializer-expr* is an assignment statement, the *initializer-expr* is evaluated by
11 executing the assignment statement.

12 If no *initializer-clause* is specified, the private variables will be initialized as follows:

- 13 • For **complex**, **real**, or **integer** types, the value 0 will be used.
- 14 • For **logical** types, the value **.false.** will be used.
- 15 • For derived types for which default initialization is specified, default initialization
16 will be used.
- 17 • Otherwise, not specifying an *initializer-clause* results in unspecified behavior.

Fortran

C/C++

18 If *reduction-identifier* is used in a **target** region then a **declare target** construct
19 must be specified for any function that can be accessed through *combiner* and
20 *initializer-expr*.

C/C++

Fortran

1 If *reduction-identifier* is used in a **target** region then a **declare target** construct
2 must be specified for any function or subroutine that can be accessed through *combiner*
3 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++

Fortran

- If the *initializer-expr* is a subroutine name with an argument list, then one of the arguments must be **omp_priv**.

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- If the **declare reduction** directive appears in a module and the corresponding **reduction** clause does not appear in the same module, the *reduction-identifier* must be a user-defined operator, one of the allowed operators or one of the allowed intrinsic procedures.
 - If the *reduction-identifier* is a user-defined operator or an extended operator, the interface for that operator must be defined in the same subprogram, or must be accessible by host or use association.
 - If the **declare reduction** directive appears in a module, any user-defined operators used in the combiner must be defined in the same subprogram, or must be accessible by host or use association. The user-defined operators must also be accessible by host or use association in the subprogram in which the corresponding **reduction** clause appears.
 - Any subroutine or function used in the **initializer** clause or *combiner* expression must be an intrinsic function, or must have an explicit interface in the same subprogram or must be accessible by host or use association.
 - 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 character type and the same *reduction-identifier* are allowed in the same scope.

Fortran

21

Cross References

22

- **reduction** clause, Section 2.14.3.6 on page 167.

2.16 Nesting of Regions

This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are as follows:

- A worksharing region may not be closely nested inside a worksharing, explicit **task**, **critical**, **ordered**, **atomic**, or **master** region.
- A **barrier** region may not be closely nested inside a worksharing, explicit **task**, **critical**, **ordered**, **atomic**, or **master** region.
- A **master** region may not be closely nested inside a worksharing, **atomic**, or explicit **task** region.
- An **ordered** region may not be closely nested inside a **critical**, **atomic**, or explicit **task** region.
- An **ordered** region must be closely nested inside a loop region (or parallel loop region) with an **ordered** clause.
- A **critical** region may not be nested (closely or otherwise) inside a **critical** region with the same name. Note that this restriction is not sufficient to prevent deadlock.
- OpenMP constructs may not be nested inside an **atomic** region.
- OpenMP constructs may not be nested inside a **simd** region.
- If a **target**, **target update**, or **target data** construct appears within a **target** region then the behavior is unspecified.
- If specified, a **teams** construct must be contained within a **target** construct. That **target** construct must contain no statements 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.
- A **distribute** construct must be closely nested in a **teams** region.
- If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a **task** construct and the **cancel** construct must be nested inside a **taskgroup** region. 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.
- A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be nested inside a **task** construct. A **cancellation point** construct for which *construct-type-clause* is not **taskgroup** must be closely nested inside an 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 sections:

4

5

- Runtime library definitions (Section 3.1 on page 188).

6

- Execution environment routines that can be used to control and to query the parallel execution environment (Section 3.2 on page 189).

7

8

- Lock routines that can be used to synchronize access to data (Section 3.3 on page 224).

9

10

- Portable timer routines (Section 3.4 on page 233).

11

12

Throughout this chapter, *true* and *false* are used as generic terms to simplify the description of the routines.

13

14

▼ C/C++ ▲
true means a nonzero integer value and *false* means an integer value of zero.

▼ Fortran ▲
true means a logical value of `.TRUE.` and *false* means a logical value of `.FALSE.`

15

16

Restrictions

17

The following restriction applies to all OpenMP runtime library routines:

18

- OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.

19

▲ Fortran ▼

3.1 Runtime Library Definitions

For each base language, a compliant implementation must supply a set of definitions for the OpenMP API runtime library routines and the special data types of their parameters. The set of definitions must contain a declaration for each OpenMP API runtime library routine and a declaration for the *simple lock*, *nestable lock*, *schedule*, and *thread affinity policy* data types. In addition, each set of definitions may specify other implementation specific values.

C/C++

The library routines are external functions with “C” linkage.

Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a header file named `omp.h`. This file defines the following:

- The prototypes of all the routines in the chapter.
- The type `omp_lock_t`.
- The type `omp_nest_lock_t`.
- The type `omp_sched_t`.
- The type `omp_proc_bind_t`.

See Section C.1 on page 288 for an example of this file.

C/C++

Fortran

The OpenMP Fortran API runtime library routines are external procedures. The return values of these routines are of default kind, unless otherwise specified.

Interface declarations for the OpenMP Fortran runtime library routines described in this chapter shall be provided in the form of a Fortran `include` file named `omp_lib.h` or a Fortran 90 `module` named `omp_lib`. It is implementation defined whether the `include` file or the `module` file (or both) is provided.

These files define the following:

- The interfaces of all of the routines in this chapter.
- The `integer` parameter `omp_lock_kind`.
- The `integer` parameter `omp_nest_lock_kind`.
- The `integer` parameter `omp_sched_kind`.
- The `integer` parameter `omp_proc_bind_kind`.

1 • The **integer parameter `openmp_version`** with a value `yyyymm` where `yyyy`
2 and `mm` are the year and month designations of the version of the OpenMP Fortran
3 API that the implementation supports. This value matches that of the C preprocessor
4 macro `_OPENMP`, when a macro preprocessor is supported (see Section 2.2 on page
5 32).

6 See Section C.2 on page 290 and Section C.3 on page 293 for examples of these files.

7 It is implementation defined whether any of the OpenMP runtime library routines that
8 take an argument are extended with a generic interface so arguments of different **KIND**
9 type can be accommodated. See Appendix C.4 for an example of such an extension.

▲────────────────── Fortran ───────────────────▲



10 **3.2 Execution Environment Routines**

11 This section describes routines that affect and monitor threads, processors, and the
12 parallel environment.

13 **3.2.1 `omp_set_num_threads`**

14 **Summary**

15 The `omp_set_num_threads` routine affects the number of threads to be used for
16 subsequent parallel regions that do not specify a `num_threads` clause, by setting the
17 value of the first element of the `nthreads-var` ICV of the current task.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

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

Constraints on Arguments

The value of the argument passed to this routine must evaluate to a positive integer, or else the behavior of this routine is implementation defined.

Binding

The binding task set for an `omp_set_num_threads` region is the generating task.

Effect

The effect of this routine is to set the value of the first element of the *nthreads-var* ICV of the current task to the value specified in the argument.

See Section 2.5.1 on page 47 for the rules governing the number of threads used to execute a `parallel` region.

Cross References

- *nthreads-var* ICV, see Section 2.3 on page 34.
- `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 239.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 192.
- `parallel` construct, see Section 2.5 on page 44.
- `num_threads` clause, see Section 2.5 on page 44.

1 3.2.2 omp_get_num_threads

2 Summary

3 The `omp_get_num_threads` routine returns the number of threads in the current
4 team.

6 Format

C/C++

```
int omp_get_num_threads(void);
```

C/C++

Fortran

```
integer function omp_get_num_threads()
```

Fortran

9 Binding

10 The binding region for an `omp_get_num_threads` region is the innermost enclosing
11 `parallel` region.

12 Effect

13 The `omp_get_num_threads` routine returns the number of threads in the team
14 executing the `parallel` region to which the routine region binds. If called from the
15 sequential part of a program, this routine returns 1.

16 See Section 2.5.1 on page 47 for the rules governing the number of threads used to
17 execute a `parallel` region.

Cross References

- `parallel` construct, see Section 2.5 on page 44.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 189.
- `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 239.

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.

1
2
3
4
5
6
7

8
9
10

11
12
13
14
15
16

17

18
19
20

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 47 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 34.
- `parallel` construct, see Section 2.5 on page 44.
- `num_threads` clause, see Section 2.5 on page 44.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 189.
- `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 239.

3.2.4 `omp_get_thread_num`

Summary

The `omp_get_thread_num` routine returns the thread number, within the current team, of the calling thread.

1

Format

▼ C/C++ ▼

```
int omp_get_thread_num(void);
```

2

▲ C/C++ ▲

▼ Fortran ▼

```
integer function omp_get_thread_num()
```

3

▲ Fortran ▲

4

Binding

5

The binding thread set for an `omp_get_thread_num` region is the current team. The binding region for an `omp_get_thread_num` region is the innermost enclosing `parallel` region.

7

8

Effect

9

The `omp_get_thread_num` routine returns the thread number of the calling thread, within the team executing the `parallel` region to which the routine region binds. The thread number is an integer between 0 and one less than the value returned by `omp_get_num_threads`, inclusive. The thread number of the master thread of the team is 0. The routine returns 0 if it is called from the sequential part of a program.

10

11

12

13

14

Note – The thread number may change during the execution of an untied task. The value returned by `omp_get_thread_num` is not generally useful during the execution of such a task region.

15

16

17

Cross References

18

- `omp_get_num_threads` routine, see Section 3.2.2 on page 191.

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
4 device.

5 Format

C/C++

```
int omp_get_num_procs(void);
```

C/C++

Fortran

```
integer function omp_get_num_procs()
```

Fortran

8 Binding

9 The binding thread set for an `omp_get_num_procs` region is all threads on a device.
10 The effect of executing this routine is not related to any specific region corresponding to
11 any construct or API routine.

12 Effect

13 The `omp_get_num_procs` routine returns the number of processors that are available
14 to the device at the time the routine is called. Note that this value may change between
15 the time that it is determined by the `omp_get_num_procs` routine and the time that it
16 is read in the calling context due to system actions outside the control of the OpenMP
17 implementation.

1 3.2.6 `omp_in_parallel`

2 Summary

3 The `omp_in_parallel` routine returns *true* if the *active-levels-var* ICV is greater
4 than zero; otherwise, it returns *false*.

5 Format

▼ C/C++ ▼

```
int omp_in_parallel(void);
```

▲ C/C++ ▲

▼ Fortran ▼

```
logical function omp_in_parallel()
```

▲ Fortran ▲

8 Binding

9 The binding task set for an `omp_in_parallel` region is the generating task.

10 Effect

11 The effect of the `omp_in_parallel` routine is to return *true* if the current task is
12 enclosed by an active `parallel` region, and the `parallel` region is enclosed by the
13 outermost initial task region on the device; otherwise it returns *false*.

14 Cross References

- 15 • *active-levels-var*, see Section 2.3 on page 34.
- 16 • `omp_get_active_level` routine, see Section 3.2.20 on page 214.

1 3.2.7 `omp_set_dynamic`

2 Summary

3 The `omp_set_dynamic` routine enables or disables dynamic adjustment of the
4 number of threads available for the execution of subsequent `parallel` regions by
5 setting the value of the *dyn-var* ICV.

6 Format

▼ C/C++ ▼

```
void omp_set_dynamic(int dynamic_threads);
```

▲ C/C++ ▲

▼ Fortran ▼

```
subroutine omp_set_dynamic (dynamic_threads)  
  logical dynamic_threads
```

▲ Fortran ▲

9 Binding

10 The binding task set for an `omp_set_dynamic` region is the generating task.

11 Effect

12 For implementations that support dynamic adjustment of the number of threads, if the
13 argument to `omp_set_dynamic` evaluates to *true*, dynamic adjustment is enabled for
14 the current task; otherwise, dynamic adjustment is disabled for the current task. For
15 implementations that do not support dynamic adjustment of the number of threads this
16 routine has no effect: the value of *dyn-var* remains *false*.

17 See Section 2.5.1 on page 47 for the rules governing the number of threads used to
18 execute a `parallel` region.

Cross References:

- *dyn-var* ICV, see Section 2.3 on page 34.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 191.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 198.
- `OMP_DYNAMIC` environment variable, see Section 4.3 on page 240.

3.2.8 `omp_get_dynamic`

Summary

The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether dynamic adjustment of the number of threads is enabled or disabled.

Format

▼ C/C++ ▼

```
int omp_get_dynamic(void);
```

▲ C/C++ ▲

▼ Fortran ▼

```
logical function omp_get_dynamic()
```

▲ Fortran ▲

Binding

The binding task set for an `omp_get_dynamic` region is the generating task.

Effect

This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current task; it returns *false*, otherwise. If an implementation does not support dynamic adjustment of the number of threads, then this routine always returns *false*.

1 See Section 2.5.1 on page 47 for the rules governing the number of threads used to
2 execute a `parallel` region.

3 **Cross References**

- 4 • *dyn-var* ICV, see Section 2.3 on page 34.
- 5 • `omp_set_dynamic` routine, see Section 3.2.7 on page 197.
- 6 • `OMP_DYNAMIC` environment variable, see Section 4.3 on page 240.

7 **3.2.9 `omp_get_cancellation`**

8 **Summary**

9 The `omp_get_cancellation` routine returns the value of the *cancel-var* ICV, which
10 controls the behavior of the `cancel` construct and cancellation points.

11 **Format**

12  C/C++ 

```
13 int omp_get_cancellation(void);
```

14  C/C++ 

15  Fortran 

```
16 logical function omp_get_cancellation()
```

17  Fortran 

18 **Binding**

19 The binding task set for an `omp_get_cancellation` region is the whole program.

1
2
3
4
5
6
7
8
9
10
11
12

Effect

This routine returns *true* if cancellation is activated. It returns *false* otherwise.

Cross References:

- *cancel-var* ICV, see Section 2.3.1 on page 35.
- `OMP_CANCELLATION` environment variable, see Section 4.11 on page 246.

3.2.10 `omp_set_nested`

Summary

The `omp_set_nested` routine enables or disables nested parallelism, by setting the *nest-var* ICV.

Format

▼ C/C++ ▼

```
void omp_set_nested(int nested);
```

▲ C/C++ ▲

▼ Fortran ▼

```
subroutine omp_set_nested (nested)  
  logical nested
```

▲ Fortran ▲

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20

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 47 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 34.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 207.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 209.
- `omp_get_nested` routine, see Section 3.2.11 on page 201.
- `OMP_NESTED` environment variable, see Section 4.6 on page 243.

3.2.11 `omp_get_nested`

Summary

The `omp_get_nested` routine returns the value of the *nest-var* ICV, which determines if nested parallelism is enabled or disabled.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

Format

▼ C/C++ ▼

```
int omp_get_nested(void);
```

▲ C/C++ ▲

▼ Fortran ▼

```
logical function omp_get_nested()
```

▲ Fortran ▲

Binding

The binding task set for an `omp_get_nested` region is the generating task.

Effect

This routine returns *true* if nested parallelism is enabled for the current task; it returns *false*, otherwise. If an implementation does not support nested parallelism, this routine always returns *false*.

See Section 2.5.1 on page 47 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 34.
- `omp_set_nested` routine, see Section 3.2.10 on page 200.
- `OMP_NESTED` environment variable, see Section 4.6 on page 243.

1 3.2.12 `omp_set_schedule`

2 Summary

3 The `omp_set_schedule` routine affects the schedule that is applied when `runtime`
4 is used as schedule kind, by setting the value of the `run-sched-var` ICV.

5 Format

6

▼ C/C++ ▼

```
void omp_set_schedule(omp_sched_t kind, int modifier);
```

7

▲ C/C++ ▲

8

▼ Fortran ▼

```
subroutine omp_set_schedule(kind, modifier)
integer (kind=omp_sched_kind) kind
integer modifier
```

▲ Fortran ▲

9

10 Constraints on Arguments

11 The first argument passed to this routine can be one of the valid OpenMP schedule kinds
12 (except for `runtime`) or any implementation specific schedule. The C/C++ header file
13 (`omp.h`) and the Fortran include file (`omp_lib.h`) and/or Fortran 90 module file
14 (`omp_lib`) define the valid constants. The valid constants must include the following,
15 which can be extended with implementation specific values:

1

◀ C/C++ ▶

```

typedef enum omp_sched_t {
    omp_sched_static = 1,
    omp_sched_dynamic = 2,
    omp_sched_guided = 3,
    omp_sched_auto = 4
} omp_sched_t;

```

2

▶ C/C++ ▶

◀ Fortran ▶

```

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

```

3

▶ Fortran ▶

4

Binding

5

The binding task set for an `omp_set_schedule` region is the generating task.

6

Effect

7

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.

8

9

10

11

12

13

14

Cross References

- *run-sched-var* ICV, see Section 2.3 on page 34.
- `omp_get_schedule` routine, see Section 3.2.13 on page 205.
- `OMP_SCHEDULE` environment variable, see Section 4.1 on page 238.
- Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 59.

3.2.13 `omp_get_schedule`

Summary

The `omp_get_schedule` routine returns the schedule that is applied when the `runtime` schedule is used.

Format

C/C++

```
void omp_get_schedule(omp_sched_t * kind, int * modifier );
```

C/C++

Fortran

```
subroutine omp_get_schedule(kind, modifier)  
integer (kind=omp_sched_kind) kind  
integer modifier
```

Fortran

Binding

The binding task set for an `omp_get_schedule` region is the generating task.

Effect

This routine returns the *run-sched-var* ICV in the task to which the routine binds. The first argument **kind** returns the schedule to be used. It can be any of the standard schedule types as defined in Section 3.2.12 on page 203, or any implementation specific schedule type. The second argument is interpreted as in the **omp_set_schedule** call, defined in Section 3.2.12 on page 203.

Cross References

- *run-sched-var* ICV, see Section 2.3 on page 34.
- **omp_set_schedule** routine, see Section 3.2.12 on page 203.
- **OMP_SCHEDULE** environment variable, see Section 4.1 on page 238.
- Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 59.

3.2.14 **omp_get_thread_limit**

Summary

The **omp_get_thread_limit** routine returns the maximum number of OpenMP threads available on the device.

Format

C/C++

```
int omp_get_thread_limit(void);
```

C/C++

Fortran

```
integer function omp_get_thread_limit()
```

Fortran

1
2
3
4

5
6
7

8
9
10

11

12
13
14

15
16

17

Binding

The binding thread set for an `omp_get_thread_limit` 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_thread_limit` routine returns the maximum number of OpenMP threads available on the device as stored in the ICV *thread-limit-var*.

Cross References

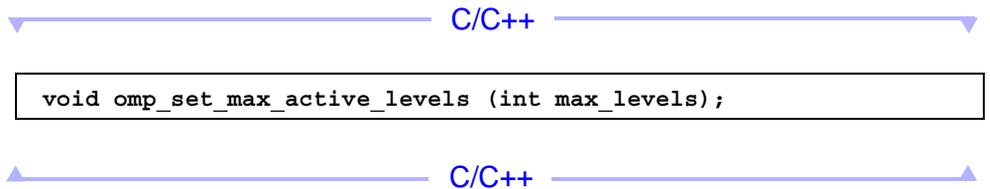
- *thread-limit-var* ICV, see Section 2.3 on page 34.
- `OMP_THREAD_LIMIT` environment variable, see Section 4.10 on page 246.

3.2.15 `omp_set_max_active_levels`

Summary

The `omp_set_max_active_levels` routine limits the number of nested active parallel regions on the device, by setting the *max-active-levels-var* ICV.

Format

A diagram showing the C/C++ signature for the `omp_set_max_active_levels` routine. It consists of a blue double-headed arrow at the top with "C/C++" in the center. Below the arrow is a rectangular box containing the signature: `void omp_set_max_active_levels (int max_levels);`. At the bottom is another blue double-headed arrow with "C/C++" in the center.

```
void omp_set_max_active_levels (int max_levels);
```

1

▼ Fortran ▼

```
subroutine omp_set_max_active_levels (max_levels)
integer max_levels
```

2

▲ Fortran ▲

3

Constraints on Arguments

4

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.

5

6

Binding

7

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.

8

9

10

11

Effect

12

The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified in the argument.

13

14

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.

15

16

17

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.

18

19

20

Cross References

21

- *max-active-levels-var* ICV, see Section 2.3 on page 34.

22

- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 209.

23

- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 245.

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-*
4 *levels-var* ICV, which determines the maximum number of nested active parallel regions
5 on the device.

6 Format

7

▼ C/C++ ▼

```
int omp_get_max_active_levels(void);
```

8

▲ C/C++ ▲

▼ Fortran ▼

```
integer function omp_get_max_active_levels()
```

9

▲ Fortran ▲

10 Binding

11 When called from a sequential part of the program, the binding thread set for an
12 `omp_get_max_active_levels` region is the encountering thread. When called
13 from within any explicit parallel region, the binding thread set (and binding region, if
14 required) for the `omp_get_max_active_levels` region is implementation defined.

15 Effect

16 The `omp_get_max_active_levels` routine returns the value of the *max-active-*
17 *levels-var* ICV, which determines the maximum number of nested active parallel regions
18 on the device.

Cross References

- *max-active-levels-var* ICV, see Section 2.3 on page 34.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 207.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 245.

3.2.17 `omp_get_level`

Summary

The `omp_get_level` routine returns the value of the *levels-var* ICV.

Format

C/C++

```
int omp_get_level(void);
```

C/C++

Fortran

```
integer function omp_get_level()
```

Fortran

Binding

The binding task set for an `omp_get_level` region is the generating task.

1
2
3
4
5

6
7
8
9

10

11
12
13

14
15

16

17

Effect

The effect of the `omp_get_level` routine is to return the number of nested `parallel` regions (whether active or inactive) 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

- *levels-var* ICV, see Section 2.3 on page 34.
- `omp_get_active_level` routine, see Section 3.2.20 on page 214.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 245.

3.2.18 `omp_get_ancestor_thread_num`

Summary

The `omp_get_ancestor_thread_num` routine returns, for a given nested level of the current thread, the thread number of the ancestor of the current thread.

Format

C/C++

```
int omp_get_ancestor_thread_num(int level);
```

C/C++

Fortran

```
integer function omp_get_ancestor_thread_num(level)  
integer level
```

Fortran

1
2
3
4

5
6
7
8
9

10
11
12

13
14
15
16

17

18
19
20

Binding

The binding thread set for an `omp_get_ancestor_thread_num` region is the encountering thread. The binding region for an `omp_get_ancestor_thread_num` region is the innermost enclosing `parallel` region.

Effect

The `omp_get_ancestor_thread_num` routine returns the thread number of the ancestor at a given nest level of the current thread or the thread number of the current thread. If the requested nest level is outside the range of 0 and the nest level of the current thread, as returned by the `omp_get_level` routine, the routine returns -1.

Note – When the `omp_get_ancestor_thread_num` routine is called with a value of `level=0`, the routine always returns 0. If `level=omp_get_level()`, the routine has the same effect as the `omp_get_thread_num` routine.

Cross References

- `omp_get_level` routine, see Section 3.2.17 on page 210.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 193.
- `omp_get_team_size` routine, see Section 3.2.19 on page 212.

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.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17

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 191.
- `omp_get_level` routine, see Section 3.2.17 on page 210.
- `omp_get_ancestor_thread_num` routine, see Section 3.2.18 on page 211.

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

Binding

The binding task set for the an `omp_get_active_level` region is the generating task.

1
2
3
4

5
6
7

8

9
10
11

12
13

14

15

16
17

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

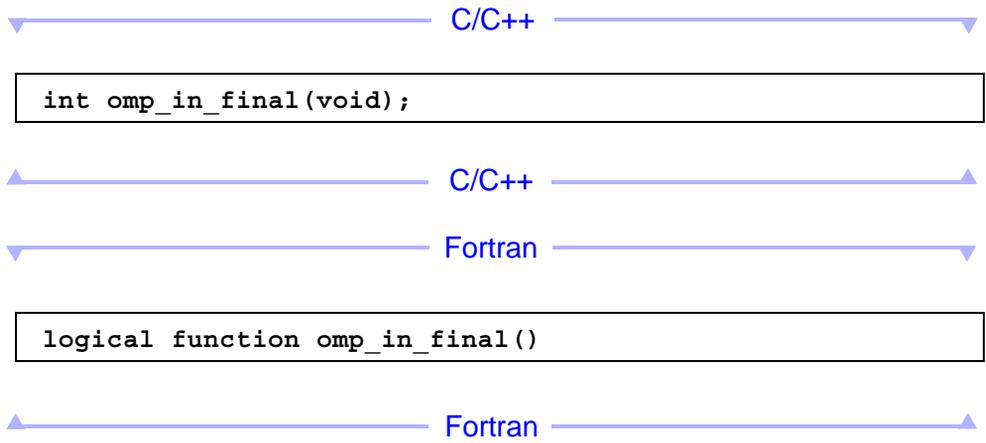
- `active-levels-var` ICV, see Section 2.3 on page 34.
- `omp_get_level` routine, see Section 3.2.17 on page 210.

3.2.21 `omp_in_final`

Summary

The `omp_in_final` routine returns *true* if the routine is executed in a final task region; otherwise, it returns *false*.

Format



Binding

The binding task set for an `omp_in_final` region is the generating task.

1
2
3

Effect

`omp_in_final` returns *true* if the enclosing task region is final. Otherwise, it returns *false*.

3.2.22 `omp_get_proc_bind`

5
6
7

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.

8
9

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 ▲

10
11

12

13

Constraints on Arguments

14
15
16
17

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:

1

C/C++

2

```
typedef enum omp_proc_bind_t {  
3     omp_proc_bind_false = 0,  
4     omp_proc_bind_true = 1,  
5     omp_proc_bind_master = 2,  
6     omp_proc_bind_close = 3,  
7     omp_proc_bind_spread = 4  
8 } omp_proc_bind_t;
```

C/C++

9

Fortran

10

```
integer (kind=omp_proc_bind_kind), &  
11     parameter :: omp_proc_bind_false = 0  
12 integer (kind=omp_proc_bind_kind), &  
13     parameter :: omp_proc_bind_true = 1  
14 integer (kind=omp_proc_bind_kind), &  
15     parameter :: omp_proc_bind_master = 2  
16 integer (kind=omp_proc_bind_kind), &  
17     parameter :: omp_proc_bind_close = 3  
18 integer (kind=omp_proc_bind_kind), &  
19     parameter :: omp_proc_bind_spread = 4
```

Fortran

20

Binding

21

The binding task set for an `omp_get_proc_bind` region is the generating task.

22

Effect

23

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 49 for the rules governing the thread affinity policy.

24

25

26

Cross References

27

- *bind-var* ICV, see Section 2.3 on page 34.

28

- `OMP_PROC_BIND` environment variable, see Section 4.4 on page 241.

29

- Controlling OpenMP thread affinity, see Section 2.5.2 on page 49.

1 3.2.23 `omp_set_default_device`

2 Summary

3 The `omp_set_default_device` routine controls the default target device by
4 assigning the value of the *default-device-var* ICV.

5 Format

6

▼ C/C++ ▼

```
void omp_set_default_device(int device_num );
```

7

▲ C/C++ ▲

▼ Fortran ▼

```
subroutine omp_set_default_device( device_num )  
integer device_num
```

▲ Fortran ▲

8

Binding

9 The binding task set for an `omp_set_default_device` region is the generating
10 task.

11 Effect

12 The effect of this routine is to set the value of the *default-device-var* ICV of the current
13 task to the value specified in the argument. When called from within a **target** region
14 the effect of this routine is unspecified.

15 Cross References:

- 16 • *default-device-var*, see Section 2.3 on page 34.
- 17 • `omp_get_default_device`, see Section 3.2.24 on page 219.
- 18 • `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 248

1 3.2.24 omp_get_default_device

2 Summary

3 The `omp_get_default_device` routine returns the default target device.

4 Format

▶ C/C++ ▶

```
int omp_get_default_device(void);
```

▲ C/C++ ▲

▶ Fortran ▶

```
integer function omp_get_default_device()
```

▲ Fortran ▲

7 Binding

8 The binding task set for an `omp_get_default_device` region is the generating
9 task.

10 Effect

11 The `omp_get_default_device` routine returns the value of the *default-device-var*
12 ICV of the current task. When called from within a **target** region the effect of this
13 routine is unspecified.

14 Cross References

- 15 • *default-device-var*, see Section 2.3 on page 34.
- 16 • `omp_set_default_device`, see Section 3.2.23 on page 218.
- 17 • `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 248.

1 3.2.25 `omp_get_num_devices`

2 Summary

3 The `omp_get_num_devices` routine returns the number of target devices.

4 Format

5  C/C++ 

```
int omp_get_num_devices(void);
```

6  C/C++ 

7  Fortran 

```
integer function omp_get_num_devices()
```

8  Fortran 

9 Binding

10 The binding task set for an `omp_get_num_devices` region is the generating task.

11 Effect

12 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.

13 Cross References:

14 None.

1 3.2.26 `omp_get_num_teams`

2 Summary

3 The `omp_get_num_teams` routine returns the number of teams in the current **teams**
4 region.

5 Format

6 C/C++

```
int omp_get_num_teams(void);
```

7 C/C++

8 Fortran

```
integer function omp_get_num_teams()
```

9 Fortran

9 Binding

10 The binding task set for an `omp_get_num_teams` region is the generating task.

11 Effect

12 The effect of this routine is to return the number of teams in the current **teams** region.
13 The routine returns 1 if it is called from outside of a **teams** region.

14 Cross References:

- 15 • **teams** construct, see Section 2.9.5 on page 86.

1 3.2.27 `omp_get_team_num`

2 Summary

3 The `omp_get_team_num` routine returns the team number of the calling thread.

4 Format

5  C/C++ 

```
int omp_get_team_num(void);
```

6  C/C++ 

7  Fortran 

```
integer function omp_get_team_num()
```

8  Fortran 

9 Binding

10 The binding task set for an `omp_get_team_num` region is the generating task.

11 Effect

12 The `omp_get_team_num` routine returns the team number of the calling thread. The
13 team number is an integer between 0 and one less than the value returned by
14 `omp_get_num_teams`, inclusive. The routine returns 0 if it is called outside of a
`teams` region.

15 Cross References:

- 16 • `teams` construct, see Section 2.9.5 on page 86.
- 17 • `omp_get_num_teams` routine, see Section 3.2.26 on page 221.

1 3.2.28 `omp_is_initial_device`

2 **Summary**

3 The `omp_is_initial_device` routine returns *true* if the current task is executing
4 on the host device; otherwise, it returns *false*.

5 **Format**

6

▼ `C/C++` ▼

```
int omp_is_initial_device(void);
```

7

▲ `C/C++` ▲

8

▼ `Fortran` ▼

```
logical function omp_is_initial_device()
```

9

▲ `Fortran` ▲

10 **Binding**

11 The binding task set for an `omp_is_initial_device` region is the generating task.

12 **Effect**

13 The effect of this routine is to return *true* if the current task is executing on the host
14 device; otherwise, it returns *false*.

15 **Cross References:**

16 • `target` construct, see Section 2.9.2 on page 79.

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.

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_destroy_lock` routine uninitialized 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`.

Fortran

The nestable lock routines are as follows:

- The `omp_init_nest_lock` routine initializes a nestable lock.
- The `omp_destroy_nest_lock` routine uninitialized a nestable lock.
- The `omp_set_nest_lock` routine waits until a nestable lock is available, and then sets it.

- 1
- 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.
- 2
3

4 Restrictions

5 OpenMP lock routines have the following restrictions:

- The use of the same OpenMP lock in different contention groups results in unspecified behavior.
- 6
7

8 3.3.1 `omp_init_lock` and `omp_init_nest_lock`

9 Summary

10 These routines provide the only means of initializing an OpenMP lock.

11 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 ▲

13

1
2
3

4
5
6

7
8

9
10

11

12

13

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_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 in the uninitialized state through either routine is
3 non-conforming. A simple lock accessed by `omp_set_lock` that is in the locked state
4 must not be owned by the task that contains the call or deadlock will result.

5 **Effect**

6 Each of these routines causes suspension of the task executing the routine until the
7 specified lock is available and then sets the lock.

8 A simple lock is available if it is unlocked. Ownership of the lock is granted to the task
9 executing the routine.

10 A nestable lock is available if it is unlocked or if it is already owned by the task
11 executing the routine. The task executing the routine is granted, or retains, ownership of
12 the lock, and the nesting count for the lock is incremented.

13

14 **3.3.4 omp_unset_lock and omp_unset_nest_lock**

15 **Summary**

16 These routines provide the means of unsetting an OpenMP lock.

1

Format



```
void omp_unset_lock(omp_lock_t *lock);
void omp_unset_nest_lock(omp_nest_lock_t *lock);
```

2



```
subroutine omp_unset_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_unset_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

3



4

Constraints on Arguments

5

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.

6

7

Effect

8

For a simple lock, the `omp_unset_lock` routine causes the lock to become unlocked.

9

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.

10

11

For either routine, if the lock becomes unlocked, and if one or more task regions were suspended because the lock was unavailable, the effect is that one task is chosen and given ownership of the lock.

12

13

1 3.3.5 omp_test_lock and omp_test_nest_lock

2 Summary

3 These routines attempt to set an OpenMP lock but do not suspend execution of the task
4 executing the routine.

5 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

8 Constraints on Arguments

9 A program that accesses a lock that is in the uninitialized state through either routine is
10 non-conforming. The behavior is unspecified if a simple lock accessed by
11 **omp_test_lock** is in the locked state and is owned by the task that contains the call.

12 Effect

13 These routines attempt to set a lock in the same manner as **omp_set_lock** and
14 **omp_set_nest_lock**, except that they do not suspend execution of the task
15 executing the routine.

16 For a simple lock, the **omp_test_lock** routine returns *true* if the lock is successfully
17 set; otherwise, it returns *false*.

1 For a nestable lock, the `omp_test_nest_lock` routine returns the new nesting count
2 if the lock is successfully set; otherwise, it returns zero.



1 3.4 Timing Routines

2 This section describes routines that support a portable wall clock timer.

3 3.4.1 `omp_get_wtime`

4 Summary

5 The `omp_get_wtime` routine returns elapsed wall clock time in seconds.

6 Format

▼ C/C++ ▼

```
double omp_get_wtime(void);
```

▲ C/C++ ▲

▼ Fortran ▼

```
double precision function omp_get_wtime()
```

▲ Fortran ▲

9 Binding

10 The binding thread set for an `omp_get_wtime` region is the encountering thread. The
11 routine's return value is not guaranteed to be consistent across any set of threads.

12 Effect

13 The `omp_get_wtime` routine returns a value equal to the elapsed wall clock time in
14 seconds since some "time in the past". The actual "time in the past" is arbitrary, but it is
15 guaranteed not to change during the execution of the application program. The time
16 returned is a "per-thread time", so it is not required to be globally consistent across all
17 the threads participating in an application.

1

Format

▼ C/C++ ▼

```
double omp_get_wtick(void);
```

2

▲ C/C++ ▲

▼ Fortran ▼

```
double precision function omp_get_wtick()
```

3

▲ Fortran ▲

4

Binding

5

The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's return value is not guaranteed to be consistent across any set of threads.

6

7

Effect

8

The `omp_get_wtick` routine returns a value equal to the number of seconds between successive clock ticks of the timer used by `omp_get_wtime`.

9

Environment Variables

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30

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 34). 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 the OpenMP program.

- 1 • **OMP_CANCELLATION** sets the *cancel-var* ICV that enables or disables cancellation.
- 2 • **OMP_DISPLAY_ENV** instructs the runtime to display the OpenMP version number
- 3 and the initial values of the ICVs, once, during initialization of the runtime.
- 4 • **OMP_DEFAULT_DEVICE** sets the *default-device-var* ICV that controls the default
- 5 device number.

6 The examples in this chapter only demonstrate how these variables might be set in Unix
7 C shell (csh) environments. In Korn shell (ksh) and DOS environments the actions are
8 similar, as follows:

- 9 • csh:

```
setenv OMP_SCHEDULE "dynamic"
```

- 10 • ksh:

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

- 11 • DOS:

```
set OMP_SCHEDULE=dynamic
```

12 4.1 OMP_SCHEDULE

13 The **OMP_SCHEDULE** environment variable controls the schedule type and chunk size
14 of all loop directives that have the schedule type **runtime**, by setting the value of the
15 *run-sched-var* ICV.

16 The value of this environment variable takes the form:

17 *type[,chunk]*

18 where

- 19 • *type* is one of **static**, **dynamic**, **guided**, or **auto**
- 20 • *chunk* is an optional positive integer that specifies the chunk size

21 If *chunk* is present, there may be white space on either side of the “,”. See Section 2.7.1
22 on page 53 for a detailed description of the schedule types.

23 The behavior of the program is implementation defined if the value of **OMP_SCHEDULE**
24 does not conform to the above format.

1 Implementation specific schedules cannot be specified in `OMP_SCHEDULE`. They can
2 only be specified by calling `omp_set_schedule`, described in Section 3.2.12 on page
3 203.

4 Example:

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

5 Cross References

- 6 • `run-sched-var` ICV, see Section 2.3 on page 34.
- 7 • Loop construct, see Section 2.7.1 on page 53.
- 8 • Parallel loop construct, see Section 2.10.1 on page 95.
- 9 • `omp_set_schedule` routine, see Section 3.2.12 on page 203.
- 10 • `omp_get_schedule` routine, see Section 3.2.13 on page 205.



11 4.2 OMP_NUM_THREADS

12 The `OMP_NUM_THREADS` environment variable sets the number of threads to use for
13 `parallel` regions by setting the initial value of the `nthreads-var` ICV. See Section 2.3
14 on page 34 for a comprehensive set of rules about the interaction between the
15 `OMP_NUM_THREADS` environment variable, the `num_threads` clause, the
16 `omp_set_num_threads` library routine and dynamic adjustment of threads, and
17 Section 2.5.1 on page 47 for a complete algorithm that describes how the number of
18 threads for a `parallel` region is determined.

19 The value of this environment variable must be a list of positive integer values. The
20 values of the list set the number of threads to use for `parallel` regions at the
21 corresponding nested levels.

22 The behavior of the program is implementation defined if any value of the list specified
23 in the `OMP_NUM_THREADS` environment variable leads to a number of threads which is
24 greater than an implementation can support, or if any value is not a positive integer.

25 Example:

```
setenv OMP_NUM_THREADS 4,3,2
```

Cross References:

- *nthreads-var* ICV, see Section 2.3 on page 34.
- `num_threads` clause, Section 2.5 on page 44.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 189.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 191.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 192.
- `omp_get_team_size` routine, see Section 3.2.19 on page 212.

4.3 OMP_DYNAMIC

The `OMP_DYNAMIC` environment variable controls dynamic adjustment of the number of threads to use for executing `parallel` regions by setting the initial value of the *dyn-var* ICV. The value of this environment variable must be `true` or `false`. If the environment variable is set to `true`, the OpenMP implementation may adjust the number of threads to use for executing `parallel` regions in order to optimize the use of system resources. If the environment variable is set to `false`, the dynamic adjustment of the number of threads is disabled. The behavior of the program is implementation defined if the value of `OMP_DYNAMIC` is neither `true` nor `false`.

Example:

```
setenv OMP_DYNAMIC true
```

Cross References:

- *dyn-var* ICV, see Section 2.3 on page 34.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 197.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 198.

1 4.4 OMP_PROC_BIND

2 The `OMP_PROC_BIND` environment variable sets the initial value of the *bind-var* ICV.
3 The value of this environment variable is either `true`, `false`, or a comma separated
4 list of `master`, `close`, or `spread`. The values of the list set the thread affinity policy
5 to be used for parallel regions at the corresponding nested level.

6 If the environment variable is set to `false`, the execution environment may move
7 OpenMP threads between OpenMP places, thread affinity is disabled, and `proc_bind`
8 clauses on `parallel` constructs are ignored.

9 Otherwise, the execution environment should not move OpenMP threads between
10 OpenMP places, thread affinity is enabled, and the initial thread is bound to the first
11 place in the OpenMP place list.

12 The behavior of the program is implementation defined if any of the values in the
13 `OMP_PROC_BIND` environment variable is not `true`, `false`, or a comma separated
14 list of `master`, `close`, or `spread`. The behavior is also implementation defined if an
15 initial thread cannot be bound to the first place in the OpenMP place list.

16 Example:

```
setenv OMP_PROC_BIND false
setenv OMP_PROC_BIND "spread, spread, close"
```

17 Cross References:

- 18 • *bind-var* ICV, see Section 2.3 on page 34.
- 19 • `proc_bind` clause, see Section 2.5.2 on page 49.
- 20 • `omp_get_proc_bind` routine, see Section 3.2.22 on page 216.

21 4.5 OMP_PLACES

22 A list of places can be specified in the `OMP_PLACES` environment variable. The *place-*
23 *partition-var* ICV obtains its initial value from the `OMP_PLACES` value, and makes the
24 list available to the execution environment. The value of `OMP_PLACES` can be one of
25 two types of values: either an abstract name describing a set of places or an explicit list
26 of places described by non-negative numbers.

1 The **OMP_PLACES** environment variable can be defined using an explicit ordered list of
2 comma-separated places. A place is defined by an unordered set of comma-separated
3 non-negative numbers enclosed by braces. The meaning of the numbers and how the
4 numbering is done are implementation defined. Generally, the numbers represent the
5 smallest unit of execution exposed by the execution environment, typically a hardware
6 thread.

7 Intervals may also be used to define places. Intervals can be specified using the *<lower-*
8 *bound> : <length> : <stride>* notation to represent the following list of numbers:
9 “*<lower-bound>, <lower-bound> + <stride>, ..., <lower-bound> + (<length>-*
10 *1)*<stride>.*” When *<stride>* is omitted, a unit stride is assumed. Intervals can specify
11 numbers within a place as well as sequences of places.

12 An exclusion operator “!” can also be used to exclude the number or place immediately
13 following the operator.

14 Alternatively, the abstract names listed in TABLE 4-1 should be understood by the
15 execution and runtime environment. The precise definitions of the abstract names are
16 implementation defined. An implementation may also add abstract names as appropriate
17 for the target platform.

18 The abstract name may be appended by a positive number in parentheses to denote the
19 length of the place list to be created, that is *abstract_name(num-places)*. When
20 requesting fewer places than available on the system, the determination of which
21 resources of type *abstract_name* are to be included in the place list is implementation
22 defined. When requesting more resources than available, the length of the place list is
23 implementation defined.

TABLE 4-1 List of defined abstract names for **OMP_PLACES**

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.

24 The behavior of the program is implementation defined when the execution environment
25 cannot map a numerical value (either explicitly defined or implicitly derived from an
26 interval) within the **OMP_PLACES** list to a processor on the target platform, or if it maps
27 to an unavailable processor. The behavior is also implementation defined when the
28 **OMP_PLACES** environment variable is defined using an abstract name.

1 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"
```

2 where each of the last three definitions corresponds to the same 4 places including the
3 smallest units of execution exposed by the execution environment numbered, in turn, 0
4 to 3, 4 to 7, 8 to 11, and 12 to 15.

5 Cross References

- 6 • *place-partition-var*, Section 2.3 on page 34.
- 7 • Controlling OpenMP thread affinity, Section 2.5.2 on page 49.

8 4.6 OMP_NESTED

9 The **OMP_NESTED** environment variable controls nested parallelism by setting the
10 initial value of the *nest-var* ICV. The value of this environment variable must be **true**
11 or **false**. If the environment variable is set to **true**, nested parallelism is enabled; if
12 set to **false**, nested parallelism is disabled. The behavior of the program is
13 implementation defined if the value of **OMP_NESTED** is neither **true** nor **false**.

14 Example:

```
setenv OMP_NESTED false
```

15 Cross References

- 16 • *nest-var* ICV, see Section 2.3 on page 34.
- 17 • **omp_set_nested** routine, see Section 3.2.10 on page 200.
- 18 • **omp_get_nested** routine, see Section 3.2.19 on page 212.

19

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` | `sizeB` | `sizeK` | `sizeM` | `sizeG`

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 34.

1 4.8 OMP_WAIT_POLICY

2 The `OMP_WAIT_POLICY` environment variable provides a hint to an OpenMP
3 implementation about the desired behavior of waiting threads by setting the *wait-policy-*
4 *var* ICV. A compliant OpenMP implementation may or may not abide by the setting of
5 the environment variable.

6 The value of this environment variable takes the form:

7 **ACTIVE** | **PASSIVE**

8 The **ACTIVE** value specifies that waiting threads should mostly be active, consuming
9 processor cycles, while waiting. An OpenMP implementation may, for example, make
10 waiting threads spin.

11 The **PASSIVE** value specifies that waiting threads should mostly be passive, not
12 consuming processor cycles, while waiting. For example, an OpenMP implementation
13 may make waiting threads yield the processor to other threads or go to sleep.

14 The details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined.

15 Examples:

```
setenv OMP_WAIT_POLICY ACTIVE
setenv OMP_WAIT_POLICY active
setenv OMP_WAIT_POLICY PASSIVE
setenv OMP_WAIT_POLICY passive
```

16 Cross References

- 17 • *wait-policy-var* ICV, see Section 2.3 on page 24.

18 4.9 OMP_MAX_ACTIVE_LEVELS

19 The `OMP_MAX_ACTIVE_LEVELS` environment variable controls the maximum number
20 of nested active `parallel` regions by setting the initial value of the *max-active-levels-*
21 *var* ICV.

1 The value of this environment variable must be a non-negative integer. The behavior of
2 the program is implementation defined if the requested value of
3 **OMP_MAX_ACTIVE_LEVELS** is greater than the maximum number of nested active
4 parallel levels an implementation can support, or if the value is not a non-negative
5 integer.

6 **Cross References**

- 7 • *max-active-levels-var* ICV, see Section 2.3 on page 34.
- 8 • **omp_set_max_active_levels** routine, see Section 3.2.15 on page 207.
- 9 • **omp_get_max_active_levels** routine, see Section 3.2.16 on page 209.

10 **4.10 OMP_THREAD_LIMIT**

11 The **OMP_THREAD_LIMIT** environment variable sets the number of OpenMP threads
12 to use for the whole OpenMP program by setting the *thread-limit-var* ICV.

13 The value of this environment variable must be a positive integer. The behavior of the
14 program is implementation defined if the requested value of **OMP_THREAD_LIMIT** is
15 greater than the number of threads an implementation can support, or if the value is not
16 a positive integer.

17 **Cross References**

- 18 • *thread-limit-var* ICV, see Section 2.3 on page 34.
- 19 • **omp_get_thread_limit** routine, see Section 3.2.14 on page 206.

20 **4.11 OMP_CANCELLATION**

21 The **OMP_CANCELLATION** environment variable sets the initial value of the *cancel-var*
22 ICV.

23 The value of this environment variable must be **true** or **false**. If set to **true**, the
24 effects of the **cancel** construct and of cancellation points are enabled and cancellation
25 is activated. If set to **false**, cancellation is disabled and the **cancel** construct and
26 cancellation points are effectively ignored.

Cross References:

- *cancel-var*, see Section 2.3.1 on page 35.
- **cancel** construct, see Section 2.13.1 on page 140.
- **cancellation point** construct, see Section 2.13.2 on page 143
- **omp_get_cancellation** routine, see Section 3.2.9 on page 199

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**, undefined, or any other value than **TRUE** or **VERBOSE**.

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
```

1 The above example causes an OpenMP implementation to generate output of the
2 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
```

3 4.13 OMP_DEFAULT_DEVICE

4 The `OMP_DEFAULT_DEVICE` environment variable sets the device number to use in
5 device constructs by setting the initial value of the *default-device-var* ICV.

6 The value of this environment variable must be a non-negative integer value.

7 Cross References:

- 8 • *default-device-var* ICV, see Section 2.3 on page 34.
- 9 • device constructs, Section 2.9 on page 77

2
3

Stubs for Runtime Library Routines

4
5
6
7
8
9
10
11
12
13
14
15

This section provides stubs for the runtime library routines defined in the OpenMP API. The stubs are provided to enable portability to platforms that do not support the OpenMP API. On these platforms, OpenMP programs must be linked with a library containing these stub routines. The stub routines assume that the directives in the OpenMP program are ignored. As such, they emulate serial semantics.

Note that the lock variable that appears in the lock routines must be accessed exclusively through these routines. It should not be initialized or otherwise modified in the user program.

In an actual implementation the lock variable might be used to hold the address of an allocated memory block, but here it is used to hold an integer value. Users should not make assumptions about mechanisms used by OpenMP implementations to implement locks based on the scheme used by the stub procedures.

Fortran

16
17
18
19

Note – In order to be able to compile the Fortran stubs file, the include file `omp_lib.h` was split into two files: `omp_lib_kinds.h` and `omp_lib.h` and the `omp_lib_kinds.h` file included 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;
47     }
48
49
50
```

```

1      void omp_set_nested(int nested)
2      {
3      }
4
5      int omp_get_nested(void)
6      {
7          return 0;
8      }
9
10     void omp_set_schedule(omp_sched_t kind, int modifier)
11     {
12     }
13
14     void omp_get_schedule(omp_sched_t *kind, int *modifier)
15     {
16         *kind = omp_sched_static;
17         *modifier = 0;
18     }
19
20     int omp_get_thread_limit(void)
21     {
22         return 1;
23     }
24
25     void omp_set_max_active_levels(int max_active_levels)
26     {
27     }
28
29     int omp_get_max_active_levels(void)
30     {
31         return 0;
32     }
33
34     int omp_get_level(void)
35     {
36         return 0;
37     }
38
39     int omp_get_ancestor_thread_num(int level)
40     {
41         if (level == 0)
42         {
43             return 0;
44         }
45         else
46         {
47             return -1;
48         }
49     }
50

```

```

1      int omp_get_team_size(int level)
2      {
3          if (level == 0)
4              {
5                  return 1;
6              }
7          else
8              {
9                  return -1;
10             }
11     }
12
13     int omp_get_active_level(void)
14     {
15         return 0;
16     }
17
18     int omp_in_final(void)
19     {
20         return 1;
21     }
22
23     omp_proc_bind_t omp_get_proc_bind(void)
24     {
25         return omp_proc_bind_false;
26     }
27
28     void omp_set_default_device(int device_num)
29     {
30     }
31
32     int omp_get_default_device(void)
33     {
34         return 0;
35     }
36
37     int omp_get_num_devices(void)
38     {
39         return 0;
40     }
41
42     int omp_get_num_teams(void)
43     {
44         return 1;
45     }
46
47     int omp_get_team_num(void)
48     {
49         return 0;
50     }
51
52
53
54

```

```

1      int omp_is_initial_device(void)
2      {
3          return 1;
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_destroy_lock(omp_lock_t *arg)
20     {
21         struct __omp_lock *lock = (struct __omp_lock *)arg;
22         lock->lock = INIT;
23     }
24
25     void omp_set_lock(omp_lock_t *arg)
26     {
27         struct __omp_lock *lock = (struct __omp_lock *)arg;
28         if (lock->lock == UNLOCKED)
29         {
30             lock->lock = LOCKED;
31         }
32         else if (lock->lock == LOCKED)
33         {
34             fprintf(stderr,
35                 "error: deadlock in using lock variable\n");
36             exit(1);
37         }
38         else
39         {
40             fprintf(stderr, "error: lock not initialized\n");
41             exit(1);
42         }
43     }
44
45
46     void omp_unset_lock(omp_lock_t *arg)
47     {
48         struct __omp_lock *lock = (struct __omp_lock *)arg;
49         if (lock->lock == LOCKED)
50         {
51             lock->lock = UNLOCKED;
52         }
53         else if (lock->lock == UNLOCKED)
54         {

```

```

1         fprintf(stderr, "error: lock not set\n");
2         exit(1);
3     }
4     else
5     {
6         fprintf(stderr, "error: lock not initialized\n");
7         exit(1);
8     }
9 }
10
11 int omp_test_lock(omp_lock_t *arg)
12 {
13     struct __omp_lock *lock = (struct __omp_lock *)arg;
14     if (lock->lock == UNLOCKED)
15     {
16         lock->lock = LOCKED;
17         return 1;
18     }
19     else if (lock->lock == LOCKED)
20     {
21         return 0;
22     }
23     else
24     {
25         fprintf(stderr, "error: lock not initialized\n");
26         exit(1);
27     }
28 }
29
30 struct __omp_nest_lock
31 {
32     short owner;
33     short count;
34 };
35
36 enum { NOOWNER = -1, MASTER = 0 };
37
38 void omp_init_nest_lock(omp_nest_lock_t *arg)
39 {
40     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
41     nlock->owner = NOOWNER;
42     nlock->count = 0;
43 }
44
45
46 void omp_destroy_nest_lock(omp_nest_lock_t *arg)
47 {
48     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
49     nlock->owner = NOOWNER;
50     nlock->count = UNLOCKED;
51 }
52

```

```

1 void omp_set_nest_lock(omp_nest_lock_t *arg)
2 {
3     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
4     if (nlock->owner == MASTER && nlock->count >= 1)
5     {
6         nlock->count++;
7     }
8     else if (nlock->owner == NOOWNER && nlock->count == 0)
9     {
10        nlock->owner = MASTER;
11        nlock->count = 1;
12    }
13    else
14    {
15        fprintf(stderr,
16            "error: lock corrupted or not initialized\n");
17        exit(1);
18    }
19 }
20
21 void omp_unset_nest_lock(omp_nest_lock_t *arg)
22 {
23     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
24     if (nlock->owner == MASTER && nlock->count >= 1)
25     {
26         nlock->count--;
27         if (nlock->count == 0)
28         {
29             nlock->owner = NOOWNER;
30         }
31     }
32     else if (nlock->owner == NOOWNER && nlock->count == 0)
33     {
34         fprintf(stderr, "error: lock not set\n");
35         exit(1);
36     }
37     else
38     {
39         fprintf(stderr,
40            "error: lock corrupted or not initialized\n");
41         exit(1);
42     }
43 }
44
45 int omp_test_nest_lock(omp_nest_lock_t *arg)
46 {
47     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
48     omp_set_nest_lock(arg);
49     return nlock->count;
50 }
51
52

```

```
1      double omp_get_wtime(void)
2      {
3      /* This function does not provide a working
4      * wallclock timer. Replace it with a version
5      * customized for the target machine.
6      */
7      return 0.0;
8      }
9
10     double omp_get_wtick(void)
11     {
12     /* This function does not provide a working
13     * clock tick function. Replace it with
14     * a version customized for the target machine.
15     */
16     return 365. * 86400.;
17     }
18
```

1 A.2 Fortran Stub Routines

```
2
3      subroutine omp_set_num_threads(num_threads)
4          integer num_threads
5          return
6      end subroutine
7
8      integer function omp_get_num_threads()
9          omp_get_num_threads = 1
10         return
11     end function
12
13     integer function omp_get_max_threads()
14         omp_get_max_threads = 1
15         return
16     end function
17
18     integer function omp_get_thread_num()
19         omp_get_thread_num = 0
20         return
21     end function
22
23     integer function omp_get_num_procs()
24         omp_get_num_procs = 1
25         return
26     end function
27
28     logical function omp_in_parallel()
29         omp_in_parallel = .false.
30         return
31     end function
32
33     subroutine omp_set_dynamic(dynamic_threads)
34         logical dynamic_threads
35         return
36     end subroutine
37
38     logical function omp_get_dynamic()
39         omp_get_dynamic = .false.
40         return
41     end function
42
43     logical function omp_get_cancellation()
44         omp_get_cancellation = .false.
45         return
46     end function
47
48
49
50
```

```

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, modifier)
12         include 'omp_lib_kinds.h'
13         integer (kind=omp_sched_kind) kind
14         integer modifier
15         return
16     end subroutine
17
18     subroutine omp_get_schedule(kind, modifier)
19         include 'omp_lib_kinds.h'
20         integer (kind=omp_sched_kind) kind
21         integer modifier
22         kind = omp_sched_static
23         modifier = 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
48         if ( level .eq. 0 ) then
49             omp_get_ancestor_thread_num = 0
50         else
51             omp_get_ancestor_thread_num = -1
52         end if
53         return
54     end function

```

```

1
2      integer function omp_get_team_size( level )
3          integer level
4          if ( level .eq. 0 ) then
5              omp_get_team_size = 1
6          else
7              omp_get_team_size = -1
8          end if
9          return
10     end function
11
12     integer function omp_get_active_level()
13         omp_get_active_level = 0
14         return
15     end function
16
17     logical function omp_in_final()
18         omp_in_final = .true.
19         return
20     end function
21
22     function omp_get_proc_bind()
23         include 'omp_lib_kinds.h'
24         integer (kind=omp_proc_bind_kind) omp_get_proc_bind
25         omp_get_proc_bind = omp_proc_bind_false
26     end function omp_get_proc_bind
27
28     subroutine omp_set_default_device(device_num)
29         integer device_num
30         return
31     end subroutine
32
33     integer function omp_get_default_device()
34         omp_get_default_device = 0
35         return
36     end function
37
38     integer function omp_get_num_devices()
39         omp_get_num_devices = 0
40         return
41     end function
42
43     integer function omp_get_num_teams()
44         omp_get_num_teams = 1
45         return
46     end function
47
48     integer function omp_get_team_num()
49         omp_get_team_num = 0
50         return
51     end function
52
53
54

```

```

1      logical function omp_is_initial_device()
2          omp_is_initial_device = .true.
3          return
4      end function
5
6      subroutine omp_init_lock(lock)
7          ! lock is 0 if the simple lock is not initialized
8          !          -1 if the simple lock is initialized but not set
9          !          1 if the simple lock is set
10         include 'omp_lib_kinds.h'
11         integer(kind=omp_lock_kind) lock
12
13         lock = -1
14         return
15     end subroutine
16
17     subroutine omp_destroy_lock(lock)
18         include 'omp_lib_kinds.h'
19         integer(kind=omp_lock_kind) lock
20
21         lock = 0
22         return
23     end subroutine
24
25     subroutine omp_set_lock(lock)
26         include 'omp_lib_kinds.h'
27         integer(kind=omp_lock_kind) lock
28
29         if (lock .eq. -1) then
30             lock = 1
31         elseif (lock .eq. 1) then
32             print *, 'error: deadlock in using lock variable'
33             stop
34         else
35             print *, 'error: lock not initialized'
36             stop
37         endif
38         return
39     end subroutine
40
41     subroutine omp_unset_lock(lock)
42         include 'omp_lib_kinds.h'
43         integer(kind=omp_lock_kind) lock
44
45         if (lock .eq. 1) then
46             lock = -1
47         elseif (lock .eq. -1) then
48             print *, 'error: lock not set'
49             stop
50         else
51             print *, 'error: lock not initialized'
52             stop
53         endif
54

```

```

1         return
2     end subroutine
3
4
5
6     logical function omp_test_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            omp_test_lock = .true.
13        elseif (lock .eq. 1) then
14            omp_test_lock = .false.
15        else
16            print *, 'error: lock not initialized'
17            stop
18        endif
19
20        return
21    end function
22
23    subroutine omp_init_nest_lock(nlock)
24        ! nlock is
25        ! 0 if the nestable lock is not initialized
26        ! -1 if the nestable lock is initialized but not set
27        ! 1 if the nestable lock is set
28        ! no use count is maintained
29        include 'omp_lib_kinds.h'
30        integer(kind=omp_nest_lock_kind) nlock
31
32        nlock = -1
33
34        return
35    end subroutine
36
37    subroutine omp_destroy_nest_lock(nlock)
38        include 'omp_lib_kinds.h'
39        integer(kind=omp_nest_lock_kind) nlock
40
41        nlock = 0
42
43        return
44    end subroutine
45

```

```

1      subroutine omp_set_nest_lock(nlock)
2          include 'omp_lib_kinds.h'
3          integer(kind=omp_nest_lock_kind) nlock
4
5          if (nlock .eq. -1) then
6              nlock = 1
7          elseif (nlock .eq. 0) then
8              print *, 'error: nested lock not initialized'
9              stop
10         else
11             print *, 'error: deadlock using nested lock variable'
12             stop
13         endif
14
15         return
16     end subroutine
17
18     subroutine omp_unset_nest_lock(nlock)
19         include 'omp_lib_kinds.h'
20         integer(kind=omp_nest_lock_kind) nlock
21
22         if (nlock .eq. 1) then
23             nlock = -1
24         elseif (nlock .eq. 0) then
25             print *, 'error: nested lock not initialized'
26             stop
27         else
28             print *, 'error: nested lock not set'
29             stop
30         endif
31
32         return
33     end subroutine
34
35     integer function omp_test_nest_lock(nlock)
36         include 'omp_lib_kinds.h'
37         integer(kind=omp_nest_lock_kind) nlock
38
39         if (nlock .eq. -1) then
40             nlock = 1
41             omp_test_nest_lock = 1
42         elseif (nlock .eq. 1) then
43             omp_test_nest_lock = 0
44         else
45             print *, 'error: nested lock not initialized'
46             stop
47         endif
48
49         return
50     end function
51
52
53
54

```

```
1      double precision function omp_get_wtime()
2          ! this function does not provide a working
3          ! wall clock timer. replace it with a version
4          ! customized for the target machine.
5
6          omp_get_wtime = 0.0d0
7
8          return
9      end function
10
11     double precision function omp_get_wtick()
12         ! this function does not provide a working
13         ! clock tick function. replace it with
14         ! a version customized for the target machine.
15         double precision one_year
16         parameter (one_year=365.d0*86400.d0)
17
18         omp_get_wtick = one_year
19
20         return
21     end function
22
23
```

This page intentionally left blank.

2 **OpenMP C and C++ Grammar**

3



4 **B.1 Notation**

5 The grammar rules consist of the name for a non-terminal, followed by a colon,
6 followed by replacement alternatives on separate lines.

7 The syntactic expression $term_{opt}$ indicates that the term is optional within the
8 replacement.

9 The syntactic expression $term_{optseq}$ is equivalent to $term-seq_{opt}$ with the following
10 additional rules:

11 *term-seq* :

12 *term*

13 *term-seq term*

14 *term-seq , term*



1 B.2 Rules

2 The notation is described in Section 6.1 of the C standard. This grammar appendix
3 shows the extensions to the base language grammar for the OpenMP C and C++
4 directives.
5

6 *statement-seq:* C++
7 *statement*
8 *openmp-directive*
9 *statement-seq statement*
10 *statement-seq openmp-directive*

▲ C++ ▲

11 *statement-list:* C90
12 *statement*
13 *openmp-directive*
14 *statement-list statement*
15 *statement-list openmp-directive*

▲ C90 ▲

16 *block-item:* C99
17 *declaration*
18 *statement*
19 *openmp-directive*

▲ C99 ▲

20

1 *statement:*

2 */* standard statements */*

3 *openmp-construct*

4 *declaration-definition:*

5 */* Any C or C++ declaration or definition statement */*

6 *function-statement:*

7 */* C or C++ function definition or declaration */*

8 *declarations-definitions-seq:*

9 *declaration-definition*

10 *declarations-definitions-seq declaration-definition*

11 *openmp-construct:*

12 *parallel-construct*

13 *for-construct*

14 *sections-construct*

15 *single-construct*

16 *simd-construct*

17 *for-simd-construct*

18 *parallel-for-simd-construct*

19 *target-data-construct*

20 *target-construct*

21 *target-update-construct*

22 *teams-construct*

23 *distribute-construct*

24 *distribute-simd-construct*

25 *distribute-parallel-for-construct*

26 *distribute-parallel-for-simd-construct*

27 *target-teams-construct*

1 *teams-distribute-construct*
2 *teams-distribute-simd-construct*
3 *target-teams-distribute-construct*
4 *target-teams-distribute-simd-construct*
5 *teams-distribute-parallel-for-construct*
6 *target-teams-distribute-parallel-for-construct*
7 *teams-distribute-parallel-for-simd-construct*
8 *target-teams-distribute-parallel-for-simd-construct*
9 *parallel-for-construct*
10 *parallel-sections-construct*
11 *task-construct*
12 *master-construct*
13 *critical-construct*
14 *atomic-construct*
15 *ordered-construct*
16 *openmp-directive:*
17 *barrier-directive*
18 *taskwait-directive*
19 *taskyield-directive*
20 *flush-directive*
21 *structured-block:*
22 *statement*
23 *parallel-construct:*
24 *parallel-directive structured-block*
25 *parallel-directive:*
26 **# pragma omp parallel** *parallel-clause_{optseq} new-line*
27
28

```

1      parallel-clause:
2          unique-parallel-clause
3          data-default-clause
4          data-privatization-clause
5          data-privatization-in-clause
6          data-sharing-clause
7          data-reduction-clause
8      unique-parallel-clause:
9          if-clause
10         num_threads ( expression )
11         copyin ( variable-list )
12     for-construct:
13         for-directive iteration-statement
14     for-directive:
15         # pragma omp for for-clauseoptseq new-line
16
17     for-clause:
18         unique-for-clause
19         data-privatization-clause
20         data-privatization-in-clause
21         data-privatization-out-clause
22         data-reduction-clause
23         nowait
24     unique-for-clause:
25         ordered
26         schedule ( schedule-kind )
27         schedule ( schedule-kind , expression )
28         collapse ( expression )

```

```

1      schedule-kind:
2          static
3          dynamic
4          guided
5          auto
6          runtime
7      sections-construct:
8          sections-directive section-scope
9      sections-directive:
10         # pragma omp sections sections-clauseoptseq new-line
11      sections-clause:
12         data-privatization-clause
13         data-privatization-in-clause
14         data-privatization-out-clause
15         data-reduction-clause
16         nowait
17      section-scope:
18         { section-sequence }
19      section-sequence:
20         section-directiveopt structured-block
21         section-sequence section-directive structured-block
22      section-directive:
23         # pragma omp section new-line
24      single-construct:
25         single-directive structured-block
26      single-directive:
27         # pragma omp single single-clauseoptseq new-line
28

```

1 *single-clause:*

2 *unique-single-clause*

3 *data-privatization-clause*

4 *data-privatization-in-clause*

5 **nowait**

6 *unique-single-clause:*

7 **copyprivate** (*variable-list*)

8 *simd-construct:*

9 *simd-directive iteration-statement*

10 *simd-directive:*

11 **# pragma omp simd** *simd-clause_{optseq} new-line*

12 *simd-clause:*

13 **collapse** (*expression*)

14 *aligned-clause*

15 *linear-clause*

16 *uniform-clause*

17 *data-reduction-clause*

18 *inbranch-clause*

19

20 *inbranch-clause:*

21 **inbranch**

22 **notinbranch**

23 *uniform-clause:*

24 **uniform** (*variable-list*)

25 *linear-clause:*

26 **linear** (*variable-list*)

27 **linear** (*variable-list : expression*)

1 *aligned-clause:*

2 **aligned** (*variable-list*)

3 **aligned** (*variable-list* : *expression*)

4 *declare-simd-construct:*

5 *declare-simd-directive-seq* *function-statement*

6 *declare-simd-directive-seq:*

7 *declare-simd-directive*

8 *declare-simd-directive-seq* *declare-simd-directive*

9 *declare-simd-directive:*

10 **# pragma omp declare simd** *declare-simd-clause*_{optseq} *new-line*

11 *declare-simd-clause:*

12 **simdlen** (*expression*)

13 *aligned-clause*

14 *linear-clause*

15 *uniform-clause*

16 *data-reduction-clause*

17 *inbranch-clause*

18 *for-simd-construct:*

19 *for-simd-directive* *iteration-statement*

20

21 *for-simd-directive:*

22 **# pragma omp for simd** *for-simd-clause*_{optseq} *new-line*

23 *for-simd-clause:*

24 *for-clause*

25 *simd-clause*

26 *parallel-for-simd-construct:*

27 *parallel-for-simd-directive* *iteration-statement*

```

1      parallel-for-simd-directive:
2          # pragma omp parallel for simd parallel-for-simd-clauseoptseq new-line
3      parallel-for-simd-clause:
4          parallel-for-clause
5          simd-clause
6      target-data-construct:
7          target-data-directive structured-block
8      target-data-directive:
9          # pragma omp target data target-data-clauseoptseq new-line
10     target-data-clause:
11         device-clause
12         map-clause
13         if-clause
14     device-clause:
15         device ( expression )
16     map-clause:
17         map ( map-typeopt variable-array-section-list )
18     map-type:
19         alloc:
20         to:
21         from:
22         tofrom:
23     target-construct:
24         target-directive structured-block
25     target-directive:
26         # pragma omp target target-clauseoptseq new-line
27

```

```

1      target-clause:
2          device-clause
3          map-clause
4          if-clause
5      target-update-construct:
6          target-update-directive structured-block
7      target-update-directive:
8          # pragma omp target update target-update-clauseseq new-line
9      target-update-clause:
10         motion-clause
11         device-clause
12         if-clause
13     motion-clause:
14         to ( variable-array-section-list )
15         from ( variable-array-section-list )
16     declare-target-construct:
17         declare-target-directive declarations-definitions-seq end-declare-target-directive
18     declare-target-directive:
19         # pragma omp declare target new-line
20     end-declare-target-directive:
21         # pragma omp end declare target new-line
22     teams-construct:
23         teams-directive structured-block
24     teams-directive:
25         # pragma omp teams teams-clauseoptseq new-line
26
27

```

1 *teams-clause:*

2 **num_teams** (*expression*)

3 **thread_limit** (*expression*)

4 *data-default-clause*

5 *data-privatization-clause*

6 *data-privatization-in-clause*

7 *data-sharing-clause*

8 *data-reduction-clause*

9 *distribute-construct:*

10 *distribute-directive iteration-statement*

11 *distribute-directive:*

12 **# pragma omp distribute** *distribute-clause*_{optseq} *new-line*

13 *distribute-clause:*

14 *data-privatization-clause*

15 *data-privatization-in-clause*

16 **collapse** (*expression*)

17 **dist_schedule** (**static**)

18 **dist_schedule** (**static** , *expression*)

19 *distribute-simd-construct:*

20 *distribute-simd-directive iteration-statement*

21 *distribute-simd-directive:*

22 **#pragma omp distribute simd** *distribute-simd-clause*_{optseq} *new-line*

23 *distribute-simd-clause:*

24 *distribute-clause*

25 *simd-clause*

26 *distribute-parallel-for-construct:*

27 *distribute-parallel-for-directive iteration-statement*

28

1 *distribute-parallel-for-directive:*
2 **#pragma omp distribute parallel for** *distribute-parallel-for-clause_{optseq}*
3 *new-line*
4 *distribute-parallel-for-clause:*
5 *distribute-clause*
6 *parallel-for-clause*
7 *distribute-parallel-for-simd-construct:*
8 *distribute-parallel-for-simd-directive iteration-statement*
9 *distribute-parallel-for-simd-directive:*
10 **#pragma omp distribute parallel for** *distribute-parallel-for-simd-*
11 *clause_{optseq} new-line*
12 *distribute-parallel-for-simd-clause:*
13 *distribute-clause*
14 *parallel-for-simd-clause*
15 *target-teams-construct:*
16 *target-teams-directive iteration-statement*
17 *target-teams-directive:*
18 **#pragma omp target teams** *target-teams-clause_{optseq} new-line*
19 *target-teams-clause:*
20 *target-clause*
21 *teams-clause*
22 *teams-distribute-construct:*
23 *teams-distribute-directive iteration-statement*
24 *teams-distribute-directive:*
25 **#pragma omp teams distribute** *teams-distribute-clause_{optseq} new-line*
26 *teams-distribute-clause:*
27 *teams-clause*
28 *distribute-clause*
29 *teams-distribute-simd-construct:*
30 *teams-distribute-simd-directive iteration-statement*
31

```

1      teams-distribute-simd-directive:
2          #pragma omp teams distribute simd teams-distribute-simd-clauseoptseq
3      new-line
4      teams-distribute-simd-clause:
5          teams-clause
6          distribute-simd-clause
7      target-teams-distribute-construct:
8          target-teams-distribute-directive iteration-statement
9      target-teams-distribute-directive:
10         #pragma omp target teams distribute target-teams-distribute-clauseoptseq
11     new-line
12     target-teams-distribute-clause:
13         target-clause
14         teams-distribute-clause
15     target-teams-distribute-simd-construct:
16         target-teams-distribute-simd-directive iteration-statement
17     target-teams-distribute-simd-directive:
18         #pragma omp target teams distribute simd target-teams-distribute-
19     simd-clauseoptseq new-line
20     target-teams-distribute-simd-clause:
21         target-clause
22         teams-distribute-simd-clause
23     teams-distribute-parallel-for-construct:
24         teams-distribute-parallel-for-directive iteration-statement
25     teams-distribute-parallel-for-directive:
26         #pragma omp teams distribute parallel for teams-distribute-
27     parallel-for-clauseoptseq new-line
28     teams-distribute-parallel-for-clause:
29         teams-clause
30         distribute-parallel-for-clause
31

```

1 *target-teams-distribute-parallel-for-construct:*

2 *target-teams-distribute-parallel-for-directive iteration-statement*

3 *target-teams-distribute-parallel-for-directive:*

4 **#pragma omp teams distribute parallel for** *target-teams-distribute-*
5 *parallel-for-clause_{optseq} new-line*

6 *target-teams-distribute-parallel-for-clause:*

7 *target-clause*

8 *teams-distribute-parallel-for-clause*

9 *teams-distribute-parallel-for-simd-construct:*

10 *teams-distribute-parallel-for-simd-directive iteration-statement*

11 *teams-distribute-parallel-for-simd-directive:*

12 **#pragma omp teams distribute parallel for simd** *teams-distribute-*
13 *parallel-for-simd-clause_{optseq} new-line*

14 *teams-distribute-parallel-for-simd-clause:*

15 *teams-clause*

16 *distribute-parallel-for-simd-clause*

17 *target-teams-distribute-parallel-for-simd-construct:*

18 *target-teams-distribute-parallel-for-simd-directive iteration-statement*

19 *target-teams-distribute-parallel-for-simd-directive:*

20 **#pragma omp target teams distribute parallel for simd** *target-*
21 *teams-distribute-parallel-for-simd-clause_{optseq} new-line*

22 *target-teams-distribute-parallel-for-simd-clause:*

23 *target-clause*

24 *teams-distribute-parallel-for-simd-clause*

25 *task-construct:*

26 *task-directive structured-block*

27 *task-directive:*

28 **# pragma omp task** *task-clause_{optseq} new-line*

29

30

31

1 *task-clause:*

2 *unique-task-clause*

3 *data-default-clause*

4 *data-privatization-clause*

5 *data-privatization-in-clause*

6 *data-sharing-clause*

7 *unique-task-clause:*

8 *if-clause*

9 **final** (*scalar-expression*)

10 **untied**

11 **mergeable**

12 **depend** (*dependence-type* : *variable-array-section-list*)

13 *dependence-type:*

14 **in**

15 **out**

16 **inout**

17 *parallel-for-construct:*

18 *parallel-for-directive iteration-statement*

19 *parallel-for-directive:*

20 **# pragma omp parallel for** *parallel-for-clause*_{optseq} *new-line*

21 *parallel-for-clause:*

22 *unique-parallel-clause*

23 *unique-for-clause*

24 *data-default-clause*

25 *data-privatization-clause*

26 *data-privatization-in-clause*

27 *data-privatization-out-clause*

1 *data-sharing-clause*

2 *data-reduction-clause*

3 *parallel-sections-construct:*

4 *parallel-sections-directive section-scope*

5 *parallel-sections-directive:*

6 **# pragma omp parallel sections** *parallel-sections-clause*_{optseq} *new-line*

7 *parallel-sections-clause:*

8 *unique-parallel-clause*

9 *data-default-clause*

10 *data-privatization-clause*

11 *data-privatization-in-clause*

12 *data-privatization-out-clause*

13 *data-sharing-clause*

14 *data-reduction-clause*

15 *master-construct:*

16 *master-directive structured-block*

17 *master-directive:*

18 **# pragma omp master** *new-line*

19 *critical-construct:*

20 *critical-directive structured-block*

21 *critical-directive:*

22 **# pragma omp critical** *region-phrase*_{opt} *new-line*

23 *region-phrase:*

24 (*identifier*)

25 *barrier-directive:*

26 **# pragma omp barrier** *new-line*

27 *taskwait-directive:*

28 **# pragma omp taskwait** *new-line*

```

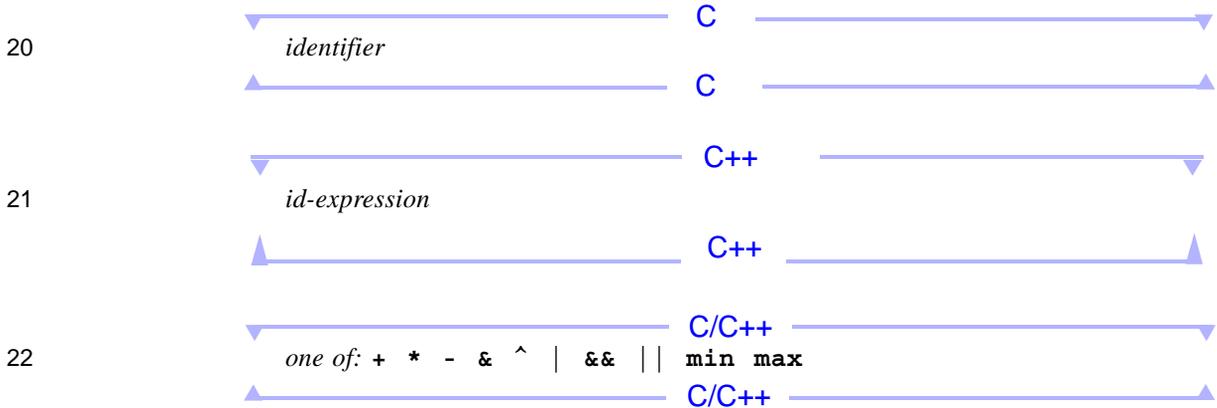
1      taskgroup-construct:
2          taskgroup-directive structured-block
3
4      taskgroup-directive:
5          # pragma omp taskgroup new-line
6
7      taskyield-directive:
8          # pragma omp taskyield new-line
9
10     atomic-construct:
11         atomic-directive expression-statement
12         atomic-directive structured block
13
14     atomic-directive:
15         # pragma omp atomic atomic-clauseopt seq_cst-clauseopt new-line
16
17     atomic-clause:
18         read
19         write
20         update
21         capture
22
23     seq_cst-clause:
24         seq_cst
25
26     flush-directive:
27         # pragma omp flush flush-varsopt new-line
28
29     flush-vars:
30         ( variable-list )
31
32     ordered-construct:
33         ordered-directive structured-block
34
35     ordered-directive:
36         # pragma omp ordered new-line
37
38     cancel-directive:
39         # pragma omp cancel construct-type-clause if-clauseopt new-line

```

```

1      construct-type-clause:
2          parallel
3          sections
4          for
5          taskgroup
6      cancellation-point-directive:
7          # pragma omp cancellation point construct-type-clause new-line
8      declaration:
9          /* standard declarations */
10     threadprivate-directive
11     declare-simd-directive
12     declare-target-construct
13     declare-reduction-directive
14     threadprivate-directive:
15         # pragma omp threadprivate ( variable-list ) new-line
16     declare-reduction-directive:
17         # pragma omp declare reduction ( reduction-identifier :
18         reduction-type-list : expression ) initializer-clauseopt new-line
19     reduction-identifier:

```



```

1      reduction-type-list:
2          type-id
3          reduction-type-list, type-id
4      initializer-clause:
5          initializer ( identifier = initializer )
6          initializer ( identifier ( argument-expression-list ) )
7
8          initializer ( identifier initializer )
9          initializer ( id-expression ( expression-list ) )
10
11      data-default-clause:
12          default ( shared )
13          default ( none )
14      data-privatization-clause:
15          private ( variable-list )
16      data-privatization-in-clause:
17          firstprivate ( variable-list )
18      data-privatization-out-clause:
19          lastprivate ( variable-list )
20      data-sharing-clause:
21          shared ( variable-list )
22      data-reduction-clause:
23          reduction ( reduction-identifier : variable-list )
24      if-clause:
25          if ( scalar-expression )

```

		C	
1	<i>array-section:</i>		
2	<i>identifier array-section-subscript</i>		
3	<i>variable-list:</i>		
4	<i>identifier</i>		
5	<i>variable-list , identifier</i>		
6	<i>variable-array-section-list:</i>		
7	<i>identifier</i>		
8	<i>array-section</i>		
9	<i>variable-array-section-list , identifier</i>		
10	<i>variable-array-section-list , array-section</i>		
		C	
		C++	
11	<i>array-section:</i>		
12	<i>id-expression array-section-subscript</i>		
13	<i>variable-list:</i>		
14	<i>id-expression</i>		
15	<i>variable-list , id-expression</i>		
16	<i>variable-array-section-list:</i>		
17	<i>id-expression</i>		
18	<i>array-section</i>		
19	<i>variable-array-section-list , id-expression</i>		
20	<i>variable-array-section-list , array-section</i>		
		C++	
21			
22			
23			
24			

1 *array-section-subscript:*
2 *array-section-subscript [expression_{opt} : expression_{opt}]*
3 *array-section-subscript [expression]*
4 *[expression_{opt} : expression_{opt}]*
5 *[expression]*

This page intentionally left blank.

2 **Interface Declarations**

3 This appendix gives examples of the C/C++ header file, the Fortran **include** file and
4 Fortran **module** that shall be provided by implementations as specified in Chapter 3. It
5 also includes an example of a Fortran 90 generic interface for a library routine. This is a
6 non-normative section, implementation files may differ.

1 C.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 schedule kinds
14  */
15 typedef enum omp_sched_t
16 {
17     omp_sched_static = 1,
18     omp_sched_dynamic = 2,
19     omp_sched_guided = 3,
20     omp_sched_auto = 4
21 } , Add vendor specific schedule constants here */
22 } omp_sched_t;
23
24 /*
25  * define the proc bind values
26  */
27 typedef enum omp_proc_bind_t
28 {
29     omp_proc_bind_false = 0,
30     omp_proc_bind_true = 1,
31     omp_proc_bind_master = 2,
32     omp_proc_bind_close = 3,
33     omp_proc_bind_spread = 4
34 } omp_proc_bind_t;
35
36 /*
37  * exported OpenMP functions
38  */
39 #ifdef __cplusplus
40 extern      "C"
41 {
42 #endif
43
44 extern void  omp_set_num_threads(int num_threads);
45 extern int   omp_get_num_threads(void);
46 extern int   omp_get_max_threads(void);
47 extern int   omp_get_thread_num(void);
48 extern int   omp_get_num_procs(void);
49 extern int   omp_in_parallel(void);
50 extern void  omp_set_dynamic(int dynamic_threads);
```

```

1      extern int      omp_get_dynamic(void);
2      extern void    omp_set_nested(int nested);
3      extern int      omp_get_cancellation(void);
4      extern int      omp_get_nested(void);
5      extern void    omp_set_schedule(omp_sched_t kind, int modifier);
6      extern void    omp_get_schedule(omp_sched_t *kind, int *modifier);
7      extern int      omp_get_thread_limit(void);
8      extern void    omp_set_max_active_levels(int max_active_levels);
9      extern int      omp_get_max_active_levels(void);
10     extern int      omp_get_level(void);
11     extern int      omp_get_ancestor_thread_num(int level);
12     extern int      omp_get_team_size(int level);
13     extern int      omp_get_active_level(void);
14     extern int      omp_in_final(void);
15     extern omp_proc_bind_t omp_get_proc_bind(void);
16     extern void    omp_set_default_device(int device_num);
17     extern int      omp_get_default_device(void);
18     extern int      omp_get_num_devices(void);
19     extern int      omp_get_num_teams(void);
20     extern int      omp_get_team_num(void);
21     extern int      omp_is_initial_device(void);
22
23     extern void    omp_init_lock(omp_lock_t *lock);
24     extern void    omp_destroy_lock(omp_lock_t *lock);
25     extern void    omp_set_lock(omp_lock_t *lock);
26     extern void    omp_unset_lock(omp_lock_t *lock);
27     extern int      omp_test_lock(omp_lock_t *lock);
28
29     extern void    omp_init_nest_lock(omp_nest_lock_t *lock);
30     extern void    omp_destroy_nest_lock(omp_nest_lock_t *lock);
31     extern void    omp_set_nest_lock(omp_nest_lock_t *lock);
32     extern void    omp_unset_nest_lock(omp_nest_lock_t *lock);
33     extern int      omp_test_nest_lock(omp_nest_lock_t *lock);
34
35     extern double  omp_get_wtime(void);
36     extern double  omp_get_wtick(void);
37
38     #ifdef __cplusplus
39     }
40     #endif
41
42     #endif

```

1 C.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     ! this selects an integer that is large enough to hold a 64 bit integer  
7     parameter ( omp_lock_kind = selected_int_kind( 10 ) )  
8     parameter ( omp_nest_lock_kind = selected_int_kind( 10 ) )  
9     integer      omp_sched_kind  
10    ! this selects an integer that is large enough to hold a 32 bit integer  
11    parameter ( omp_sched_kind = selected_int_kind( 8 ) )  
12    integer ( omp_sched_kind ) omp_sched_static  
13    parameter ( omp_sched_static = 1 )  
14    integer ( omp_sched_kind ) omp_sched_dynamic  
15    parameter ( omp_sched_dynamic = 2 )  
16    integer ( omp_sched_kind ) omp_sched_guided  
17    parameter ( omp_sched_guided = 3 )  
18    integer ( omp_sched_kind ) omp_sched_auto  
19    parameter ( omp_sched_auto = 4 )  
20    integer omp_proc_bind_kind  
21    parameter ( omp_proc_bind_kind = selected_int_kind( 8 ) )  
22    integer ( omp_proc_bind_kind ) omp_proc_bind_false  
23    parameter ( omp_proc_bind_false = 0 )  
24    integer ( omp_proc_bind_kind ) omp_proc_bind_true  
25    parameter ( omp_proc_bind_true = 1 )  
26    integer ( omp_proc_bind_kind ) omp_proc_bind_master  
27    parameter ( omp_proc_bind_master = 2 )  
28    integer ( omp_proc_bind_kind ) omp_proc_bind_close  
29    parameter ( omp_proc_bind_close = 3 )  
30    integer ( omp_proc_bind_kind ) omp_proc_bind_spread  
31    parameter ( omp_proc_bind_spread = 4 )
```

```
32 omp_lib.h:  
33 ! default integer type assumed below  
34 ! default logical type assumed below  
35 ! OpenMP API v4.0  
36  
37     include 'omp_lib_kinds.h'  
38     integer      openmp_version  
39     parameter ( openmp_version = 201307 )  
40  
41     external omp_set_num_threads  
42     external omp_get_num_threads  
43     integer omp_get_num_threads  
44     external omp_get_max_threads  
45     integer omp_get_max_threads  
46     external omp_get_thread_num  
47     integer omp_get_thread_num  
48     external omp_get_num_procs
```

```

1         integer    omp_get_num_procs
2         external  omp_in_parallel
3         logical   omp_in_parallel
4         external  omp_set_dynamic
5         external  omp_get_dynamic
6         logical   omp_get_dynamic
7         external  omp_get_cancellation
8         integer   omp_get_cancellation
9         external  omp_set_nested
10        external  omp_get_nested
11        logical   omp_get_nested
12        external  omp_set_schedule
13        external  omp_get_schedule
14        external  omp_get_thread_limit
15        integer   omp_get_thread_limit
16        external  omp_set_max_active_levels
17        external  omp_get_max_active_levels
18        integer   omp_get_max_active_levels
19        external  omp_get_level
20        integer   omp_get_level
21        external  omp_get_ancestor_thread_num
22        integer   omp_get_ancestor_thread_num
23        external  omp_get_team_size
24        integer   omp_get_team_size
25        external  omp_get_active_level
26        integer   omp_get_active_level
27        external  omp_set_default_device
28        external  omp_get_default_device
29        integer   omp_get_default_device
30        external  omp_get_num_devices
31        integer   omp_get_num_devices
32        external  omp_get_num_teams
33        integer   omp_get_num_teams
34        external  omp_get_team_num
35        integer   omp_get_team_num
36        external  omp_is_initial_device
37        logical   omp_is_initial_device
38
39        external  omp_in_final
40        logical   omp_in_final
41
42        integer ( omp_proc_bind_kind ) omp_get_proc_bind
43        external  omp_get_proc_bind
44
45        external  omp_init_lock
46        external  omp_destroy_lock
47        external  omp_set_lock
48        external  omp_unset_lock
49        external  omp_test_lock
50        logical   omp_test_lock
51
52        external  omp_init_nest_lock
53        external  omp_destroy_nest_lock
54        external  omp_set_nest_lock

```

```
1         external omp_unset_nest_lock
2         external omp_test_nest_lock
3         integer  omp_test_nest_lock
4
5         external omp_get_wtick
6         double precision  omp_get_wtick
7         external omp_get_wtime
8         double precision  omp_get_wtime
9
```

1 C.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_sched_kind = selected_int_kind( 8 )
10        integer(kind=omp_sched_kind), parameter ::
11        & omp_sched_static = 1
12        integer(kind=omp_sched_kind), parameter ::
13        & omp_sched_dynamic = 2
14        integer(kind=omp_sched_kind), parameter ::
15        & omp_sched_guided = 3
16        integer(kind=omp_sched_kind), parameter ::
17        & omp_sched_auto = 4
18        integer, parameter :: omp_proc_bind_kind = selected_int_kind( 8 )
19        integer (kind=omp_proc_bind_kind), parameter ::
20        & omp_proc_bind_false = 0
21        integer (kind=omp_proc_bind_kind), parameter ::
22        & omp_proc_bind_true = 1
23        integer (kind=omp_proc_bind_kind), parameter ::
24        & omp_proc_bind_master = 2
25        integer (kind=omp_proc_bind_kind), parameter ::
26        & omp_proc_bind_close = 3
27        integer (kind=omp_proc_bind_kind), parameter ::
28        & omp_proc_bind_spread = 4
29        end module omp_lib_kinds
30
31    module omp_lib
32
33        use omp_lib_kinds
34        !                                     OpenMP API v4.0
35        integer, parameter :: openmp_version = 201307
36
37        interface
38
39            subroutine omp_set_num_threads (number_of_threads_expr)
40                integer, intent(in) :: number_of_threads_expr
41            end subroutine omp_set_num_threads
42
43            function omp_get_num_threads ()
44                integer :: omp_get_num_threads
45            end function omp_get_num_threads
46
47            function omp_get_max_threads ()
48                integer :: omp_get_max_threads
49            end function omp_get_max_threads
```

```

1
2      function omp_get_thread_num ()
3          integer :: omp_get_thread_num
4      end function omp_get_thread_num
5
6
7      function omp_get_num_procs ()
8          integer :: omp_get_num_procs
9      end function omp_get_num_procs
10
11     function omp_in_parallel ()
12         logical :: omp_in_parallel
13     end function omp_in_parallel
14
15     subroutine omp_set_dynamic (enable_expr)
16         logical, intent(in) :: enable_expr
17     end subroutine omp_set_dynamic
18
19     function omp_get_dynamic ()
20         logical :: omp_get_dynamic
21     end function omp_get_dynamic
22
23     function omp_get_cancellation ()
24         integer :: omp_get_cancellation
25     end function omp_get_cancellation
26
27     subroutine omp_set_nested (enable_expr)
28         logical, intent(in) :: enable_expr
29     end subroutine omp_set_nested
30
31     function omp_get_nested ()
32         logical :: omp_get_nested
33     end function omp_get_nested
34
35     subroutine omp_set_schedule (kind, modifier)
36         use omp_lib_kinds
37         integer(kind=omp_sched_kind), intent(in) :: kind
38         integer, intent(in) :: modifier
39     end subroutine omp_set_schedule
40
41     subroutine omp_get_schedule (kind, modifier)
42         use omp_lib_kinds
43         integer(kind=omp_sched_kind), intent(out) :: kind
44         integer, intent(out)::modifier
45     end subroutine omp_get_schedule
46
47     function omp_get_thread_limit()
48         integer :: omp_get_thread_limit
49     end function omp_get_thread_limit
50
51     subroutine omp_set_max_active_levels(var)
52         integer, intent(in) :: var
53     end subroutine omp_set_max_active_levels
54

```

```

1      function omp_get_max_active_levels()
2          integer :: omp_get_max_active_levels
3      end function omp_get_max_active_levels
4
5      function omp_get_level()
6          integer :: omp_get_level
7      end function omp_get_level
8
9      function omp_get_ancestor_thread_num(level)
10         integer, intent(in) :: level
11         integer :: omp_get_ancestor_thread_num
12     end function omp_get_ancestor_thread_num
13
14     function omp_get_team_size(level)
15         integer, intent(in) :: level
16         integer :: omp_get_team_size
17     end function omp_get_team_size
18
19     function omp_get_active_level()
20         integer :: omp_get_active_level
21     end function omp_get_active_level
22
23     function omp_in_final()
24         logical omp_in_final
25     end function omp_in_final
26
27     function omp_get_proc_bind( )
28         include 'omp_lib_kinds.h'
29         integer (kind=omp_proc_bind_kind) omp_get_proc_bind
30         omp_get_proc_bind = omp_proc_bind_false
31     end function omp_get_proc_bind
32
33     subroutine omp_set_default_device (device_num)
34         integer :: device_num
35     end subroutine omp_set_default_device
36
37     function omp_get_default_device ( )
38         integer :: omp_get_default_device
39     end function omp_get_default_device
40
41     function omp_get_num_devices ( )
42         integer :: omp_get_num_devices
43     end function omp_get_num_devices
44
45     function omp_get_num_teams ( )
46         integer :: omp_get_num_teams
47     end function omp_get_num_teams
48
49     function omp_get_team_num ( )
50         integer :: omp_get_team_num
51     end function omp_get_team_num
52
53
54

```

```

1      function omp_is_initial_device ()
2          logical :: omp_is_initial_device
3      end function omp_is_initial_device
4
5      subroutine omp_init_lock (var)
6          use omp_lib_kinds
7          integer (kind=omp_lock_kind), intent(out) :: var
8      end subroutine omp_init_lock
9
10     subroutine omp_destroy_lock (var)
11         use omp_lib_kinds
12         integer (kind=omp_lock_kind), intent(inout) :: var
13     end subroutine omp_destroy_lock
14
15     subroutine omp_set_lock (var)
16         use omp_lib_kinds
17         integer (kind=omp_lock_kind), intent(inout) :: var
18     end subroutine omp_set_lock
19
20     subroutine omp_unset_lock (var)
21         use omp_lib_kinds
22         integer (kind=omp_lock_kind), intent(inout) :: var
23     end subroutine omp_unset_lock
24
25     function omp_test_lock (var)
26         use omp_lib_kinds
27         logical :: omp_test_lock
28         integer (kind=omp_lock_kind), intent(inout) :: var
29     end function omp_test_lock
30
31     subroutine omp_init_nest_lock (var)
32         use omp_lib_kinds
33         integer (kind=omp_nest_lock_kind), intent(out) :: var
34     end subroutine omp_init_nest_lock
35
36     subroutine omp_destroy_nest_lock (var)
37         use omp_lib_kinds
38         integer (kind=omp_nest_lock_kind), intent(inout) :: var
39     end subroutine omp_destroy_nest_lock
40
41     subroutine omp_set_nest_lock (var)
42         use omp_lib_kinds
43         integer (kind=omp_nest_lock_kind), intent(inout) :: var
44     end subroutine omp_set_nest_lock
45
46     subroutine omp_unset_nest_lock (var)
47         use omp_lib_kinds
48         integer (kind=omp_nest_lock_kind), intent(inout) :: var
49     end subroutine omp_unset_nest_lock
50
51     function omp_test_nest_lock (var)
52         use omp_lib_kinds
53         integer :: omp_test_nest_lock
54         integer (kind=omp_nest_lock_kind), intent(inout) :: var

```

```
1         end function omp_test_nest_lock
2
3
4
5         function omp_get_wtick ()
6             double precision :: omp_get_wtick
7         end function omp_get_wtick
8
9         function omp_get_wtime ()
10            double precision :: omp_get_wtime
11        end function omp_get_wtime
12
13        end interface
14
15    end module omp_lib
```

1 C.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
4 with a generic 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
6 as follows:

```
interface omp_set_num_threads

    subroutine omp_set_num_threads_4(number_of_threads_expr)
        use omp_lib_kinds
        integer(4), intent(in) :: number_of_threads_expr
    end subroutine omp_set_num_threads_4

    subroutine omp_set_num_threads_8(number_of_threads_expr)
        use omp_lib_kinds
        integer(8), intent(in) :: number_of_threads_expr
    end subroutine omp_set_num_threads_8

end interface omp_set_num_threads
```

7

8

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 17).
- **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 20).
- **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 36).
- **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 47).
- **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

1 can be fulfilled is implementation defined. If not, the number of threads in the team
2 and their mapping to places become implementation defined (see Section 2.5.2 on
3 page 49).

- 4 • **Loop directive:** the integer type (or kind, for Fortran) used to compute the iteration
5 count of a collapsed loop is implementation defined. The effect of the
6 **schedule(runtime)** clause when the *run-sched-var* ICV is set to **auto** is
7 implementation defined. See Section 2.7.1 on page 53.
- 8 • **sections construct:** the method of scheduling the structured blocks among threads
9 in the team is implementation defined (see Section 2.7.2 on page 60).
- 10 • **single construct:** the method of choosing a thread to execute the structured block
11 is implementation defined (see Section 2.7.3 on page 63)
- 12 • **simd construct:** the integer type (or kind, for Fortran) used to compute the iteration
13 count for the collapsed loop is implementation defined. The number of iterations that
14 are executed concurrently at any given time is implementation defined. If the
15 **aligned** clause is not specified, the assumed alignment is implementation defined
16 (see Section 2.8.1 on page 68).
- 17 • **declare simd construct:** if the **simdlen** clause is not specified, the number of
18 concurrent arguments for the function is implementation defined. If the **aligned**
19 clause is not specified, the assumed alignment is implementation defined (see
20 Section 2.8.2 on page 72).
- 21 • **teams construct:** the number of teams that are created is implementation defined but
22 less than or equal to the value of the **num_teams** clause if specified. The maximum
23 number of threads participating in the contention group that each team initiates is
24 implementation defined but less than or equal to the value of the **thread_limit**
25 clause if specified (see Section 2.9.5 on page 86).
- 26 • If no **dist_schedule** clause is specified then the schedule for the **distribute**
27 construct is implementation defined (see Section 2.9.6 on page 88).
- 28 • **atomic construct:** a compliant implementation may enforce exclusive access
29 between **atomic** regions that update different storage locations. The circumstances
30 under which this occurs are implementation defined. If the storage location
31 designated by *x* is not size-aligned (that is, if the byte alignment of *x* is not a multiple
32 of the size of *x*), then the behavior of the atomic region is implementation defined
33 (see Section 2.12.6 on page 127).
- 34 • **omp_set_num_threads routine:** if the argument is not a positive integer the
35 behavior is implementation defined (see Section 3.2.1 on page 189).
- 36 • **omp_set_schedule routine:** for implementation specific schedule types, the
37 values and associated meanings of the second argument are implementation defined.
38 (see Section 3.2.12 on page 203).

- 1 • **omp_set_max_active_levels routine**: when called from within any explicit
2 parallel region the binding thread set (and binding region, if required) for the
3 **omp_set_max_active_levels** region is implementation defined and the
4 behavior is implementation defined. If the argument is not a non-negative integer
5 then the behavior is implementation defined (see Section 3.2.15 on page 207).
- 6 • **omp_get_max_active_levels routine**: when called from within any explicit
7 parallel region the binding thread set (and binding region, if required) for the
8 **omp_get_max_active_levels** region is implementation defined (see
9 Section 3.2.16 on page 209).
- 10 • **OMP_SCHEDULE environment variable**: if the value of the variable does not
11 conform to the specified format then the result is implementation defined (see
12 Section 4.1 on page 238).
- 13 • **OMP_NUM_THREADS environment variable**: if any value of the list specified in the
14 **OMP_NUM_THREADS** environment variable leads to a number of threads that is
15 greater than the implementation can support, or if any value is not a positive integer,
16 then the result is implementation defined (see Section 4.2 on page 239).
- 17 • **OMP_PROC_BIND environment variable**: if the value is not **true**, **false**, or a
18 comma separated list of **master**, **close**, or **spread**, the behavior is
19 implementation defined. The behavior is also implementation defined if an initial
20 thread cannot be bound to the first place in the OpenMP place list (see Section 4.4 on
21 page 241).
- 22 • **OMP_DYNAMIC environment variable**: if the value is neither **true** nor **false** the
23 behavior is implementation defined (see Section 4.3 on page 240).
- 24 • **OMP_NESTED environment variable**: if the value is neither **true** nor **false** the
25 behavior is implementation defined (see Section on page 241).
- 26 • **OMP_STACKSIZE environment variable**: if the value does not conform to the
27 specified format or the implementation cannot provide a stack of the specified size
28 then the behavior is implementation defined (see Section 4.7 on page 244).
- 29 • **OMP_WAIT_POLICY environment variable**: the details of the **ACTIVE** and
30 **PASSIVE** behaviors are implementation defined (see Section 4.8 on page 245).
- 31 • **OMP_MAX_ACTIVE_LEVELS environment variable**: if the value is not a non-
32 negative integer or is greater than the number of parallel levels an implementation
33 can support then the behavior is implementation defined (see Section 4.9 on page
34 245).
- 35 • **OMP_THREAD_LIMIT environment variable**: if the requested value is greater than
36 the number of threads an implementation can support, or if the value is not a positive
37 integer, the behavior of the program is implementation defined (see Section 4.10 on
38 page 246).
- 39 • **OMP_PLACES environment variable**: the meaning of the numbers specified in the
40 environment variable and how the numbering is done are implementation defined.
41 The precise definitions of the abstract names are implementation defined. An

1 implementation may add implementation-defined abstract names as appropriate for
2 the target platform. When creating a place list of n elements by appending the
3 number n to an abstract name, the determination of which resources to include in the
4 place list is implementation defined. When requesting more resources than available,
5 the length of the place list is also implementation defined. The behavior of the
6 program is implementation defined when the execution environment cannot map a
7 numerical value (either explicitly defined or implicitly derived from an interval)
8 within the **OMP_PLACES** list to a processor on the target platform, or if it maps to an
9 unavailable processor. The behavior is also implementation defined when the
10 **OMP_PLACES** environment variable is defined using an abstract name (see
11 Section 4.5 on page 241).

- 12 • **Thread affinity policy:** if the affinity request for a **parallel** construct cannot be
13 fulfilled, the behavior of the thread affinity policy is implementation defined for that
14 **parallel** construct.

Fortran

- 15 • **threadprivate directive:** if the conditions for values of data in the threadprivate
16 objects of threads (other than an initial thread) to persist between two consecutive
17 active parallel regions do not all hold, the allocation status of an allocatable variable
18 in the second region is implementation defined (see Section 2.14.2 on page 150).
- 19 • **shared clause:** passing a shared variable to a non-intrinsic procedure may result in
20 the value of the shared variable being copied into temporary storage before the
21 procedure reference, and back out of the temporary storage into the actual argument
22 storage after the procedure reference. Situations where this occurs other than those
23 specified are implementation defined (see Section 2.14.3.2 on page 157).
- 24 • **Runtime library definitions:** it is implementation defined whether the include file
25 **omp_lib.h** or the module **omp_lib** (or both) is provided. It is implementation
26 defined whether any of the OpenMP runtime library routines that take an argument
27 are extended with a generic interface so arguments of different **KIND** type can be
28 accommodated (see Section 3.1 on page 188).

Fortran

2

Features History

3 This appendix summarizes the major changes between recent versions of the OpenMP
4 API since version 2.5.

5

E.1 Version 3.1 to 4.0 Differences

- 6 • Various changes throughout the specification were made to provide initial support of
7 Fortran 2003 (see Section 1.6 on page 22).
- 8 • C/C++ array syntax was extended to support array sections (see Section 2.4 on page
9 42).
- 10 • The `proc_bind` clause (see Section 2.5.2 on page 49), the `OMP_PLACES`
11 environment variable (see Section 4.5 on page 241), and the `omp_get_proc_bind`
12 runtime routine (see Section 3.2.22 on page 216) were added to support thread
13 affinity policies.
- 14 • SIMD constructs were added to support SIMD parallelism (see Section 2.8 on page
15 68).
- 16 • Device constructs (see Section 2.9 on page 77), the `OMP_DEFAULT_DEVICE`
17 environment variable (see Section 4.13 on page 248), the
18 `omp_set_default_device`, `omp_get_default_device`,
19 `omp_get_num_devices`, `omp_get_num_teams`, `omp_get_team_num`, and
20 `omp_is_initial_device` routines were added to support execution on devices.
- 21 • Implementation defined task scheduling points for untied tasks were removed (see
22 Section 2.11.3 on page 118).
- 23 • The `depend` clause (see Section 2.11.1.1 on page 116) was added to support task
24 dependences.
- 25 • The `taskgroup` construct (see Section 2.12.5 on page 126) was added to support
26 more flexible deep task synchronization.

- The **reduction** clause (see Section 2.14.3.6 on page 167) was extended and the **declare reduction** construct (see Section 2.15 on page 180) was added to support user defined reductions.
- The **atomic** construct (see Section 2.12.6 on page 127) 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 2.13.1 on page 140), the **cancellation point** construct (see Section 2.13.2 on page 143), the **omp_get_cancellation** runtime routine (see Section 3.2.9 on page 199) and the **OMP_CANCELLATION** environment variable (see Section 4.11 on page 246) were added to support the concept of cancellation.
- The **OMP_DISPLAY_ENV** environment variable (see Section 4.12 on page 247) was added to display the value of ICVs associated with the OpenMP environment variables.
- Examples (previously Appendix A) were moved to a separate document.

E.2 Version 3.0 to 3.1 Differences

- The **final** and **mergeable** clauses (see Section 2.11.1 on page 113) were added to the **task** construct to support optimization of task data environments.
- The **taskyield** construct (see Section 2.11.2 on page 117) was added to allow user-defined task scheduling points.
- The **atomic** construct (see Section 2.12.6 on page 127) 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 2.14.3.4 on page 162).
- Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 2.14.3.4 on page 162) and **lastprivate** (see Section 2.14.3.5 on page 164).
- New reduction operators **min** and **max** were added for C and C++
- The nesting restrictions in Section 2.16 on page 186 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 215) was added to support specialization of final task regions.

- 1 • The *nthreads-var* ICV has been modified to be a list of the number of threads to use
2 at each nested parallel region level. The value of this ICV is still set with the
3 **OMP_NUM_THREADS** environment variable (see Section 4.2 on page 239), but the
4 algorithm for determining the number of threads used in a parallel region has been
5 modified to handle a list (see Section 2.5.1 on page 47).
- 6 • The *bind-var* ICV has been added, which controls whether or not threads are bound
7 to processors (see Section 2.3.1 on page 35). The value of this ICV can be set with
8 the **OMP_PROC_BIND** environment variable (see Section 4.4 on page 241).
- 9 • Descriptions of examples (see Appendix A on page 221) were expanded and clarified.
- 10 • Replaced incorrect use of **omp_integer_kind** in Fortran interfaces (see
11 Section C.3 on page 293 and Section C.4 on page 298) with
12 **selected_int_kind(8)**.



13 E.3 Version 2.5 to 3.0 Differences

- 14 The concept of tasks has been added to the OpenMP execution model (see Section 1.2.4
15 on page 8 and Section 1.3 on page 14).
- 16 • The **task** construct (see Section 2.11 on page 113) has been added, which provides
17 a mechanism for creating tasks explicitly.
 - 18 • The **taskwait** construct (see Section 2.12.4 on page 125) has been added, which
19 causes a task to wait for all its child tasks to complete.
 - 20 • The OpenMP memory model now covers atomicity of memory accesses (see
21 Section 1.4.1 on page 17). The description of the behavior of **volatile** in terms of
22 **flush** was removed.
 - 23 • In Version 2.5, there was a single copy of the *nest-var*, *dyn-var*, *nthreads-var* and
24 *run-sched-var* internal control variables (ICVs) for the whole program. In Version
25 3.0, there is one copy of these ICVs per task (see Section 2.3 on page 34). As a result,
26 the **omp_set_num_threads**, **omp_set_nested** and **omp_set_dynamic**
27 runtime library routines now have specified effects when called from inside a
28 **parallel** region (see Section 3.2.1 on page 189, Section 3.2.7 on page 197 and
29 Section 3.2.10 on page 200).
 - 30 • The definition of active **parallel** region has been changed: in Version 3.0 a
31 **parallel** region is active if it is executed by a team consisting of more than one
32 thread (see Section 1.2.2 on page 2).
 - 33 • The rules for determining the number of threads used in a **parallel** region have
34 been modified (see Section 2.5.1 on page 47).
 - 35 • In Version 3.0, the assignment of iterations to threads in a loop construct with a
36 **static** schedule kind is deterministic (see Section 2.7.1 on page 53).

- 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 53).
- 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 53).
- 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 53).
- Fortran assumed-size arrays now have predetermined data-sharing attributes (see Section 2.14.1.1 on page 146).
- In Fortran, **firstprivate** is now permitted as an argument to the **default** clause (see Section 2.14.3.1 on page 156).
- 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 2.14.3.3 on page 159).
- In Version 3.0, Fortran allocatable arrays may appear in **private**, **firstprivate**, **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses. (see Section 2.14.2 on page 150, Section 2.14.3.3 on page 159, Section 2.14.3.4 on page 162, Section 2.14.3.5 on page 164, Section 2.14.3.6 on page 167, Section 2.14.4.1 on page 173 and Section 2.14.4.2 on page 175).
- In Version 3.0, static class members variables may appear in a **threadprivate** directive (see Section 2.14.2 on page 150).
- Version 3.0 makes clear where, and with which arguments, constructors and destructors of private and threadprivate class type variables are called (see Section 2.14.2 on page 150, Section 2.14.3.3 on page 159, Section 2.14.3.4 on page 162, Section 2.14.4.1 on page 173 and Section 2.14.4.2 on page 175)
- 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 203 and Section 3.2.13 on page 205).
- 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 35, Section 3.2.14 on page 206 and Section 4.10 on page 246).
- 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

- 1 with the `omp_get_max_active_levels` runtime library routine (see
2 Section 2.3.1 on page 35, Section 3.2.15 on page 207, Section 3.2.16 on page 209 and
3 Section 4.9 on page 245).
- 4 • The `stacksize-var` ICV has been added, which controls the stack size for threads that
5 the OpenMP implementation creates. The value of this ICV can be set with the
6 `OMP_STACKSIZE` environment variable (see Section 2.3.1 on page 35 and
7 Section 4.7 on page 244).
 - 8 • The `wait-policy-var` ICV has been added, which controls the desired behavior of
9 waiting threads. The value of this ICV can be set with the `OMP_WAIT_POLICY`
10 environment variable (see Section 2.3.1 on page 35 and Section 4.8 on page 245).
 - 11 • The `omp_get_level` runtime library routine has been added, which returns the
12 number of nested `parallel` regions enclosing the task that contains the call (see
13 Section 3.2.17 on page 210).
 - 14 • The `omp_get_ancestor_thread_num` runtime library routine has been added,
15 which returns, for a given nested level of the current thread, the thread number of the
16 ancestor (see Section 3.2.18 on page 211).
 - 17 • The `omp_get_team_size` runtime library routine has been added, which returns,
18 for a given nested level of the current thread, the size of the thread team to which the
19 ancestor belongs (see Section 3.2.19 on page 212).
 - 20 • The `omp_get_active_level` runtime library routine has been added, which
21 returns the number of nested, active `parallel` regions enclosing the task that
22 contains the call (see Section 3.2.20 on page 214).
 - 23 • In Version 3.0, locks are owned by tasks, not by threads (see Section 3.3 on page
24 224).

Index

Symbols

`_OPENMP` macro, 2-32

A

array sections, 2-42

`atomic`, 2-127

`atomic` construct, 8-300

attributes, data-sharing, 2-146

`auto`, 2-57

B

`barrier`, 2-123

C

`cancel`, 2-140

cancellation constructs

`cancel`, 2-140

`cancellation point`, 2-143

`cancellation point`, 2-143

`capture`, `atomic`, 2-127

clauses

`collapse`, 2-55

`copyin`, 2-173

`copyprivate`, 2-175

 data-sharing, 2-155

`default`, 2-156

`depend`, 2-116

`firstprivate`, 2-162

`lastprivate`, 2-164

`map`, 2-177

`private`, 2-159

`reduction`, 2-167

`schedule`, 2-56

`shared`, 2-157

`collapse`, 2-55

compliance, 1-21

conditional compilation, 2-32

constructs

`atomic`, 2-127

`barrier`, 2-123

`cancel`, 2-140

`cancellation point`, 2-143

`critical`, 2-122

`declare simd`, 2-72

`declare target`, 2-83

`distribute`, 2-88

`distribute parallel do`, 2-92

`distribute parallel do simd`, 2-94

`distribute parallel for`, 2-92

`distribute parallel for simd`, 2-94

 distribute parallel loop, 2-92

`distribute simd`, 2-91

`do`, *Fortran*, 2-54

`flush`, 2-134

`for`, *C/C++*, 2-54

loop, 2-53

 Loop SIMD, 2-76

`master`, 2-120

`ordered`, 2-138

`parallel`, 2-44

`parallel for`, *C/C++*, 2-95

`parallel sections`, 2-97

`parallel workshare`, *Fortran*, 2-99

`sections`, 2-60

`simd`, 2-68

- single**, 2-63
- target**, 2-79
- target data**, 2-77
- target teams**, 2-101
- target teams distribute**, 2-102, 2-105
- target update**, 2-81
- task**, 2-113
- taskgroup**, 2-126
- taskwait**, 2-125
- taskyield**, 2-117
- teams**, 2-86
- teams distribute**, 2-102
- workshare**, 2-65
 - worksharing*, 2-53

copyin, 2-173

copyprivate, 2-175

critical, 2-122

D

data sharing, 2-146

data-sharing clauses, 2-155

declare reduction, 2-180

declare simd construct, 2-72

declare target, 2-83

default, 2-156

depend, 2-116

device constructs

- declare target**, 2-83
- distribute**, 2-88
- target**, 2-79
- target data**, 2-77
- target update**, 2-81
- teams**, 2-86

device data environments, 1-18

directives, 2-25

- format, 2-26
- threadprivate**, 2-150
- see also constructs

distribute, 2-88

distribute parallel do, 2-92

distribute parallel do simd, 2-94

distribute parallel for, 2-92

distribute parallel for simd, 2-94

distribute simd, 2-91

do simd, 2-76

do, *Fortran*, 2-54

dynamic, 2-57

dynamic thread adjustment, 8-299

E

environment variables, 4-237

- modifying ICV's, 2-36
- OMP_CANCELLATION**, 4-246
- OMP_DEFAULT_DEVICE**, 4-248
- OMP_DISPLAY_ENV**, 4-247
- OMP_DYNAMIC**, 4-240
- OMP_MAX_ACTIVE_LEVELS**, 4-245
- OMP_NESTED**, 4-243
- OMP_NUM_THREADS**, 4-239
- OMP_SCHEDULE**, 4-238
- OMP_STACKSIZE**, 4-244
- OMP_THREAD_LIMIT**, 4-246
- OMP_WAIT_POLICY**, 4-245

execution model, 1-14

F

firstprivate, 2-162

flush, 2-134

flush operation, 1-19

for simd, 2-76

for, *C/C++*, 2-54

G

glossary, 1-2

grammar rules, 6-266

guided, 2-57

H

header files, 3-188, 7-287

I

ICVs (internal control variables), 2-34

implementation, 8-299

include files, 3-188, 7-287

internal control variables, 8-299

internal control variables (ICVs), 2-34

L

lastprivate, 2-164

loop directive, 8-300

loop SIMD construct, 2-76

loop, scheduling, 2-59

M

map, 2-177

master, 2-120

memory model, 1-17, 8-299

model

 execution, 1-14

 memory, 1-17

N

nested parallelism, 1-15, 2-34, 3-200

nesting, 2-186

number of threads, 2-47

O

OMP_CANCELLATION, 4-246

OMP_DEFAULT_DEVICE, 4-248

omp_destroy_lock, 3-227

omp_destroy_nest_lock, 3-227

OMP_DISPLAY_ENV, 4-247

OMP_DYNAMIC, 4-240, 8-301

omp_get_active_level, 3-214

omp_get_ancestor_thread_num, 3-211

omp_get_cancellation, 3-199

omp_get_default_device, 3-219

omp_get_dynamic, 3-198

omp_get_level, 3-210

omp_get_max_active_levels, 3-209, 8-301

omp_get_max_threads, 3-192

omp_get_nested, 3-201

omp_get_num_devices, 3-220

omp_get_num_procs, 3-195

omp_get_num_teams, 3-221

omp_get_num_threads, 3-191

omp_get_proc_bind, 3-216

omp_get_schedule, 3-205

omp_get_team_num, 3-222

omp_get_team_size, 3-212

omp_get_thread_limit, 3-206

omp_get_thread_num, 3-193

omp_get_wtick, 3-234

omp_get_wtime, 3-233

omp_in_final, 3-215

omp_in_parallel, 3-196

omp_init_lock, 3-226

omp_init_nest_lock, 3-226

omp_is_initial_device, 3-223

omp_lock_kind, 3-225

omp_lock_t, 3-225

OMP_MAX_ACTIVE_LEVELS, 4-245, 8-301

omp_nest_lock_kind, 3-225

omp_nest_lock_t, 3-225

OMP_NESTED, 4-243, 8-301

OMP_NUM_THREADS, 4-239, 8-301

OMP_PLACES, 4-241, 8-301

OMP_PROC_BIND, 8-301

OMP_SCHEDULE, 4-238, 8-301

omp_set_default_device, 3-218

omp_set_dynamic, 3-197

omp_set_lock, 3-228

omp_set_max_active_levels, 3-207, 8-301

omp_set_nest_lock, 3-228

omp_set_nested, 3-200

omp_set_num_threads, 3-189, 8-300

omp_set_schedule, 3-203, 8-300

OMP_STACKSIZE, 4-244, 8-301

omp_test_lock, 3-231

omp_test_nest_lock, 3-231

OMP_THREAD_LIMIT, 4-246, 8-301

omp_unset_lock, 3-229

omp_unset_nest_lock, 3-229

OMP_WAIT_POLICY, 4-245, 8-301

OpenMP

 compliance, 1-21

 features history, 9-303

 implementation, 8-299

ordered, 2-138

P

parallel, 2-44

parallel do, 2-96

parallel do simd, 2-100

parallel for simd, 2-100

parallel for, C/C++, 2-95

parallel loop SIMD construct, 2-100

parallel sections, 2-97

parallel workshare, *Fortran*, 2-99

pragmas
 see constructs
private, 2-159

R

read, atomic, 2-127
reduction, 2-167
references, 1-22
regions, nesting, 2-186
runtime, 2-58
runtime library
 interfaces and prototypes, 3-188
runtime library definitions, 8-302

S

schedule, 2-56
scheduling
 loop, 2-59
 tasks, 2-118
sections, 2-60
sections construct, 8-300
shared, 2-157
shared clause, 8-302
simd, 2-68
simd construct, 2-68
SIMD lanes, 1-15
SIMD loop, 2-68
SIMD loop construct, 2-76
SIMD parallel loop construct, 2-100
single, 2-63
single construct, 8-300
static, 2-57
stubs for runtime library routines
 C/C++, 5-250
 Fortran, 5-257
synchronization, locks
 constructs, 2-120
 routines, 3-224

T

target, 2-79
target data, 2-77
target teams, 2-101
target teams distribute, 2-102, 2-105

target update, 2-81
task
 scheduling, 2-118
task, 2-113
taskgroup, 2-126
tasking, 2-113
taskwait, 2-125
taskyield, 2-117
teams, 2-86
teams distribute, 2-102
terminology, 1-2
thread affinity policy, 8-302
threadprivate, 2-150, 8-302
timer, 3-233
timing routines, 3-233

U

update, atomic, 2-127

V

variables, environment, 4-237

W

wall clock timer, 3-233
website
 www.openmp.org
workshare, 2-65
worksharing
 constructs, 2-53
 parallel, 2-95
 scheduling, 2-59
write, atomic, 2-127