

# OpenMP Fortran Application Program Interface

Version 1.1–November–1999



# Contents

---

	<i>Page</i>
<b>Introduction [1]</b>	<b>1</b>
Scope . . . . .	1
Execution Model . . . . .	2
Compliance . . . . .	2
Organization . . . . .	3
<b>Directives [2]</b>	<b>5</b>
OpenMP directive format . . . . .	5
Directive sentinels . . . . .	6
Fixed source form directive sentinels . . . . .	6
Free source form directive sentinel . . . . .	7
Conditional compilation . . . . .	7
Fixed source form conditional compilation sentinels . . . . .	8
Free source form conditional compilation sentinel . . . . .	8
Parallel region construct . . . . .	9
Work-sharing constructs . . . . .	11
DO directive . . . . .	11
SECTIONS directive . . . . .	14
SINGLE directive . . . . .	16
Combined parallel work-sharing constructs . . . . .	16
PARALLEL DO directive . . . . .	17
PARALLEL SECTIONS directive . . . . .	17
Synchronization constructs and the MASTER directive . . . . .	18
MASTER directive . . . . .	19
CRITICAL directive . . . . .	19
<b>Version</b>	<b>i</b>
<b>1.1–November–1999</b>	

	<i>Page</i>
BARRIER directive . . . . .	20
ATOMIC directive . . . . .	20
FLUSH directive . . . . .	21
ORDERED directive . . . . .	23
<b>Data environment constructs . . . . .</b>	<b>24</b>
THREADPRIVATE directive . . . . .	24
<b>Data scope attribute clauses . . . . .</b>	<b>25</b>
PRIVATE clause . . . . .	26
SHARED clause . . . . .	27
DEFAULT clause . . . . .	27
FIRSTPRIVATE . . . . .	28
LASTPRIVATE clause . . . . .	28
REDUCTION clause . . . . .	29
COPYIN clause . . . . .	31
Data environment rules . . . . .	32
Directive binding . . . . .	34
Directive nesting . . . . .	35
 <b>Run-time Library Routines [3]</b>	 <b>37</b>
<b>Execution Environment Routines . . . . .</b>	<b>37</b>
OMP_SET_NUM_THREADS Subroutine . . . . .	37
OMP_GET_NUM_THREADS Function . . . . .	38
OMP_GET_MAX_THREADS Function . . . . .	38
OMP_GET_THREAD_NUM Function . . . . .	39
OMP_GET_NUM_PROCS Function . . . . .	39
OMP_IN_PARALLEL Function . . . . .	40
OMP_SET_DYNAMIC Subroutine . . . . .	40
OMP_GET_DYNAMIC Function . . . . .	41
OMP_SET_NESTED Subroutine . . . . .	41
OMP_GET_NESTED Function . . . . .	42
<b>Lock Routines . . . . .</b>	<b>42</b>

	<i>Page</i>
OMP_INIT_LOCK Subroutine . . . . .	42
OMP_DESTROY_LOCK Subroutine . . . . .	43
OMP_SET_LOCK Subroutine . . . . .	43
OMP_UNSET_LOCK Subroutine . . . . .	43
OMP_TEST_LOCK Function . . . . .	44
<b>Environment Variables [4]</b>	<b>45</b>
OMP_SCHEDULE Environment Variable . . . . .	45
OMP_NUM_THREADS Environment Variable . . . . .	45
OMP_DYNAMIC Environment Variable . . . . .	46
OMP_NESTED Environment Variable . . . . .	46
<b>Appendix A Examples</b>	<b>47</b>
Executing a Simple Loop in Parallel . . . . .	47
Specifying Conditional Compilation . . . . .	47
Using Parallel Regions . . . . .	48
Using the NOWAIT Clause . . . . .	48
Using the CRITICAL Directive . . . . .	48
Using the LASTPRIVATE Clause . . . . .	49
Using the REDUCTION Clause . . . . .	49
Specifying Parallel Sections . . . . .	50
Using SINGLE Directives . . . . .	50
Specifying Sequential Ordering . . . . .	50
Specifying a Fixed Number of Threads . . . . .	51
Using the ATOMIC Directive . . . . .	51
Using the FLUSH Directive . . . . .	52
Determining the Number of Threads Used . . . . .	52
Using Locks . . . . .	53
Nested DO Directives . . . . .	54
<b>Version</b>	<b>iii</b>
<b>1.1–November–1999</b>	

	<i>Page</i>
Examples Showing Incorrect Nesting of Work-sharing Directives . . . . .	55
Binding of BARRIER Directives . . . . .	57
Scoping Variables with the PRIVATE Clause . . . . .	58
Examples of Invalid Storage Association . . . . .	59
Examples of Syntax of Parallel DO Loops . . . . .	60
Examples of the ATOMIC Directive . . . . .	61
Examples of the ORDERED Directive . . . . .	63
Examples of THREADPRIVATE COMMON Data . . . . .	64
Examples of the Data Attribute Clauses: SHARED and PRIVATE . . . . .	66
 <b>Appendix B Stubs for Run-time Library Routines</b>	 <b>69</b>
 <b>Appendix C Using the SCHEDULE Clause</b>	 <b>73</b>
 <b>Tables</b>	
Table 1. Initialization Values . . . . .	30

Copyright © 1997-99 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.





# Introduction [1]

---

This document specifies a collection of compiler directives, library routines, and environment variables that can be used to specify shared memory parallelism in Fortran programs. The functionality described in this document is collectively known as the *OpenMP Fortran Application Program Interface (API)*. The goal of this specification is to provide a model for parallel programming that is portable across shared memory architectures from different vendors. The OpenMP Fortran API will be supported by compilers from numerous vendors. More information about OpenMP can be found at the following web site:

<http://www.openmp.org>

The directives, library routines, and environment variables defined in this document will allow users to create and manage parallel programs while ensuring portability. The directives extend the Fortran sequential programming model with single-program multiple data (SPMD) constructs, work-sharing constructs, synchronization constructs, and provide support for the sharing and privatization of data. The library routines and environment variables provide the functionality to control the run-time execution environment. The directive sentinels are structured so that the directives are treated as Fortran comments. Compilers that support the OpenMP Fortran API will include a command line option that activates and allows interpretation of all OpenMP compiler directives.

## 1.1 Scope

This specification describes only user-directed parallelization, wherein the user explicitly specifies the actions to be taken by the compiler and run-time system in order to execute the program in parallel. OpenMP Fortran implementations are not required to check for dependencies, conflicts, deadlocks, race conditions or other problems that result in incorrect program execution. The user is responsible for ensuring that the application using the OpenMP Fortran API constructs execute correctly.

Compiler-generated automatic parallelization and directives to the compiler to assist such parallelization are not included in this specification.

## 1.2 Execution Model

The OpenMP Fortran API uses the fork-join model of parallel execution. A program that is written with the OpenMP Fortran API begins execution as a single process, called the *master thread* of execution. The master thread executes sequentially until the first parallel construct is encountered. In the OpenMP Fortran API, the `PARALLEL` and `END PARALLEL` directive pair constitutes the parallel construct. When a parallel construct is encountered, the master thread creates a *team* of threads, and the master thread becomes the master of the team. The statements in the program that are enclosed by the parallel construct, including routines called from within the enclosed statements, are executed in parallel by each thread in the team. The statements enclosed lexically within a construct define the *static* extent of the construct. The *dynamic* extent further includes the routines called from within the construct.

Upon completion of the parallel construct, the threads in the team synchronize and only the master thread continues execution. Any number of parallel constructs can be specified in a single program. As a result, a program may fork and join many times during execution.

The OpenMP Fortran API allows programmers to use directives in routines called from within parallel constructs. Directives that do not appear in the lexical extent of the parallel construct but lie in the dynamic extent are called *orphaned* directives. Orphaned directives allow users to execute major portions of their program in parallel with only minimal changes to the sequential program. With this functionality, users can code parallel constructs at the top levels of the program call tree and use directives to control execution in any of the called routines.

## 1.3 Compliance

An implementation of the OpenMP Fortran API is *OpenMP compliant* if it recognizes and preserves the semantics of all the elements of this specification as laid out in chapters 2, 3, and 4. The appendixes are for information purposes only and are not part of the specification.

The OpenMP Fortran API is an extension to the base language that is supported by an implementation. If the base language does not support a language construct or extension that appears in this document, the OpenMP implementation is not required to support it.

All standard Fortran intrinsics and library routines and Fortran 90 `ALLOCATE` and `DEALLOCATE` statements must be thread-safe. Unsynchronized use of such intrinsics and routines by different threads in a parallel region must produce

correct results (though not necessarily the same as serial execution results, as in the case of random number generation intrinsics, for example).

Unsynchronized use of Fortran output statements to the same unit may result in output in which data written by different threads is interleaved. Similarly, unsynchronized input statements from the same unit may read data in an interleaved fashion. Unsynchronized use of Fortran I/O, such that each thread accesses a different unit, produces the same results as serial execution of the I/O statements.

## 1.4 Organization

The rest of this document is organized into the following chapters:

- Chapter 2, page 5, describes the compiler directives.
- Chapter 3, page 37, describes the run-time library routines.
- Chapter 4, page 45, describes the environment variables.
- Appendix A, page 47, contains examples.
- Appendix B, page 69, describes stub library routines.
- Appendix C, page 73, has information about using the `SCHEDULE` clause.



# Directives [2]

---

Directives are special Fortran comments that are identified with a unique *sentinel*. The directive sentinels are structured so that the directives are treated as Fortran comments. Compilers that support the OpenMP Fortran API will include a command line option that activates and allows interpretation of all OpenMP compiler directives. In the remainder of this document, the phrase *OpenMP compilation* is used to mean that OpenMP directives are interpreted during compilation.

This chapter addresses the following topics:

- Section 2.1, page 5, describes the directive format.
- Section 2.1.1, page 6, describes directive sentinels for both fixed source form and free source form.
- Section 2.1.2, page 7, describes conditional compilation.
- Section 2.2, page 9, describes the parallel region construct.
- Section 2.3, page 11, describes work-sharing constructs.
- Section 2.4, page 16, describes the combined parallel work-sharing constructs.
- Section 2.5, page 18, describes synchronization constructs and the `MASTER` directive.
- Section 2.6, page 24, describes the data environment, which includes directives and clauses that affect the data environment.
- Section 2.7, page 34, describes directive binding.
- Section 2.8, page 35, describes directive nesting.

## 2.1 OpenMP directive format

The format of an OpenMP directive is as follows:

<code><i>sentinel</i> directive_name [clause[, ] clause] . . .</code>
-----------------------------------------------------------------------

All OpenMP compiler directives must begin with a directive *sentinel*. Directives are case insensitive. Clauses can appear in any order after the directive name.

Clauses on directives can be repeated as needed, subject to the restrictions listed in the description of each clause. Directives cannot be embedded within continued statements, and statements cannot be embedded within directives. Comments cannot appear on the same line as a directive.

The following sections contain more information on directive sentinels and describe conditional compilation.

### 2.1.1 Directive sentinels

The directive sentinels accepted by an OpenMP-compliant compiler differ depending on the Fortran source form being used. The `!$OMP` sentinel is accepted when compiling either fixed source form files or free source form files. The `C$OMP` and `*$OMP` sentinels are accepted only when compiling fixed source form files.

The following sections contain more information on using the different sentinels.

#### 2.1.1.1 Fixed source form directive sentinels

The OpenMP Fortran API accepts the following sentinels in fixed source form files:

<code>!\$OMP</code>   <code>C\$OMP</code>   <code>*\$OMP</code>
-----------------------------------------------------------------

Sentinels must start in column one and appear as a single word with no intervening white space. Fortran fixed form line length, case sensitivity, white space, continuation, and column rules apply to the directive line. Initial directive lines must have a space or zero in column six, and continuation directive lines must have a character other than a space or a zero in column six.

Example: The following formats for specifying directives are equivalent (the first line represents the position of the first 9 columns):

```
C23456789
!$OMP PARALLEL DO SHARED(A,B,C)

C$OMP PARALLEL DO
C$OMP+SHARED(A,B,C)

C$OMP PARALLELDOSHARED(A,B,C)
```

### 2.1.1.2 Free source form directive sentinel

The OpenMP Fortran API accepts the following sentinel in free source form files:

```
!$OMP
```

The sentinel can appear in any column as long as it is preceded only by white space. It must appear as a single word with no intervening white space. Fortran free form line length, case sensitivity, white space, and continuation rules apply to the directive line. Initial directive lines must have a space after the sentinel. Continued directive lines must have an ampersand as the last nonblank character on the line. Continuation directive lines can have an ampersand after the directive sentinel with optional white space before and after the ampersand.

Example: The following formats for specifying directives are equivalent (the first line represents the position of the first 9 columns):

```
!23456789
    !$OMP PARALLEL DO &
        !$OMP SHARED(A,B,C)

!$OMP PARALLEL &
    !$OMP&DO SHARED(A,B,C)

!$OMP PARALLEL DO SHARED(A,B,C)
```

In order to simplify the presentation, the remainder of this document uses the !\$OMP sentinel.

### 2.1.2 Conditional compilation

The OpenMP Fortran API permits Fortran lines to be compiled conditionally. The directive sentinels for conditional compilation that are accepted by an OpenMP-compliant compiler depend on the Fortran source form being used. The !\$ sentinel is accepted when compiling either fixed source form files or free source form files. The C\$ and \*\$ sentinels are accepted only when compiling fixed source form.

During OpenMP compilation, the sentinel is replaced by two spaces, and the rest of the line is treated as a normal Fortran line.

In addition to the Fortran conditional compilation sentinels, a C preprocessor macro, `_OPENMP`, can be used for conditional compilation. OpenMP-compliant compilers will define this macro during OpenMP compilation.

The following sections contain more information on using the different sentinels for conditional compilation. (See Section A.2, page 47 for an example.)

#### 2.1.2.1 Fixed source form conditional compilation sentinels

The OpenMP Fortran API accepts the following conditional compilation sentinels in fixed source form files:

!\$   C\$   *\$
-----------------

The sentinels must start in column one and appear as a single word with no intervening white space. Fortran fixed form line length, case sensitivity, white space, continuation, and column rules apply to the line. Initial lines must have a space or zero in column six, and continuation lines must have a character other than a space or zero in column six.

Example: The following forms for specifying conditional compilation are equivalent:

```
C23456789
!$ 10 IAM = OMP_GET_THREAD_NUM() +
!$   &           INDEX

#ifdef _OPENMP
    10 IAM = OMP_GET_THREAD_NUM() +
        &           INDEX
#endif
```

#### 2.1.2.2 Free source form conditional compilation sentinel

The OpenMP Fortran API accepts the following conditional compilation sentinel in free source form files:

!\$
-----

This sentinel can appear in any column as long as it is preceded only by white space. It must appear as a single word with no intervening white space. Fortran free source form line length, case sensitivity, white space, and



continuation rules apply to the line. Initial lines must have a space after the sentinel. Continued lines must have an ampersand as the last nonblank character on the line. Continuation lines can have an ampersand after the sentinel, with optional white space before and after the ampersand.

## 2.2 Parallel region construct

The `PARALLEL` and `END PARALLEL` directives define a *parallel region*. A parallel region is a block of code that is to be executed by multiple threads in parallel. This is the fundamental parallel construct in OpenMP that starts parallel execution. These directives have the following format:

```
!$OMP PARALLEL [clause [, clause] . . .]  
  
block  
  
!$OMP END PARALLEL
```

The *clause* can be one of the following:

- `PRIVATE(list)`
- `SHARED(list)`
- `DEFAULT(PRIVATE | SHARED | NONE)`
- `FIRSTPRIVATE(list)`
- `REDUCTION ( {operator | intrinsic } : list )`
- `IF(scalar_logical_expression)`
- `COPYIN(list)`

For information on the `PRIVATE`, `SHARED`, `DEFAULT`, `FIRSTPRIVATE`, `REDUCTION`, and `COPYIN` clauses, see Section 2.6.2, page 25. For an example of how to implement coarse-grain parallelism using these directives, see Section A.3, page 48.

When a thread encounters a parallel region, it creates a team of threads, and it becomes the master of the team. The master thread is a member of the team and it has a thread number of 0 within the team. The number of threads in the team is controlled by environment variables and/or library calls. For more

information on environment variables, see Chapter 4. For more information on library routines, see Chapter 3, page 37.

The number of physical processors actually hosting the threads at any given time is implementation dependent. Once created, the number of threads in the team remains constant for the duration of that parallel region, but it can be changed either explicitly by the user or automatically by the run-time system from one parallel region to another. The `OMP_SET_DYNAMIC` library routine and the `OMP_DYNAMIC` environment variable can be used to enable and disable the automatic adjustment of the number of threads. For more information on the `OMP_SET_DYNAMIC` library routine, see Section 3.1.7, page 40. For more information on the `OMP_DYNAMIC` environment variable, see Section 4.3, page 46.

The *block* denotes a structured block of Fortran statements. It is illegal to branch into or out of the block. The code contained within the dynamic extent of the parallel region is executed on each thread, and the code path can be different for different threads.

The `END PARALLEL` directive denotes the end of the parallel region. There is an implied barrier at this point. Only the master thread of the team continues execution at the end of a parallel region.

If a thread in a team executing a parallel region encounters another parallel region, it creates a new team, and it becomes the master of that new team. By default, nested parallel regions are serialized; that is, they are executed by a team composed of one thread. This default behavior can be changed by using either the `OMP_SET_NESTED` run-time library routine or the `OMP_NESTED` environment variable. For more information on the `OMP_SET_NESTED` library routine, see Section 3.1.9, page 41. For more information on the `OMP_NESTED` environment variable, see Section 4.4, page 46.

If an `IF` clause is present, the enclosed code region is executed in parallel only if the *scalar\_logical\_expression* evaluates to `.TRUE.`. Otherwise, the parallel region is serialized. The expression must be a scalar Fortran logical expression. In the absence of an `IF` clause, the region is executed as if an `IF(.TRUE.)` clause were specified.

The following restrictions apply to parallel regions:

- The `PARALLEL/END PARALLEL` directive pair must appear in the same routine in the executable section of the code.
- The code contained by these two directives must be a structured block. It is illegal to branch into or out of a parallel region.

- Only a single `IF` clause can appear on the directive. The `IF` expression is evaluated outside the context of the parallel region. Results are undefined if the `IF` expression contains a function reference that has side effects.

Unsynchronized use of Fortran I/O statements by multiple threads on the same unit has undefined behavior. (See Section A.3, page 48 for a related example.)

## 2.3 Work-sharing constructs

A work-sharing construct divides the execution of the enclosed code region among the members of the team that encounter it. A work-sharing construct must be enclosed dynamically within a parallel region in order for the directive to execute in parallel. The work-sharing directives do not launch new threads, and there is no implied barrier on entry to a work-sharing construct.

The following restrictions apply to the work-sharing directives:

- Work-sharing constructs and `BARRIER` directives must be encountered by all threads in a team or by none at all.
- Work-sharing constructs and `BARRIER` directives must be encountered in the same order by all threads in a team.

The following sections describe the work-sharing directives:

- Section 2.3.1, page 11, describes the `DO` and `END DO` directives.
- Section 2.3.2, page 14, describes the `SECTIONS`, `SECTION`, and `END SECTIONS` directives.
- Section 2.3.3, page 16, describes the `SINGLE` and `END SINGLE` directives.

### 2.3.1 DO directive

The `DO` directive specifies that the iterations of the immediately following `DO` loop must be executed in parallel. The loop that follows a `DO` directive cannot be a `DO WHILE` or a `DO` loop without loop control. The iterations of the `DO` loop are distributed across threads that already exist.

The format of this directive is as follows:

```
!$OMP DO [clause [, ] clause . . . ]  
  
do_loop  
  
[!$OMP END DO [NOWAIT]]
```

The *do\_loop* may be a *do\_construct*, an *outer\_shared\_do\_construct*, or an *inner\_shared\_do\_construct*. A DO construct that contains several DO statements that share the same DO termination statement syntactically consists of a sequence of *outer\_shared\_do\_constructs*, followed by a single *inner\_shared\_do\_construct*. If an END DO directive follows such a DO construct, a DO directive can only be specified for the first (i.e., the outermost) *outer\_shared\_do\_construct*. (See examples in Section A.21, page 60.)

The *clause* can be one of the following:

- PRIVATE(*list*)
- FIRSTPRIVATE(*list*)
- LASTPRIVATE(*list*)
- REDUCTION( {*operator* | *intrinsic* } : *list*)
- SCHEDULE(*type* [ , *chunk*])
- ORDERED

The SCHEDULE and ORDERED clauses are described in this section. The PRIVATE, FIRSTPRIVATE, LASTPRIVATE, and REDUCTION clauses are described in Section 2.6.2, page 25.

If ordered sections are contained in the dynamic extent of the DO directive, the ORDERED clause must be present. For more information on ordered sections, see the ORDERED directive in Section 2.5.6, page 23.

The SCHEDULE clause specifies how iterations of the DO loop are divided among the threads of the team. The *chunk* expression is evaluated outside the context of the DO construct. Results are undefined if the *chunk* expression contains a function reference that has side effects. Within the SCHEDULE(*type* [ , *chunk*]) clause syntax, *type* can be one of the following:

<u><i>type</i></u>	<u>Effect</u>
STATIC	When SCHEDULE(STATIC, <i>chunk</i> ) is specified, iterations are divided into pieces of a size specified by <i>chunk</i> . The pieces are

statically assigned to threads in the team in a round-robin fashion in the order of the thread number. *chunk* must be a scalar integer expression.

When no *chunk* is specified, the iterations are divided among threads in contiguous pieces, and one piece is assigned to each thread.

DYNAMIC When `SCHEDULE(DYNAMIC, chunk)` is specified, the iterations are broken into pieces of a size specified by *chunk*. As each thread finishes a piece of the iteration space, it dynamically obtains the next set of iterations.

When no *chunk* is specified, it defaults to 1.

GUIDED When `SCHEDULE(GUIDED, chunk)` is specified, the iteration space is divided into pieces such that the size of each successive piece is exponentially decreasing. *chunk* specifies the size of the smallest piece, except possibly the last. The size of the initial piece is implementation dependent. As each thread finishes a piece of the iteration space, it dynamically obtains the next available piece.

When no *chunk* is specified, it defaults to 1.

RUNTIME When `SCHEDULE(RUNTIME)` is specified, the decision regarding scheduling is deferred until run time. The schedule type and chunk size can be chosen at run time by setting the `OMP_SCHEDULE` environment variable. If this environment variable is not set, the resulting schedule is implementation-dependent. For more information on the `OMP_SCHEDULE` environment variable, see Section 4.1, page 45.

When `SCHEDULE(RUNTIME)` is specified, it is illegal to specify a *chunk*.

In the absence of the `SCHEDULE` clause, the default schedule is implementation dependent. An OpenMP-compliant program should not rely on a particular schedule for correct execution. Users should not rely on a particular implementation of a schedule type for correct execution, because it is possible to have variations in the implementations of the same schedule type across different compilers.

If an `END DO` directive is not specified, an `END DO` directive is assumed at the end of the `DO` loop. If `NOWAIT` is specified on the `END DO` directive, threads do not synchronize at the end of the parallel loop. Threads that finish early proceed straight to the instructions following the loop without waiting for the

other members of the team to finish the `DO` directive. (See Section A.4, page 48 for an example.)

Parallel `DO` loop control variables are block-level entities within the `DO` loop. If the loop control variable also appears in the `LASTPRIVATE` list of the parallel `DO`, it is copied out to a variable of the same name in the enclosing `PARALLEL` region. The variable in the enclosing `PARALLEL` region must be `SHARED` if it is specified on the `LASTPRIVATE` list of a `DO` directive.

The following restrictions apply to the `DO` directives:

- It is illegal to branch out of a `DO` loop associated with a `DO` directive.
- The values of the loop control parameters of the `DO` loop associated with a `DO` directive must be the same for all the threads in the team.
- The `DO` loop iteration variable must be of type integer.
- If used, the `END DO` directive must appear immediately after the end of the loop.
- Only a single `SCHEDULE` clause can appear on a `DO` directive.
- Only a single `ORDERED` clause can appear on a `DO` directive.
- The value of the *chunk* parameter must be the same for all of the threads in the team.

### 2.3.2 `SECTIONS` directive

The `SECTIONS` directive is a non-iterative work-sharing construct that specifies that the enclosed sections of code are to be divided among threads in the team. Each section is executed once by a thread in the team.

The format of this directive is as follows:

```

!$OMP SECTIONS [clause[,] clause... ]

[!$OMP SECTION]

block

[!$OMP SECTION

block]

. . .

!$OMP END SECTIONS [NOWAIT]

```

The *clause* can be one of the following:

- PRIVATE(*list*)
- FIRSTPRIVATE(*list*)
- LASTPRIVATE(*list*)
- REDUCTION( { *operator* | *intrinsic* } : *list* )

The PRIVATE, FIRSTPRIVATE, LASTPRIVATE, and REDUCTION clauses are described in Section 2.6.2, page 25.

Each section is preceded by a SECTION directive, though the SECTION directive is optional for the first section. The SECTION directives must appear within the lexical extent of the SECTIONS/END SECTIONS directive pair. The last section ends at the END SECTIONS directive. Threads that complete execution of their sections wait at a barrier at the END SECTIONS directive unless a NOWAIT is specified.

The following restrictions apply to the SECTIONS directive:

- The code enclosed in a SECTIONS/END SECTIONS directive pair must be a structured block. In addition, each constituent section must also be a structured block. It is illegal to branch into or out of the constituent section blocks.
- It is illegal for a SECTION directive to be outside the lexical extent of the SECTIONS/END SECTIONS directive pair. (See Section A.8, page 50 for an example that uses these directives.)

### 2.3.3 SINGLE directive

The `SINGLE` directive specifies that the enclosed code is to be executed by only one thread in the team. Threads in the team that are not executing the `SINGLE` directive wait at the `END SINGLE` directive unless `NOWAIT` is specified.

The format of this directive is as follows:

```
!$OMP SINGLE [clause[[,] clause] . . . ]  
  
block  
  
!$OMP END SINGLE [NOWAIT]
```

The *clause* can be one of the following:

- `PRIVATE( list )`
- `FIRSTPRIVATE( list )`

The `PRIVATE` and `FIRSTPRIVATE` clauses are described in Section 2.6.2, page 25.

The following restriction applies to a `SINGLE` directive:

- The code enclosed in a `SINGLE/END SINGLE` directive pair must be a structured block. It is illegal to branch into or out of the block.

See Section A.9, page 50 for an example of the `SINGLE` directive.

## 2.4 Combined parallel work-sharing constructs

The combined parallel work-sharing constructs are shortcuts for specifying a parallel region that contains only one work-sharing construct. The semantics of these directives are identical to that of explicitly specifying a `PARALLEL` directive followed by a single work-sharing construct.

The following sections describe the combined parallel work-sharing directives:

- Section 2.4.1, page 17, describes the `PARALLEL DO` and `END PARALLEL DO` directives.
- Section 2.4.2, page 17, describes the `PARALLEL SECTIONS` and `END PARALLEL SECTIONS` directives.



### 2.4.1 PARALLEL DO directive

The PARALLEL DO directive provides a shortcut form for specifying a parallel region that contains a single DO directive. (See also Section A.1, page 47 for an example.)

The format of this directive is as follows:

```
!$OMP PARALLEL DO [clause[[,] clause] . . .]  
  
do_loop  
  
[!$OMP END PARALLEL DO]
```

The *do\_loop* may be a *do\_construct*, an *outer\_shared\_do\_construct*, or an *inner\_shared\_do\_construct*. A DO construct that contains several DO statements that share the same DO termination statement syntactically consists of a sequence of *outer\_shared\_do\_constructs*, followed by a single *inner\_shared\_do\_construct*. If an END PARALLEL DO directive follows such a DO construct, a PARALLEL DO directive can only be specified for the first (i.e., the outermost) *outer\_shared\_do\_construct*. (See Section A.21, page 60 for examples.)

The *clause* can be one of the *clauses* accepted by the PARALLEL and DO directives. For information on the PARALLEL directive and the IF clause, see Section 2.2, page 9. For information on the DO directive and the SCHEDULED and ORDERED clauses, see Section 2.3.1, page 11. For information on the remaining clauses, see Section 2.6.2, page 25.

If the END PARALLEL DO directive is not specified, the PARALLEL DO is assumed to end with the DO loop that immediately follows the PARALLEL DO directive. If used, the END PARALLEL DO directive must appear immediately after the end of the DO loop.

The semantics are identical to explicitly specifying a PARALLEL directive immediately followed by a DO directive.

### 2.4.2 PARALLEL SECTIONS directive

The PARALLEL SECTIONS directive provides a shortcut form for specifying a parallel region that contains a single SECTIONS directive. The semantics are identical to explicitly specifying a PARALLEL directive immediately followed by a SECTIONS directive.

The format of this directive is as follows:

```
!$OMP PARALLEL SECTIONS [clause[,] clause. . . ]  
  
[!$OMP SECTION ]  
  
block  
  
[!$OMP SECTION  
  
block]  
  
. . .  
  
!$OMP END PARALLEL SECTIONS
```

The *clause* can be one of the *clauses* accepted by the PARALLEL and SECTIONS directives. For more information on the PARALLEL directive, see Section 2.2, page 9. For more information on the SECTIONS directive, see Section 2.3.2, page 14. The PRIVATE, FIRSTPRIVATE, LASTPRIVATE, and REDUCTION clauses are described in Section 2.6.2, page 25.

The last section ends at the END PARALLEL SECTIONS directive.

## 2.5 Synchronization constructs and the MASTER directive

The following sections describe the synchronization constructs and the MASTER directive:

- Section 2.5.1, page 19, describes the MASTER and END MASTER directives.
- Section 2.5.2, page 19, describes the CRITICAL and END CRITICAL directives.
- Section 2.5.3, page 20, describes the BARRIER directive.
- Section 2.5.4, page 20, describes the ATOMIC directive.
- Section 2.5.5, page 21, describes the FLUSH directive.
- Section 2.5.6, page 23, describes the ORDERED and END ORDERED directives.

### 2.5.1 MASTER directive

The code enclosed within `MASTER` and `END MASTER` directives is executed by the master thread of the team.

These directives have the following format:

```
!$OMP MASTER  
  
block  
  
!$OMP END MASTER
```

The other threads in the team skip the enclosed section of code and continue execution. There is no implied barrier either on entry to or exit from the master section.

This directive has the following restriction:

- The section of code enclosed by `MASTER` and `END MASTER` directives must be a structured block. It is illegal to branch into or out of the block.

### 2.5.2 CRITICAL directive

The `CRITICAL` and `END CRITICAL` directives restrict access to the enclosed code to only one thread at a time.

These directives have the following format:

```
!$OMP CRITICAL [(name)]  
  
block  
  
!$OMP END CRITICAL [(name)]
```

The optional *name* argument identifies the critical section.

A thread waits at the beginning of a critical section until no other thread in the team is executing a critical section with the same name. All unnamed `CRITICAL` directives map to the same name. Critical section names are global entities of the program. If a name conflicts with any other entity, the behavior of the program is undefined.

The following restrictions apply to the `CRITICAL` directive:

- The section of code enclosed by the `CRITICAL` and `END CRITICAL` directive pair must be a structured block. It is illegal to branch into or out of the block.
- If a *name* is specified on a `CRITICAL` directive, the same *name* must also be specified on the `END CRITICAL` directive. If no *name* appears on the `CRITICAL` directive, no *name* can appear on the `END CRITICAL` directive.

See Section A.5, page 48 for an example that uses named `CRITICAL` sections.

### 2.5.3 BARRIER directive

The `BARRIER` directive synchronizes all the threads in a team. When encountered, each thread waits until all of the others threads in that team have reached this point.

This directive has the following format:

```
!$OMP BARRIER
```

### 2.5.4 ATOMIC directive

The `ATOMIC` directive ensures that a specific memory location is to be updated atomically, rather than exposing it to the possibility of multiple, simultaneous writing threads.

This directive has the following format:

```
!$OMP ATOMIC
```

This directive applies only to the immediately following statement, which must have one of the following forms:

```
x = x operator expr  
x = expr operator x  
x = intrinsic (x, expr)  
x = intrinsic (expr, x)
```

In the preceding statements:

- *x* is a scalar variable of intrinsic type.
- *expr* is a scalar expression that does not reference *x*.
- *intrinsic* is one of MAX, MIN, IAND, IOR, or IEOR.
- *operator* is one of +, \*, -, /, .AND., .OR., .EQV., or .NEQV. .

This directive permits optimization beyond that of the necessary critical section around the assignment. An implementation can replace all ATOMIC directives by enclosing the statement in a critical section. All of these critical sections must use the same unique name.

Only the load and store of *x* are atomic; the evaluation of *expr* is not atomic. To avoid race conditions, all updates of the location in parallel must be protected with the ATOMIC directive, except those that are known to be free of race conditions. The function *intrinsic*, the operator *operator*, and the assignment must be the intrinsic function, operator, and assignment.

The following restriction applies to the use of ATOMIC directives:

- All atomic references to the storage location of variable *x* throughout the program are required to have the same type and type parameters.

Example:

```
!$OMP ATOMIC
      Y(INDEX(I)) = Y(INDEX(I)) + B
```

See Section A.12, page 51 and Section A.22, page 61 for more examples using the ATOMIC directive.

### 2.5.5 FLUSH directive

The FLUSH directive, whether explicit or implied, identifies a cross-thread sequence point at which the implementation is required to ensure that each thread in the team has a consistent view of certain variables in memory.

A consistent view requires that all memory operations (both reads and writes) that occur before the FLUSH directive in the program be performed before the sequence point in the executing thread; similarly, all memory operations that occur after the FLUSH must be performed after the sequence point in the executing thread.

Thread-visible variables include the following data items:

- Globally visible variables (in common blocks and in modules).
- Local variables that do not have the `SAVE` attribute but have had their address taken and saved or have had their address passed to another subprogram.
- Local variables that do not have the `SAVE` attribute that are declared shared in a parallel region within the subprogram.
- Dummy arguments.
- All pointer dereferences.

Implementations must ensure that modifications made to thread-visible variables within the executing thread are made visible to all other threads at the sequence point. For example, compilers must restore values from registers to memory, and hardware may need to flush write buffers. Furthermore, implementations must assume that thread-visible variables may have been updated by other threads at the sequence point and must be retrieved from memory before their first use past the sequence point.

Finally, the `FLUSH` directive only provides consistency between operations within the executing thread and global memory. To achieve a globally consistent view across all threads, each thread must execute a `FLUSH` operation.

This directive has the following format:

```
!$OMP FLUSH [( list )]
```

This directive must appear at the precise point in the code at which the synchronization is required. The optional *list* argument consists of a comma-separated list of variables that need to be flushed in order to avoid flushing all variables. The *list* should contain only named variables (see Section A.13, page 52). The `FLUSH` directive is implied for the following directives:

- `BARRIER`
- `CRITICAL` and `END CRITICAL`
- `END DO`
- `END PARALLEL`
- `END SECTIONS`

- END SINGLE
- ORDERED and END ORDERED

The directive is not implied if a NOWAIT clause is present.

### 2.5.6 ORDERED directive

The code enclosed within ORDERED and END ORDERED directives is executed in the order in which iterations would be executed in a sequential execution of the loop.

These directives have the following format:

```
!$OMP ORDERED  
  
block  
  
!$OMP END ORDERED
```

An ORDERED directive can appear only in the dynamic extent of a DO or PARALLEL DO directive. The DO directive to which the ordered section binds must have the ORDERED clause specified (see Section 2.3.1, page 11). One thread is allowed in an ordered section at a time. Threads are allowed to enter in the order of the loop iterations. No thread can enter an ordered section until it is guaranteed that all previous iterations have completed or will never execute an ordered section. This sequentializes and orders code within ordered sections while allowing code outside the section to run in parallel. ORDERED sections that bind to different DO directives are independent of each other.

The following restrictions apply to the ORDERED directive:

- The code enclosed by the ORDERED and END ORDERED directives must be a structured block. It is illegal to branch into or out of the block.
- An ORDERED directive cannot bind to a DO directive that does not have the ORDERED clause specified.
- An iteration of a loop with a DO directive must not execute the same ORDERED directive more than once, and it must not execute more than one ORDERED directive.

See Section A.10, page 50 and Section A.23, page 63 for examples using the ORDERED directive.

## 2.6 Data environment constructs

This section presents constructs for controlling the data environment during the execution of parallel constructs:

- Section 2.6.1, page 24, describes the `THREADPRIVATE` directive, which makes common blocks local to a thread.
- Section 2.6.2, page 25, describes directive clauses that affect the data environment.
- Section 2.6.3, page 32 describes the data environment rules.

### 2.6.1 `THREADPRIVATE` directive

The `THREADPRIVATE` directive makes named common blocks private to a thread but global within the thread.

This directive must appear in the declaration section of the scoping unit after the declaration of the listed common blocks. 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. Each thread gets its own copy of the common block, so data written to the common block by one thread is not directly visible to other threads. During serial portions and `MASTER` sections of the program, accesses are to the master thread's copy of the common block. (See Section A.24, page 64 for examples.)

On entry to the first parallel region, data in the `THREADPRIVATE` common blocks should be assumed to be undefined unless a `COPYIN` clause is specified on the `PARALLEL` directive. When a common block that is initialized using `DATA` statements appears in a `THREADPRIVATE` directive, each thread's copy is initialized once prior to its first use. For subsequent parallel regions, the data in the `THREADPRIVATE` common blocks is guaranteed to persist only if the dynamic threads mechanism has been disabled and if the number of threads is the same for all parallel regions. For more information on dynamic threads, see the `OMP_SET_DYNAMIC` library routine, Section 3.1.7, page 40, and the `OMP_DYNAMIC` environment variable, Section 4.3, page 46.

The format of this directive is as follows:

```
!$OMP THREADPRIVATE ( /cb/[ , /cb/ ] . . . )
```



*cb* is the name of the common block to be made private to a thread.

The following restrictions apply to the `THREADPRIVATE` directive:

- The `THREADPRIVATE` directive must appear after every declaration of a thread private common block.
- Only named common blocks can be made thread private.
- It is illegal for a `THREADPRIVATE` common block or its constituent variables to appear in any clause other than a `COPYIN` clause. As a result, they are not permitted in a `PRIVATE`, `FIRSTPRIVATE`, `LASTPRIVATE`, `SHARED`, or `REDUCTION` clause. They are not affected by the `DEFAULT` clause.

## 2.6.2 Data scope attribute clauses

Several directives accept clauses that allow a user to control the scope attributes of variables for the duration of the construct. Not all of the following clauses are allowed on all directives, but the clauses that are valid on a particular directive are included with the description of the directive. If no data scope clauses are specified for a directive, the default scope for variables affected by the directive is `SHARED`. (See Section 2.6.3, page 32, for exceptions.)

Scope attribute clauses that appear on a `PARALLEL` directive indicate how the specified variables are to be treated with respect to the parallel region associated with the `PARALLEL` directive. They do not indicate the scope attributes of these variables for any enclosing parallel regions, if they exist.

In determining the appropriate scope attribute for a variable used in the lexical extent of a parallel region, all references and definitions of the variable must be considered, including references and definitions which occur in any nested parallel regions.

Each clause accepts an argument *list*, which is a comma-separated list of named variables or named common blocks that are accessible in the scoping unit. Subobjects cannot be specified as items in any of the lists. When common blocks appear in a list, their names must appear between slashes.

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 data scope attribute clause must be declared to be a common block in the same scoping unit in which the data scope attribute clause appears.

The following sections describe the data scope attribute clauses:

- Section 2.6.2.1, page 26, describes the `PRIVATE` clause.

- Section 2.6.2.2, page 27, describes the `SHARED` clause.
- Section 2.6.2.3, page 27, describes the `DEFAULT` clause.
- Section 2.6.2.4, page 28, describes the `FIRSTPRIVATE` clause.
- Section 2.6.2.5, page 28, describes the `LASTPRIVATE` clause.
- Section 2.6.2.6, page 29, describes the `REDUCTION` clause.
- Section 2.6.2.7, page 31, describes the `COPYIN` clause.

### 2.6.2.1 `PRIVATE` clause

The `PRIVATE` clause declares the variables in *list* to be private to each thread in a team.

This clause has the following format:

<code>PRIVATE ( <i>list</i> )</code>
--------------------------------------

The behavior of a variable declared in a `PRIVATE` clause is as follows:

1. A new object of the same type is declared once for each thread in the team. One thread in the team is permitted, but not required, to re-use the existing storage as the storage for the new object. For all other threads, new storage is created for the new object.
2. All references to the original object in the lexical extent of the directive construct are replaced with references to the private object.
3. Variables defined as `PRIVATE` are undefined for each thread on entering the construct, and the corresponding shared variable is undefined on exit from a parallel construct.
4. A variable declared as `PRIVATE` may be storage-associated with other variables when the `PRIVATE` clause is encountered. Storage association may exist because of constructs such as `EQUIVALENCE`, `COMMON`, etc. If *a* is a variable appearing in a `PRIVATE` clause and *b* is a variable which was storage-associated with *a*, then:
  - a. The contents, allocation, and association status of *b* are undefined on entry to the parallel construct.
  - b. Any definition of *a*, or of its allocation or association status, causes the contents, allocation, and association status of *b* to become undefined.

- c. Any definition of `b`, or of its allocation or association status, causes the contents, allocation, and association status of `a` to become undefined.

See Section A.19, page 58 and Section A.20, page 59, for examples.

### 2.6.2.2 SHARED clause

The `SHARED` clause makes variables that appear in the *list* shared among all the threads in a team. All threads within a team access the same storage area for `SHARED` data.

This clause has the following format:

```
SHARED ( list )
```

### 2.6.2.3 DEFAULT clause

The `DEFAULT` clause allows the user to specify a `PRIVATE`, `SHARED`, or `NONE` scope attribute for all variables in the lexical extent of any parallel region. Variables in `THREADPRIVATE` common blocks are not affected by this clause.

This clause has the following format:

```
DEFAULT ( PRIVATE | SHARED | NONE )
```

The `PRIVATE`, `SHARED`, and `NONE` specifications have the following effects:

- Specifying `DEFAULT ( PRIVATE )` makes all named objects in the lexical extent of the parallel region, including common block variables but excluding `THREADPRIVATE` variables, private to a thread as if each variable were listed explicitly in a `PRIVATE` clause.
- Specifying `DEFAULT ( SHARED )` makes all named objects in the lexical extent of the parallel region shared among the threads in a team, as if each variable were listed explicitly in a `SHARED` clause. In the absence of an explicit `DEFAULT` clause, the default behavior is the same as if `DEFAULT ( SHARED )` were specified.
- Specifying `DEFAULT ( NONE )` requires that each variable used in the lexical extent of the parallel region be explicitly listed in a data scope attribute clause on the parallel region, unless it is one of the following:
  - `THREADPRIVATE`.

- A Cray pointee.
- A loop iteration variable used only as a loop iteration variable for sequential loops in the lexical extent of the region or parallel DO loops that bind to the region.
- Only used in work-sharing constructs that bind to the region, and is specified in a data scope attribute clause for each such construct.

Only one DEFAULT clause can be specified on a PARALLEL directive.

Variables can be exempted from a defined default using the PRIVATE, SHARED, FIRSTPRIVATE, LASTPRIVATE, and REDUCTION clauses. As a result, the following example is legal:

```
!$OMP PARALLEL DO DEFAULT(PRIVATE), FIRSTPRIVATE(I), SHARED(X),  
!$OMP& SHARED(R) LASTPRIVATE(I)
```

#### 2.6.2.4 FIRSTPRIVATE

The FIRSTPRIVATE clause provides a superset of the functionality provided by the PRIVATE clause.

This clause has the following format:

FIRSTPRIVATE ( <i>list</i> )
------------------------------

Variables that appear in the *list* are subject to PRIVATE clause semantics described in Section 2.6.2.1, page 26. In addition, private copies of the variables are initialized from the original object existing before the construct.

#### 2.6.2.5 LASTPRIVATE clause

The LASTPRIVATE clause provides a superset of the functionality provided by the PRIVATE clause.

This clause has the following format:

LASTPRIVATE ( <i>list</i> )
-----------------------------

Variables that appear in the *list* are subject to the PRIVATE clause semantics described in Section 2.6.2.1, page 26. When the LASTPRIVATE clause appears on a DO directive, the thread that executes the sequentially last iteration updates

the version of the object it had before the construct (see Section A.6, page 49 for an example). When the `LASTPRIVATE` clause appears in a `SECTIONS` directive, the thread that executes the lexically last `SECTION` updates the version of the object it had before the construct. Subobjects that are not assigned a value by the last iteration of the `DO` or the lexically last `SECTION` of the `SECTIONS` directive are undefined after the construct.

#### 2.6.2.6 REDUCTION clause

This clause performs a reduction on the variables that appear in *list*, with the operator *operator* or the intrinsic *intrinsic*, where *operator* is one of the following: `+`, `*`, `-`, `.AND.`, `.OR.`, `.EQV.`, or `.NEQV.`, and *intrinsic* is one of the following: `MAX`, `MIN`, `IAND`, `IOR`, or `IEOR`.

This clause has the following format:

<code>REDUCTION( { <i>operator</i>   <i>intrinsic</i> } : <i>list</i> )</code>
--------------------------------------------------------------------------------

Variables in *list* must be named scalar variables of intrinsic type. Since the intermediate values of the `REDUCTION` variables may be combined in random order, there is no guarantee that bit-identical results will be obtained for floating point reductions from one parallel run to another.

Variables that appear in a `REDUCTION` clause must be `SHARED` in the enclosing context. A private copy of each variable in *list* is created for each thread as if the `PRIVATE` clause had been used. The private copy is initialized according to the operator. See Table 1, page 28, for more information.

At the end of the `REDUCTION`, the shared variable is updated to reflect the result of combining the original value of the (shared) reduction variable with the final value of each of the private copies using the operator specified. The reduction operators are all associative (except for subtraction), and the compiler can freely reassociate the computation of the final value (the partial results of a subtraction reduction are added to form the final value).

The value of the shared variable becomes undefined when the first thread reaches the containing clause, and it remains so until the reduction computation is complete. Normally, the computation is complete at the end of the `REDUCTION` construct; however, if the `REDUCTION` clause is used on a construct to which `NOWAIT` is also applied, the shared variable remains undefined until a barrier synchronization has been performed to ensure that all the threads have completed the `REDUCTION` clause.

The REDUCTION clause is intended to be used on a region or work-sharing construct in which the reduction variable is used only in reduction statements with one of the following forms:

```
x = x operator expr
x = expr operator x (except for subtraction)
x = intrinsic (x, expr)
x = intrinsic (expr, x)
```

Some reductions can be expressed in other forms. For instance, a MAX reduction might be expressed as follows:

```
IF (x .LT. expr) x = expr
```

Alternatively, the reduction might be hidden inside a subroutine call. The user should be careful that the operator specified in the REDUCTION clause matches the reduction operation.

The following table lists the operators and intrinsics that are valid and their canonical initialization values. The actual initialization value will be consistent with the data type of the reduction variable.

Table 1. Initialization Values

<u>Operator/Intrinsic</u>	<u>Initialization</u>
+	0
*	1
-	0
.AND.	.TRUE.
.OR.	.FALSE.
.EQV.	.TRUE.
.NEQV.	.FALSE.
MAX	Smallest representable number
MIN	Largest representable number
IAND	All bits on
IOR	0

IEOR 0

See Section A.7, page 49 for an example that uses the + operator.

Any number of reduction clauses can be specified on the directive, but a variable can appear only once in a REDUCTION clause for that directive.

Example:

```
!$OMP DO REDUCTION(+: A, Y) REDUCTION(.OR.: AM)
```

### 2.6.2.7 COPYIN clause

The COPYIN clause applies only to common blocks that are declared as THREADPRIVATE. A COPYIN clause on a parallel region specifies that the data in the master thread of the team be copied to the thread private copies of the common block at the beginning of the parallel region.

This clause has the following format:

COPYIN( <i>list</i> )
-----------------------

It is not necessary to specify a whole common block to be copied in. Named variables appearing in the THREADPRIVATE common block can be specified in the *list*.

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 COPYIN clause must be declared to be a common block in the same scoping unit in which the COPYIN clause appears. See Section A.24, page 64 for more information.

In the following example, the common blocks BLK1 and FIELDS are specified as thread private, but only one of the variables in common block FIELDS is specified to be copied in.

```
COMMON /BLK1/ SCRATCH
COMMON /FIELDS/ XFIELD, YFIELD, ZFIELD
!$OMP THREADPRIVATE(/BLK1/, /FIELDS/)
!$OMP PARALLEL DEFAULT(PRIVATE) COPYIN(/BLK1/, ZFIELD)
```

An OpenMP-compliant implementation is required to ensure that the value of each threadprivate copy is the same as the value of the master thread copy when the master thread reached the directive containing the COPYIN clause.

### 2.6.3 Data environment rules

A program that conforms to the OpenMP Fortran API must adhere to the following rules and restrictions with respect to data scope:

1. Sequential DO loop control variables in the lexical extent of a PARALLEL region that would otherwise be SHARED based on default rules are automatically made private on the PARALLEL directive. Sequential DO loop control variables with no enclosing PARALLEL region are not classified automatically. It is up to the user to guarantee that these indexes are private if the containing procedures are called from a PARALLEL region.

All implied DO loop control variables and FORALL indexes are automatically made private at the enclosing implied DO or FORALL construct.

2. Variables that are privatized in a parallel region cannot be privatized again on an enclosed work-sharing directive. As a result, variables that appear in the PRIVATE, FIRSTPRIVATE, LASTPRIVATE, and REDUCTION clauses on a work-sharing directive must have shared scope in the enclosing parallel region.
3. A variable that appears in a PRIVATE, FIRSTPRIVATE, LASTPRIVATE, or REDUCTION clause must be definable.
4. Assumed-size and assumed-shape arrays cannot be specified as PRIVATE, FIRSTPRIVATE, or LASTPRIVATE. Array dummy arguments that are explicitly shaped (including variably dimensioned) can be declared in any scoping clause.
5. Fortran pointers and allocatable arrays can be declared as PRIVATE or SHARED but not as FIRSTPRIVATE or LASTPRIVATE.

Within a parallel region, the initial status of a private pointer is undefined. Private pointers that become allocated during the execution of a parallel region should be explicitly deallocated by the program prior to the end of the parallel region to avoid memory leaks.

The association status of a SHARED pointer becomes undefined upon entry to and on exit from the parallel construct if it is associated with a target or a subobject of a target that is PRIVATE, FIRSTPRIVATE, LASTPRIVATE, or REDUCTION inside the parallel construct. An allocatable array declared PRIVATE must have an allocation status of “not currently allocated” on entry to and on exit from the construct.

6. PRIVATE or SHARED attributes can be declared for a Cray pointer but not for the pointee. The scope attribute for the pointee is determined at the



---

point of pointer definition. It is illegal to declare a scope attribute for a pointee. Cray pointers may not be specified in `FIRSTPRIVATE` or `LASTPRIVATE` clauses.

7. Scope clauses apply only to variables in the static extent of the directive on which the clause appears, with the exception of variables passed as actual arguments. Local variables in called routines that do not have the `SAVE` attribute are `PRIVATE`. Common blocks and modules in called routines in the dynamic extent of a parallel region always have an implicit `SHARED` attribute, unless they are `THREADPRIVATE` common blocks. Local variables in called routines that have the `SAVE` attribute are `SHARED`.
8. When a named common block is specified in a `PRIVATE`, `FIRSTPRIVATE`, or `LASTPRIVATE` clause of a directive, none of its constituent elements may be declared in another scope attribute clause in that directive. It should be noted that when individual members of a common block are privatized, the storage of the specified variables is no longer associated with the storage of the common block itself. (See Section A.24, page 64 for examples.)
9. Variables that are not allowed in the `PRIVATE` and `SHARED` clauses are not affected by `DEFAULT(PRIVATE)` or `DEFAULT(SHARED)` clauses, respectively.
10. Clauses can be repeated as needed, but each variable and each named common block can appear explicitly in only one clause per directive, with the following exceptions:
  - A variable can be specified as both `FIRSTPRIVATE` and `LASTPRIVATE`.
  - Variables affected by the `DEFAULT` clause can be listed explicitly in a clause to override the default specification.
11. Variables that are specified as `LASTPRIVATE` for a work-sharing directive for which `NOWAIT` appears, must not be used prior to a barrier.
12. Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions must not be specified in `PRIVATE`, `FIRSTPRIVATE`, or `LASTPRIVATE` clauses.
13. The shared variables that are specified in `REDUCTION` or `LASTPRIVATE` clauses become defined at the end of the construct. Any concurrent uses or definitions of those variables must be synchronized with the definition that occurs at the end of the construct to avoid race conditions.
14. If the following three conditions hold regarding an actual argument in a reference to a non-intrinsic procedure, then any references to (or definitions

of) the shared storage that is associated with the dummy argument by any other thread must be synchronized with the procedure reference to avoid possible race conditions:

- a. The actual argument is one of the following:
  - A `SHARED` variable
  - A subobject of a `SHARED` variable
  - An object associated with a `SHARED` variable
  - An object associated with a subobject of a `SHARED` variable
- b. The actual argument is also one of the following:
  - An array section with a vector subscript
  - An array section
  - An assumed-shape array
  - A pointer array
- c. The associated dummy argument for this actual argument is an explicit-shape array or an assumed-size array.

The situations described above may result in the value of the shared variable being copied into temporary storage before the procedure reference, and back out of the temporary storage into the actual argument storage after the procedure reference. This effectively results in references to and definitions of the storage during the procedure reference.

15. An implementation that conforms to the OpenMP Fortran API must adhere to the following rule:
  - If a variable is specified as `FIRSTPRIVATE` and `LASTPRIVATE`, the implementation must ensure that the update required for `LASTPRIVATE` occurs after all initializations for `FIRSTPRIVATE`.

## 2.7 Directive binding

An implementation that conforms to the OpenMP Fortran API must adhere to the following rules with respect to the dynamic binding of directives:

- A parallel region is available for binding purposes, whether it is serialized or executed in parallel.

- The `DO`, `SECTIONS`, `SINGLE`, `MASTER`, and `BARRIER` directives bind to the dynamically enclosing `PARALLEL` directive, if one exists. (See Section A.18, page 57 for an example.)
- The `ORDERED` directive binds to the dynamically enclosing `DO` directive.
- The `ATOMIC` directive enforces exclusive access with respect to `ATOMIC` directives in all threads, not just the current team.
- The `CRITICAL` directive enforces exclusive access with respect to `CRITICAL` directives in all threads, not just the current team.
- A directive can never bind to any directive outside the closest enclosing `PARALLEL`.

## 2.8 Directive nesting

An implementation that conforms to the OpenMP Fortran API must adhere to the following rules with respect to the dynamic nesting of directives:

- A `PARALLEL` directive dynamically inside another `PARALLEL` directive logically establishes a new team, which is composed of only the current thread, unless nested parallelism is enabled.
- `DO`, `SECTIONS`, and `SINGLE` directives that bind to the same `PARALLEL` directive are not allowed to be nested one inside the other.
- `DO`, `SECTIONS`, and `SINGLE` directives are not permitted in the dynamic extent of `CRITICAL` and `MASTER` directives.
- `BARRIER` directives are not permitted in the dynamic extent of `DO`, `SECTIONS`, `SINGLE`, `MASTER`, and `CRITICAL` directives.
- `MASTER` directives are not permitted in the dynamic extent of `DO`, `SECTIONS`, and `SINGLE` directives.
- `ORDERED` sections are not allowed in the dynamic extent of `CRITICAL` sections.
- Any directive set that is legal when executed dynamically inside a `PARALLEL` region is also legal when executed outside a parallel region. When executed dynamically outside a user-specified parallel region, the directive is executed with respect to a team composed of only the master thread.

See Section A.16, page 54 for legal examples of directive nesting, and Section A.17, page 55 for invalid examples.



# Run-time Library Routines [3]

---

This section describes the OpenMP Fortran API run-time library routines that can be used to control and query the parallel execution environment. A set of general purpose lock routines are also provided.

OpenMP Fortran API run-time library routines are external procedures. In the following descriptions, *scalar\_integer\_expression* is a default scalar integer expression, and *scalar\_logical\_expression* is a default scalar logical expression. The return values of these routines are also of default kind.

## 3.1 Execution Environment Routines

The following sections describe the execution environment routines:

- Section 3.1.1, page 37, describes the `OMP_SET_NUM_THREADS` subroutine.
- Section 3.1.2, page 38, describes the `OMP_GET_NUM_THREADS` function.
- Section 3.1.3, page 38, describes the `OMP_GET_MAX_THREADS` function.
- Section 3.1.4, page 39, describes the `OMP_GET_THREAD_NUM` function.
- Section 3.1.5, page 39, describes the `OMP_GET_NUM_PROCS` function.
- Section 3.1.6, page 40, describes the `OMP_IN_PARALLEL` function.
- Section 3.1.7, page 40, describes the `OMP_SET_DYNAMIC` subroutine.
- Section 3.1.8, page 41, describes the `OMP_GET_DYNAMIC` function.
- Section 3.1.9, page 41, describes the `OMP_SET_NESTED` subroutine.
- Section 3.1.10, page 42, describes the `OMP_GET_NESTED` function.

### 3.1.1 `OMP_SET_NUM_THREADS` Subroutine

The `OMP_SET_NUM_THREADS` subroutine sets the number of threads to use for the next parallel region.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_SET_NUM_THREADS(scalar_integer_expression)
```

The *scalar\_integer\_expression* is evaluated, and its value is used as the number of threads to use. This function has effect only when called from serial portions of the program. If this function is called from a portion of the program where the `OMP_IN_PARALLEL` function returns `.TRUE.`, the behavior of the function is undefined. When dynamic adjustment of the number of threads is enabled, calls to `OMP_SET_NUM_THREADS` set the maximum number of threads to use for the next parallel region. For additional information on this subject, see the `OMP_SET_DYNAMIC()` subroutine described in Section 3.1.7, page 40, and the `OMP_GET_DYNAMIC()` function described in Section 3.1.8, page 41, and the example in Section A.11, page 51.

This call has precedence over the `OMP_NUM_THREADS` environment variable.

### 3.1.2 `OMP_GET_NUM_THREADS` Function

The `OMP_GET_NUM_THREADS` function returns the number of threads currently in the team executing the parallel region from which it is called.

This function has the following format:

INTEGER FUNCTION <code>OMP_GET_NUM_THREADS()</code>
-----------------------------------------------------

The `OMP_SET_NUM_THREADS()` call and the `OMP_NUM_THREADS` environment variable control the number of threads in a team. For more information on the `OMP_SET_NUM_THREADS()` call, see Section 3.1.1, page 37.

If the number of threads has not been explicitly set by the user, the default is implementation dependent. This function binds to the closest enclosing `PARALLEL` directive. For more information on the `PARALLEL` directive, see Section 2.2, page 9.

If this call is made from the serial portion of a program, or from a nested parallel region that is serialized, this function returns 1. (See Section A.14, page 52 for an example.)

### 3.1.3 `OMP_GET_MAX_THREADS` Function

The `OMP_GET_MAX_THREADS` function returns the maximum value that can be returned by calls to the `OMP_GET_NUM_THREADS()` function. For more information on `OMP_GET_NUM_THREADS()`, see Section 3.1.2, page 38.

This function has the following format:

```
INTEGER FUNCTION OMP_GET_MAX_THREADS ( )
```

If `OMP_SET_NUM_THREADS ( )` is used to change the number of threads, subsequent calls to `OMP_GET_MAX_THREADS ( )` will return the new value. This function can be used to allocate maximum sized per-thread data structures when the `OMP_SET_DYNAMIC ( )` subroutine is set to `.TRUE..` For more information on `OMP_SET_DYNAMIC ( )`, see Section 3.1.7, page 40.

This function has global scope and returns the maximum value whether executing from a serial region or a parallel region.

### 3.1.4 OMP\_GET\_THREAD\_NUM Function

The `OMP_GET_THREAD_NUM` function returns the thread number, within the team, that lies between 0 and `OMP_GET_NUM_THREADS ( ) - 1`, inclusive. (See the second example in Section A.14, page 52.) The master thread of the team is thread 0.

The format of this function is as follows:

```
INTEGER FUNCTION OMP_GET_THREAD_NUM ( )
```

This function binds to the closest enclosing `PARALLEL` directive. For more information on the `PARALLEL` directive, see Section 2.2, page 9.

When called from a serial region, `OMP_GET_THREAD_NUM` returns 0. When called from within a nested parallel region that is serialized, this function returns 0.

### 3.1.5 OMP\_GET\_NUM\_PROCS Function

The `OMP_GET_NUM_PROCS` function returns the number of processors that are available to the program.

The format of this function is as follows:

```
INTEGER FUNCTION OMP_GET_NUM_PROCS ( )
```

### 3.1.6 OMP\_IN\_PARALLEL Function

The `OMP_IN_PARALLEL` function returns `.TRUE.` if it is called from the dynamic extent of a region executing in parallel, and `.FALSE.` otherwise. A parallel region that is serialized is not considered to be a region executing in parallel.

The format of this function is as follows:

```
LOGICAL FUNCTION OMP_IN_PARALLEL ( )
```

This function has global scope. As a result, it will always return `.TRUE.` within the dynamic extent of a region executing in parallel, regardless of nested regions that are serialized.

### 3.1.7 OMP\_SET\_DYNAMIC Subroutine

The `OMP_SET_DYNAMIC` subroutine enables or disables dynamic adjustment of the number of threads available for execution of parallel regions.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_SET_DYNAMIC ( scalar_logical_expression )
```

If *scalar\_logical\_expression* evaluates to `.TRUE.`, the number of threads that are used for executing subsequent parallel regions can be adjusted automatically by the run-time environment to obtain the best use of system resources. As a consequence, the number of threads specified by the user is the maximum thread count. The number of threads always remains fixed over the duration of each parallel region and is reported by the `OMP_GET_NUM_THREADS ( )` function. For more information on the `OMP_GET_NUM_THREADS ( )` function, see Section 3.1.2, page 38.

If *scalar\_logical\_expression* evaluates to `.FALSE.`, dynamic adjustment is disabled. (See Section A.11, page 51 for an example.)

A call to `OMP_SET_DYNAMIC` has precedence over the `OMP_DYNAMIC` environment variable. For more information on the `OMP_DYNAMIC` environment variable, see Section 4.3, page 46.

The default for dynamic thread adjustment is implementation dependent. As a result, user codes that depend on a specific number of threads for correct



execution should explicitly disable dynamic threads. Implementations are not required to provide the ability to dynamically adjust the number of threads, but they are required to provide the interface in order to support portability across platforms.

### 3.1.8 OMP\_GET\_DYNAMIC Function

The OMP\_GET\_DYNAMIC function returns `.TRUE.` if dynamic thread adjustment is enabled and returns `.FALSE.` otherwise. For more information on dynamic thread adjustment, see Section 3.1.7, page 40.

The format of this function is as follows:

```
LOGICAL FUNCTION OMP_GET_DYNAMIC ( )
```

If the implementation does not implement dynamic adjustment of the number of threads, this function always returns `.FALSE.`

### 3.1.9 OMP\_SET\_NESTED Subroutine

The OMP\_SET\_NESTED subroutine enables or disables nested parallelism.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_SET_NESTED ( scalar_logical_expression )
```

If *scalar\_logical\_expression* evaluates to `.FALSE.`, which is the default, nested parallelism is disabled, and nested parallel regions are serialized and executed by the current thread. If set to `.TRUE.`, nested parallelism is enabled, and parallel regions that are nested can deploy additional threads to form the team.

This call has precedence over the OMP\_NESTED environment variable. For more information on the OMP\_NESTED environment variable, see Section 4.4, page 46.

When nested parallelism is enabled, the number of threads used to execute nested parallel regions is implementation dependent. As a result, OpenMP-compliant implementations are allowed to serialize nested parallel regions even when nested parallelism is enabled.

### 3.1.10 OMP\_GET\_NESTED Function

The `OMP_GET_NESTED` function returns `.TRUE.` if nested parallelism is enabled and `.FALSE.` if nested parallelism is disabled. For more information on nested parallelism, see Section 3.1.9, page 41.

The format of this function is as follows:

```
LOGICAL FUNCTION OMP_GET_NESTED( )
```

If an implementation does not implement nested parallelism, this function always returns `.FALSE.`.

## 3.2 Lock Routines

The OpenMP run-time library includes a set of general-purpose locking routines. The lock variable, *var*, must be accessed only through the routines described in this section. For all these routines, *var* should be of type integer and of a `KIND` large enough to hold an address. For example, on 64-bit addressable systems, the *var* may be declared as `INTEGER( SELECTED_INT_KIND( 18 ) )`.

The lock control routines are as follows:

- Section 3.2.1, page 42, describes the `OMP_INIT_LOCK` subroutine.
- Section 3.2.2, page 43, describes the `OMP_DESTROY_LOCK` subroutine.
- Section 3.2.3, page 43, describes the `OMP_SET_LOCK` subroutine.
- Section 3.2.4, page 43, describes the `OMP_UNSET_LOCK` subroutine.
- Section 3.2.5, page 44, describes the `OMP_TEST_LOCK` function.

See Section A.15, page 53 for an example using these subroutines.

### 3.2.1 OMP\_INIT\_LOCK Subroutine

The `OMP_INIT_LOCK` subroutine initializes a lock associated with lock variable *var* for use in subsequent calls.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_INIT_LOCK( var )
```

The initial state is unlocked. It is illegal to call this routine with a lock variable that is already associated with a lock.

### 3.2.2 OMP\_DESTROY\_LOCK Subroutine

The OMP\_DESTROY\_LOCK subroutine disassociates the given lock variable *var* from any locks.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_DESTROY_LOCK ( var )
```

It is illegal to call this routine with a lock variable that has not been initialized.

### 3.2.3 OMP\_SET\_LOCK Subroutine

The OMP\_SET\_LOCK subroutine forces the executing thread to wait until the specified lock is available.

The format of this subroutine is as follows:

```
SUBROUTINE OMP_SET_LOCK ( var )
```

The thread is granted ownership of the lock when it is available. It is illegal to call this routine with a lock variable that has not been initialized.

### 3.2.4 OMP\_UNSET\_LOCK Subroutine

The OMP\_UNSET\_LOCK subroutine releases the executing thread from ownership of the lock.

The format of this subroutine as follows:

```
SUBROUTINE OMP_UNSET_LOCK ( var )
```

The behavior is undefined if the thread does not own that lock. It is illegal to call this routine with a lock variable that has not been initialized.

### 3.2.5 OMP\_TEST\_LOCK Function

The OMP\_TEST\_LOCK function tries to set the lock associated with the lock variable *var*.

The format of this function is as follows:

LOGICAL FUNCTION OMP_TEST_LOCK( <i>var</i> )
----------------------------------------------

This function returns `.TRUE.` if the lock was set successfully, otherwise it returns `.FALSE.` It is illegal to call this routine with a lock variable that has not been initialized.

# Environment Variables [4]

---

This chapter describes the OpenMP Fortran API environment variables (or equivalent platform-specific mechanisms) that control the execution of parallel code. The names of environment variables must be uppercase. The values assigned to them are case insensitive.

## 4.1 OMP\_SCHEDULE Environment Variable

This variable applies only to DO and PARALLEL DO directives that have the schedule type RUNTIME. For more information on the DO directive, see Section 2.3.1, page 11. For more information on the PARALLEL DO directive, see Section 2.4.1, page 17.

The schedule type and chunk size for all such loops can be set at run time by setting this environment variable to any of the recognized schedule types and to an optional chunk size. For DO and PARALLEL DO directives that have a schedule type other than RUNTIME, this environment variable is ignored. The default value for this environment variable is implementation dependent. If the optional chunk size is not set, a chunk size of 1 is assumed, except in the case of a STATIC schedule. For a STATIC schedule, the default chunk size is set to the loop iteration space divided by the number of threads applied to the loop.

Examples:

```
setenv OMP_SCHEDULE "GUIDED,4"  
setenv OMP_SCHEDULE "dynamic"
```

## 4.2 OMP\_NUM\_THREADS Environment Variable

The OMP\_NUM\_THREADS environment variable sets the number of threads to use during execution, unless that number is explicitly changed by calling the OMP\_SET\_NUM\_THREADS( ) subroutine. For more information on the OMP\_SET\_NUM\_THREADS( ) subroutine, see Section 3.1.1, page 37.

When dynamic adjustment of the number of threads is enabled, the value of this environment variable is the maximum number of threads to use. The default value is implementation dependent.

Example:

```
setenv OMP_NUM_THREADS 16
```

### 4.3 OMP\_DYNAMIC Environment Variable

The OMP\_DYNAMIC environment variable enables or disables dynamic adjustment of the number of threads available for execution of parallel regions. For more information on parallel regions, see Section 2.2, page 9.

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

If set to FALSE, dynamic adjustment is disabled. The default condition is implementation dependent. For more information, see the OMP\_SET\_DYNAMIC subroutine described in Section 3.1.7, page 40.

Example:

```
setenv OMP_DYNAMIC TRUE
```

### 4.4 OMP\_NESTED Environment Variable

The OMP\_NESTED environment variable enables or disables nested parallelism. If set to TRUE, nested parallelism is enabled; if it is set to FALSE, it is disabled. The default value is FALSE. See also Section 3.1.9, page 41.

Example:

```
setenv OMP_NESTED TRUE
```

The following are examples of the constructs defined in this document.

## A.1 Executing a Simple Loop in Parallel

The following example shows how to parallelize a simple loop using the `PARALLEL DO` directive (specified in Section 2.4.1, page 17). The loop iteration variable is private by default, so it is not necessary to declare it explicitly.

```
!$OMP PARALLEL DO
      DO I=2,N
          B(I) = (A(I) + A(I-1)) / 2.0
      ENDDO
!$OMP END PARALLEL DO
```

The `END PARALLEL DO` directive is optional.

## A.2 Specifying Conditional Compilation

The following example illustrates the use of the conditional compilation sentinel (specified in Section 2.1.2, page 7). Assuming Fortran fixed source form, the following statement is illegal when using OpenMP constructs:

```
C234567890
!$ X(I) = X(I) + XLOCAL
```

With OpenMP compilation, the conditional compilation sentinel `!$` is treated as two spaces. As a result, the statement infringes on the statement label field. To be legal, the statement should begin after column 6, like any other fixed source form statement:

```
C234567890
!$   X(I) = X(I) + XLOCAL
```

In other words, conditionally compiled statements need to meet all applicable language rules when the sentinel is replaced with two spaces.

### A.3 Using Parallel Regions

The `PARALLEL` directive (specified in Section 2.2, page 9) can be used in coarse-grain parallel programs. In the following example, each thread in the parallel region decides what part of the global array `X` to work on based on the thread number:

```
!$OMP PARALLEL DEFAULT(PRIVATE) SHARED(X,NPOINTS)
    IAM = OMP_GET_THREAD_NUM()
    NP = OMP_GET_NUM_THREADS()
    IPOINTS = NPOINTS/NP
    CALL SUBDOMAIN(X, IAM, IPOINTS)
!$OMP END PARALLEL
```

### A.4 Using the `NOWAIT` Clause

If there are multiple independent loops within a parallel region, you can use the `NOWAIT` clause (specified in Section 2.3.1, page 11) to avoid the implied `BARRIER` at the end of the `DO` directive, as follows:

```
!$OMP PARALLEL
!$OMP DO
    DO I=2,N
        B(I) = (A(I) + A(I-1)) / 2.0
    ENDDO
!$OMP END DO NOWAIT
!$OMP DO
    DO I=1,M
        Y(I) = SQRT(Z(I))
    ENDDO
!$OMP END DO NOWAIT
!$OMP END PARALLEL
```

### A.5 Using the `CRITICAL` Directive

The following example (for Section 2.5.2, page 19) includes several `CRITICAL` directives. The example illustrates a queuing model in which a task is dequeued and worked on. To guard against multiple threads dequeuing the same task, the dequeuing operation must be in a critical section. Because there are two independent queues in this example, each queue is protected by `CRITICAL` directives with different names, `XAXIS` and `YAXIS`, respectively.



```

!$OMP PARALLEL DEFAULT(PRIVATE) SHARED(X,Y)
!$OMP CRITICAL(XAXIS)
      CALL DEQUEUE(IX_NEXT, X)
!$OMP END CRITICAL(XAXIS)
      CALL WORK(IX_NEXT, X)
!$OMP CRITICAL(YAXIS)
      CALL DEQUEUE(IY_NEXT, Y)
!$OMP END CRITICAL(YAXIS)
      CALL WORK(IY_NEXT, Y)
!$OMP END PARALLEL

```

## A.6 Using the LASTPRIVATE Clause

Correct execution sometimes depends on the value that the last iteration of a loop assigns to a variable. Such programs must list all such variables as arguments to a LASTPRIVATE clause (specified in Section 2.6.2.5, page 28) so that the values of the variables are the same as when the loop is executed sequentially.

```

!$OMP PARALLEL
!$OMP DO LASTPRIVATE(I)
      DO I=1,N
          A(I) = B(I) + C(I)
      ENDDO
!$OMP END PARALLEL
      CALL REVERSE(I)

```

In the preceding example, the value of I at the end of the parallel region will equal N+1, as in the sequential case.

## A.7 Using the REDUCTION Clause

The following example (for Section 2.6.2.6, page 29) shows how to use the REDUCTION clause:

```

!$OMP PARALLEL DO DEFAULT(PRIVATE) REDUCTION(+: A,B)
      DO I=1,N
          CALL WORK(ALOCAL, BLOCAL)
          A = A + ALOCAL
          B = B + BLOCAL
      ENDDO
!$OMP END PARALLEL DO

```

## A.8 Specifying Parallel Sections

In the following example (for Section 2.3.2, page 14), subroutines XAXIS, YAXIS, and ZAXIS can be executed concurrently. The first SECTION directive is optional. Note that all SECTION directives need to appear in the lexical extent of the PARALLEL SECTIONS/END PARALLEL SECTIONS construct.

```
!$OMP PARALLEL SECTIONS
!$OMP SECTION
    CALL XAXIS
!$OMP SECTION
    CALL YAXIS
!$OMP SECTION
    CALL ZAXIS
!$OMP END PARALLEL SECTIONS
```

## A.9 Using SINGLE Directives

The first thread that encounters the SINGLE directive (specified in Section 2.3.3, page 16) executes subroutines OUTPUT and INPUT. The user must not make any assumptions as to which thread will execute the SINGLE section. All other threads will skip the SINGLE section and stop at the barrier at the END SINGLE construct. If other threads can proceed without waiting for the thread executing the SINGLE section, a NOWAIT clause can be specified on the END SINGLE directive.

```
!$OMP PARALLEL DEFAULT(SHARED)
    CALL WORK(X)
!$OMP BARRIER
!$OMP SINGLE
    CALL OUTPUT(X)
    CALL INPUT(Y)
!$OMP END SINGLE
    CALL WORK(Y)
!$OMP END PARALLEL
```

## A.10 Specifying Sequential Ordering

ORDERED sections (specified in Section 2.5.6, page 23) are useful for sequentially ordering the output from work that is done in parallel. Assuming that a reentrant I/O library exists, the following program prints out the indexes in sequential order:

```

!$OMP DO ORDERED SCHEDULE(DYNAMIC)
  DO I=LB,UB,ST
    CALL WORK(I)
  END DO
  . . .
  SUBROUTINE WORK(K)
!$OMP ORDERED
  WRITE(*,*) K
!$OMP END ORDERED
  END

```

## A.11 Specifying a Fixed Number of Threads

Some programs rely on a fixed, prespecified number of threads to execute correctly. Because the default setting for the dynamic adjustment of the number of threads is implementation-dependent, such programs can choose to turn off the dynamic threads capability and set the number of threads explicitly to ensure portability. The following example (for Section 3.1.1, page 37) shows how to do this:

```

      CALL OMP_SET_DYNAMIC(.FALSE.)
      CALL OMP_SET_NUM_THREADS(16)
!$OMP PARALLEL DEFAULT(PRIVATE) SHARED(X,NPOINTS)
      IAM = OMP_GET_THREAD_NUM()
      IPOINTS = NPOINTS/16
      CALL DO_BY_16(X,IAM,IPOINTS)
!$OMP END PARALLEL

```

In this example, the program executes correctly only if it is executed by 16 threads. Note that the number of threads executing a parallel region remains constant during a parallel region, regardless of the dynamic threads setting. The dynamic threads mechanism determines the number of threads to use at the start of the parallel region and keeps it constant for the duration of the region.

## A.12 Using the ATOMIC Directive

The following example (for Section 2.5.4, page 20) avoids race conditions by protecting all simultaneous updates of the location, by multiple threads, with the ATOMIC directive:

```

!$OMP PARALLEL DO DEFAULT(PRIVATE) SHARED(X,Y,INDEX,N)
  DO I=1,N

```

```
        CALL WORK(XLOCAL, YLOCAL)
!$OMP ATOMIC
        X(INDEX(I)) = X(INDEX(I)) + XLOCAL
        Y(I) = Y(I) + YLOCAL
    ENDDO
```

Note that the `ATOMIC` directive applies only to the Fortran statement immediately following it. As a result, `Y` is not updated atomically in this example.

### A.13 Using the `FLUSH` Directive

The following example (for Section 2.5.5, page 21) uses the `FLUSH` directive for point-to-point synchronization between pairs of threads:

```
!$OMP PARALLEL DEFAULT(PRIVATE) SHARED(ISYNC)
    IAM = OMP_GET_THREAD_NUM()
    ISYNC(IAM) = 0
!$OMP BARRIER
    CALL WORK()
C    I AM DONE WITH MY WORK, SYNCHRONIZE WITH MY NEIGHBOR
    ISYNC(IAM) = 1
!$OMP FLUSH(ISYNC)
C    WAIT TILL NEIGHBOR IS DONE
    DO WHILE (ISYNC(NEIGH) .EQ. 0)
!$OMP FLUSH(ISYNC)
    END DO
!$OMP END PARALLEL
```

### A.14 Determining the Number of Threads Used

Consider the following incorrect example:

```
        NP = OMP_GET_NUM_THREADS()
!$OMP PARALLEL DO SCHEDULE(STATIC)
    DO I = 0, NP-1
        CALL WORK(I)
    ENDDO
!$OMP END PARALLEL DO
```

The `OMP_GET_NUM_THREADS()` call (specified in Section 3.1.2, page 38) returns 1 in the serial section of the code, so `NP` will always be equal to 1 in the

preceding example. To determine the number of threads that will be deployed for the parallel region, the call should be inside the parallel region.

The following example shows how to rewrite this program without including a query for the number of threads:

```
!$OMP PARALLEL PRIVATE(I)
      I = OMP_GET_THREAD_NUM()
      CALL WORK(I)
!$OMP END PARALLEL
```

## A.15 Using Locks

This is an example of the use of the lock routines (specified in Section 3.2, page 42).

In the following program, note that the argument to the lock routines should be of type `INTEGER` and of a `KIND` large enough to hold an address:

```
PROGRAM LOCK_USAGE
EXTERNAL OMP_TEST_LOCK
LOGICAL OMP_TEST_LOCK

INTEGER LCK          ! THIS VARIABLE SHOULD BE POINTER SIZED

CALL OMP_INIT_LOCK(LCK)
!$OMP PARALLEL SHARED(LCK) PRIVATE(ID)
      ID = OMP_GET_THREAD_NUM()
      CALL OMP_SET_LOCK(LCK)
      PRINT *, 'MY THREAD ID IS ', ID
      CALL OMP_UNSET_LOCK(LCK)
```

```
DO WHILE ( .NOT. OMP_TEST_LOCK(LCK) )
  CALL SKIP(ID)      ! WE DO NOT YET HAVE THE LOCK
                    ! SO WE MUST DO SOMETHING ELSE
END DO

CALL WORK(ID)       ! WE NOW HAVE THE LOCK
                    ! AND CAN DO THE WORK
CALL OMP_UNSET_LOCK( LCK )
!$OMP END PARALLEL

CALL OMP_DESTROY_LOCK( LCK )

END
```

## A.16 Nested DO Directives

The following example of directive nesting (specified in Section 2.8, page 35) is legal because the inner and outer DO directives bind to different PARALLEL regions:

```
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
!$OMP PARALLEL SHARED(I,N)
!$OMP DO
  DO J = 1, N
    CALL WORK(I,J)
  END DO
!$OMP END PARALLEL
  END DO
!$OMP END PARALLEL
```

The following variation of the preceding example is also legal:

```
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
    CALL SOME_WORK(I,N)
  END DO
!$OMP END PARALLEL

SUBROUTINE SOME_WORK(I,N)
!$OMP PARALLEL DEFAULT(SHARED)
```

```
!$OMP DO
  DO J = 1, N
    CALL WORK(I,J)
  END DO
!$OMP END PARALLEL
RETURN
END
```

## A.17 Examples Showing Incorrect Nesting of Work-sharing Directives

The examples in this section illustrate the directive nesting rules (specified in Section 2.8, page 35).

The following example is illegal because the inner and outer DO directives are nested and bind to the same PARALLEL directive:

```
PROGRAM WRONG1
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
!$OMP DO
  DO J = 1, N
    CALL WORK(I,J)
  END DO
  END DO
!$OMP END PARALLEL
END
```

The following dynamically nested version of the preceding example is also illegal:

```
PROGRAM WRONG2
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
    CALL SOME_WORK(I,N)
  END DO
!$OMP END PARALLEL
END
SUBROUTINE SOME_WORK(I,N)
!$OMP DO
  DO J = 1, N
    CALL WORK(I,J)
  END DO
RETURN
END
```

The following example is illegal because the DO and SINGLE directives are nested, and they bind to the same PARALLEL region:

```
PROGRAM WRONG3
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
!$OMP SINGLE
  CALL WORK(I)
!$OMP END SINGLE
  END DO
!$OMP END PARALLEL
END
```

The following example is illegal because a BARRIER directive inside a SINGLE or a DO can result in deadlock:

```
PROGRAM WRONG3
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
  DO I = 1, N
    CALL WORK(I)
!$OMP BARRIER
    CALL MORE_WORK(I)
  END DO
```



```
!$OMP END PARALLEL
      END
```

The following example is illegal because the `BARRIER` results in deadlock due to the fact that only one thread at a time can enter the critical section:

```
      PROGRAM WRONG4
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP CRITICAL
      CALL WORK(N,1)
!$OMP BARRIER
      CALL MORE_WORK(N,2)
!$OMP END CRITICAL
!$OMP END PARALLEL
      END
```

The following example is illegal because the `BARRIER` results in deadlock due to the fact that only one thread executes the `SINGLE` section:

```
      PROGRAM WRONG5
!$OMP PARALLEL DEFAULT(SHARED)
      CALL SETUP(N)
!$OMP SINGLE
      CALL WORK(N,1)
!$OMP BARRIER
      CALL MORE_WORK(N,2)
!$OMP END SINGLE
      CALL FINISH(N)
!$OMP END PARALLEL
      END
```

## A.18 Binding of `BARRIER` Directives

The directive binding rules call for a `BARRIER` directive to bind to the closest enclosing `PARALLEL` directive. (For more information, see Section 2.7, page 34.)

In the following example, the call from `MAIN` to `SUB2` is legal because the `BARRIER` (in `SUB3`) binds to the `PARALLEL` region in `SUB2`. The call from `MAIN` to `SUB1` is legal because the `BARRIER` binds to the `PARALLEL` region in subroutine `SUB2`.

```
      PROGRAM MAIN
      CALL SUB1(2)
      CALL SUB2(2)
```

```
END

SUBROUTINE SUB1(N)
!$OMP PARALLEL PRIVATE(I) SHARED(N)
!$OMP DO
DO I = 1, N
CALL SUB2(I)
END DO
!$OMP END PARALLEL
END

SUBROUTINE SUB2(K)
!$OMP PARALLEL SHARED(K)
CALL SUB3(K)
!$OMP END PARALLEL
END

SUBROUTINE SUB3(N)
CALL WORK(N)
!$OMP BARRIER
CALL WORK(N)
END
```

## A.19 Scoping Variables with the `PRIVATE` Clause

The values of `I` and `J` in the following example are undefined on exit from the parallel region:

```
INTEGER I, J
I = 1
J = 2
!$OMP PARALLEL PRIVATE(I) FIRSTPRIVATE(J)
I = 3
J = J+ 2
!$OMP END PARALLEL
PRINT *, I, J
```

(For more information, see Section 2.6.2.1, page 26.)

## A.20 Examples of Invalid Storage Association

The following examples illustrate the implications of the `PRIVATE` clause rules (see Section 2.6.2.1, page 26, rule 4) with regard to storage association:

### Example 1: Invalid Example

```

COMMON /BLOCK/ X
X = 1.0
!$OMP PARALLEL PRIVATE (X)
X = 2.0
CALL SUB()
...
!$OMP END PARALLEL
...
SUBROUTINE SUB()
COMMON /BLOCK/ X
...
PRINT *,X                ! X is undefined. The result of the
                        ! print is undefined.
...
END SUBROUTINE SUB
END PROGRAM

```

### Example 2: Invalid Example

```

COMMON /BLOCK/ X
X = 1.0
!$OMP PARALLEL PRIVATE (X)
X = 2.0
CALL SUB()
...
!$OMP END PARALLEL
...
CONTAINS
SUBROUTINE SUB()
COMMON /BLOCK/ Y
...
PRINT *,X                ! X is undefined.
PRINT *,Y                ! Y is undefined.
...
END SUBROUTINE SUB
END PROGRAM

```

**Example 3: Invalid Example**

```
      EQUIVALENCE (X,Y)
      X = 1.0
!$OMP PARALLEL PRIVATE(X)
      ...
      PRINT *,Y           ! Y is undefined.
      Y = 10
      PRINT *,X           ! X is undefined.
!$OMP END PARALLEL
```

**Example 4: Invalid Example**

```
      INTEGER A(100), B(100)
      EQUIVALENCE (A(51), B(1))

!$OMP PARALLEL DO DEFAULT(PRIVATE) PRIVATE(I,J) LASTPRIVATE(A)
      DO I=1,100
      DO J=1,100
      B(J) = J - 1
      ENDDO

      DO J=1,100
      A(J) = J           ! B becomes undefined at this point
      ENDDO
      DO J=1,50
      B(J) = B(J) + 1   ! Reference to B is not defined. A
                        ! becomes undefined at this point.
      ENDDO
      ENDDO
!$OMP END PARALLEL DO   ! The LASTPRIVATE write for A has
                        ! undefined results.

      PRINT *, B         ! B is undefined since the LASTPRIVATE
                        ! write of A was not defined.
END
```

**A.21 Examples of Syntax of Parallel DO Loops**

Both block-do and non-block-do are permitted with `PARALLEL DO` and `work-sharing DO` directives. However, if a user specifies an `ENDDO` directive for a non-block-do construct with shared termination, then the matching `DO`

directive must precede the outermost DO. (For more information, see Section 2.3.1, page 11 and Section 2.4.1, page 17.)

The following are some examples:

**Example 1:**

```

      DO 100 I = 1,10
!$OMP  DO
      DO 100 J = 1,10
      ...
100    CONTINUE

```

**Example 2:**

```

!$OMP DO
      DO 100 J = 1,10
      ...
100    A(I) = I + 1
!$OMP ENDDO

```

**Example 3:**

```

!$OMP DO
      DO 100 I = 1,10
      DO 100 J = 1,10
      ...
100    CONTINUE
!$OMP ENDDO

```

**Example 4: Invalid Example**

```

      DO 100 I = 1,10
!$OMP  DO
      DO 100 J = 1,10
      ...
100    CONTINUE
!$OMP ENDDO

```

## A.22 Examples of the ATOMIC Directive

All atomic references to the storage location of a variable that appears on the left-hand side of an ATOMIC assignment statement throughout the program are

required to have the same type and type parameters. (For more information, see Section 2.5.4, page 20.)

The following are some examples:

**Example 1: Invalid Example**

```
      INTEGER:: I
      REAL:: R
      EQUIVALENCE(I,R)
!$OMP PARALLEL
      ...
!$OMP ATOMIC
      I = I + 1
      ...
!$OMP ATOMIC
      R = R + 1.0
!$OMP END PARALLEL
```

**Example 2: Invalid Example**

```
      SUBROUTINE FRED()
      COMMON /BLK/ I
      INTEGER:: I
!$OMP PARALLEL
      ...
!$OMP ATOMIC
      I = I + 1
      ...
      CALL SUB()
!$OMP END PARALLEL

      SUBROUTINE SUB()
      COMMON /BLK/ R
      REAL:: R
      ...
!$OMP ATOMIC
      R = R + 1.0

      END SUBROUTINE
```

### Example 3: Invalid Example

Although the following example might work on some implementation, this is considered a non-conforming example.

```
        INTEGER:: I
        REAL:: R
        EQUIVALENCE(I,R)
!OMP PARALLEL
    ...
!OMP ATOMIC
    I = I + 1
!OMP END PARALLEL
    ...
!OMP PARALLEL
    ...
!OMP ATOMIC
    R = R + 1.0
!OMP END PARALLEL
```

## A.23 Examples of the ORDERED Directive

It is possible to have multiple ORDERED sections within a DO specified with the ORDERED clause. However, the following example is invalid, because the API states the following:

An iteration of a loop with a DO directive must not execute the same ORDERED directive more than once, and it must not execute more than one ORDERED directive.

For more information, see Section 2.5.6, page 23.

**Example 1: Invalid Example**

In this example, all iterations execute 2 ORDERED sections:

```
!$OMP DO
    DO I = 1, N
        ...
!$OMP ORDERED
        ...
!$OMP END ORDERED
        ...
!$OMP ORDERED
        ...
!$OMP END ORDERED
        ...
    END DO
```

**Example 2:**

This is a valid example of a DO with more than one ORDERED section:

```
!$OMP DO ORDERED
    DO I = 1,N
        ...
        IF ( I <= 10 ) THEN
            ...
!$OMP ORDERED
            WRITE(4,*) I
!$OMP END ORDERED
        ENDIF
        ...
        IF ( I > 10 ) THEN
            ...
!$OMP ORDERED
            WRITE(3,*) I
!$OMP END ORDERED
        ENDIF
    ENDDO
```

**A.24 Examples of THREADPRIVATE COMMON Data**

The following examples show two invalid uses and one correct use of the THREADPRIVATE directive. For more information, see Section 2.6.1, page 24 and Section 2.6.2.7, page 31.



**Example 1: Invalid Example**

```
MODULE FOO
COMMON /T/ A
END MODULE FOO

SUBROUTINE BAR
USE FOO
!$OMP THREADPRIVATE(/T/)
!$OMP PARALLEL
...
!$OMP END PARALLEL
END SUBROUTINE BAR
```

**Example 2: Invalid Example:**

```
COMMON /T/ A
!$OMP THREADPRIVATE(/T/)
...
CONTAINS
SUBROUTINE BAR
!$OMP PARALLEL COPYIN(/T/)
...
!$OMP END PARALLEL
END SUBROUTINE BAR
END PROGRAM
```

**Example 3: Correct Rewrite of Previous Example**

```
COMMON /T/ A
!$OMP THREADPRIVATE(/T/)
...
CONTAINS
SUBROUTINE BAR
COMMON /T/ A
!$OMP THREADPRIVATE(/T/)
!$OMP PARALLEL COPYIN(/T/)
...
!$OMP END PARALLEL
END SUBROUTINE BAR
END PROGRAM
```

## A.25 Examples of the Data Attribute Clauses: `SHARED` and `PRIVATE`

When a named common block is specified in a `PRIVATE`, `FIRSTPRIVATE` or `LASTPRIVATE` clause of a directive, none of its constituent elements may be declared in another scope attribute clause in that directive. The following examples, both valid and invalid, illustrate this point. (For more information, see item 8 of Section 2.6.3, page 32.)

### Example 1:

```
COMMON /C/ X,Y
!$OMP PARALLEL PRIVATE (/C/)
...
!$OMP END PARALLEL
...
!$OMP PARALLEL SHARED (X,Y)
...
!$OMP END PARALLEL
```

### Example 2:

```
COMMON /C/ X,Y
!$OMP PARALLEL
...
!$OMP DO PRIVATE(/C/)
...
!$OMP END DO
!
!$OMP DO PRIVATE(X)
...
!$OMP END DO
...
!$OMP END PARALLEL
```

### Example 3: Invalid Example

```
COMMON /C/ X,Y
!$OMP PARALLEL PRIVATE(/C/), SHARED(X)
...
!$OMP END PARALLEL
```

**Example 4:**

```
COMMON /C/ X,Y
!$OMP PARALLEL PRIVATE (/C/)
...
!$OMP END PARALLEL
...
!$OMP PARALLEL SHARED (/C/)
...
!$OMP END PARALLEL
```

**Example 5: Invalid Example**

```
COMMON /C/ X,Y
!$OMP PARALLEL PRIVATE(/C/), SHARED(/C/)
...
!$OMP END PARALLEL
```



# Stubs for Run-time Library Routines [B]

---

This section provides stubs for the runtime library routines defined in the OpenMP Fortran API. The stubs are provided to enable portability to platforms that do not support the OpenMP Fortran 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:** 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. Users should not make assumptions about mechanisms used by OpenMP Fortran implementations to implement locks based on the scheme used by the stub routines.

```
SUBROUTINE OMP_SET_NUM_THREADS(NP)
  INTEGER NP
  RETURN
END

INTEGER FUNCTION OMP_GET_NUM_THREADS()
  OMP_GET_NUM_THREADS = 1
  RETURN
END

INTEGER FUNCTION OMP_GET_MAX_THREADS()
  OMP_GET_MAX_THREADS = 1
  RETURN
END

INTEGER FUNCTION OMP_GET_THREAD_NUM()
  OMP_GET_THREAD_NUM = 0
  RETURN
END

INTEGER FUNCTION OMP_GET_NUM_PROCS()
  OMP_GET_NUM_PROCS = 1
  RETURN
END

SUBROUTINE OMP_SET_DYNAMIC(FLAG)
  LOGICAL FLAG
```

```
RETURN
END

LOGICAL FUNCTION OMP_GET_DYNAMIC( )
OMP_GET_DYNAMIC = .FALSE.
RETURN
END

LOGICAL FUNCTION OMP_IN_PARALLEL( )
OMP_IN_PARALLEL = .FALSE.
RETURN
END

SUBROUTINE OMP_SET_NESTED(FLAG)
LOGICAL FLAG
RETURN
END

LOGICAL FUNCTION OMP_GET_NESTED( )
OMP_GET_NESTED = .FALSE.
RETURN
END

SUBROUTINE OMP_INIT_LOCK(LOCK)
POINTER (LOCK,IL)
INTEGER IL
LOCK = -1
RETURN
END

SUBROUTINE OMP_DESTROY_LOCK(LOCK)
POINTER (LOCK,IL)
INTEGER IL
LOCK = 0
RETURN
END

SUBROUTINE OMP_SET_LOCK(LOCK)
POINTER (LOCK,IL)
INTEGER IL

IF(LOCK .EQ. 0) THEN
  PRINT*, 'ERROR: LOCK NOT INITIALIZED'
```

```
        STOP
    ELSEIF(LOCK .EQ. 1) THEN
        PRINT*, 'ERROR: DEADLOCK IN USING LOCK VARIABLE'
        STOP
    ELSE
        LOCK = 1
    ENDIF
    RETURN
END

SUBROUTINE OMP_UNSET_LOCK(LOCK)
    POINTER (LOCK,IL)
    INTEGER IL
    IF(LOCK .EQ. 0) THEN
        PRINT*, 'ERROR: LOCK NOT INITIALIZED'
        STOP
    ELSEIF(LOCK .EQ. 1) THEN
        LOCK = -1
    ELSE
        PRINT*, 'ERROR: LOCK NOT SET'
        STOP
    ENDIF
    RETURN
END

LOGICAL FUNCTION OMP_TEST_LOCK(LOCK)
    POINTER (LOCK,IL)
    INTEGER IL
    IF (LOCK .EQ. -1) THEN
        LOCK = 1
        OMP_TEST_LOCK = .TRUE.
    ELSEIF(LOCK .EQ. 1) THEN
        OMP_TEST_LOCK = .FALSE.
    ELSE
        PRINT*, 'ERROR: LOCK NOT INITIALIZED'
        STOP
    ENDIF
    RETURN
END
```





# Using the SCHEDULE Clause [C]

---

A parallel region has at least one barrier, at its end, and may have additional barriers within it. At each barrier, the other members of the team must wait for the last thread to arrive. To minimize this wait time, shared work should be distributed so that all threads arrive at the barrier at about the same time. If some of that shared work is contained in DO constructs, the SCHEDULE clause can be used for this purpose.

When there are repeated references to the same objects, the choice of schedule for a DO construct may be determined primarily by characteristics of the memory system, such as the presence and size of caches and whether memory access times are uniform or nonuniform. Such considerations may make it preferable to have each thread consistently refer to the same set of elements of an array in a series of loops, even if some threads are assigned relatively less work in some of the loops. This can be done by using the STATIC schedule with the same bounds for all the loops. In the following example, note that 1 is used as the lower bound in the second loop, even though K would be more natural if the schedule were not important.

```
!$OMP PARALLEL
!$OMP DO SCHEDULE(STATIC)
  DO I=1,N
    A(I) = WORK1(I)
  ENDDO
!$OMP DO SCHEDULE(STATIC)
  DO I=1,N
    IF(I .GE. K) A(I) = A(I) + WORK2(I)
  ENDDO
!$OMP END PARALLEL
ENDDO
```

In the remaining examples, it is assumed that memory access is not the dominant consideration, and, unless otherwise stated, that all threads receive comparable computational resources. In these cases, the choice of schedule for a DO construct depends on all the shared work that is to be performed between the nearest preceding barrier and either the implied closing barrier or the nearest subsequent barrier, if there is a NOWAIT clause. For each kind of schedule, a short example shows how that schedule kind is likely to be the best choice. A brief discussion follows each example.

The `STATIC` schedule is also appropriate for the simplest case, a parallel region containing a single `DO` construct, with each iteration requiring the same amount of work.

```
!$OMP PARALLEL DO SCHEDULE(STATIC)
  DO I=1,N
    CALL INVARIANT_AMOUNT_OF_WORK(I)
  ENDDO
```

The `STATIC` schedule is characterized by the properties that each thread gets approximately the same number of iterations as any other thread, and each thread can independently determine the iterations assigned to it. Thus no synchronization is required to distribute the work, and, under the assumption that each iteration requires the same amount of work, all threads should finish at about the same time.

For a team of  $P$  threads, let  $\text{CEILING}(N/P)$  be the integer  $Q$ , which satisfies  $N = P*Q - R$  with  $0 \leq R < P$ . One implementation of the `STATIC` schedule for this example would assign  $Q$  iterations to the first  $P-1$  threads, and  $Q-R$  iterations to the last thread. Another acceptable implementation would assign  $Q$  iterations to the first  $P-R$  threads, and  $Q-1$  iterations to the remaining  $R$  threads. This illustrates why a program should not rely on the details of a particular implementation.

The `DYNAMIC` schedule is appropriate for the case of a `DO` construct with the iterations requiring varying, or even unpredictable, amounts of work.

```
!$OMP PARALLEL DO SCHEDULE(DYNAMIC)
  DO I=1,N
    CALL UNPREDICTABLE_AMOUNT_OF_WORK(I)
  ENDDO
```

The `DYNAMIC` schedule is characterized by the property that no thread waits at the barrier for longer than it takes another thread to execute its final iteration. This requires that iterations be assigned one at a time to threads as they become available, with synchronization for each assignment. The synchronization overhead can be reduced by specifying a minimum chunk size  $K$  greater than 1, so that each thread is assigned  $K$  iterations at a time until fewer than  $K$  iterations remain. This guarantees that no thread waits at the barrier longer than it takes another thread to execute its final chunk of (at most)  $K$  iterations.

The `DYNAMIC` schedule can be useful if the threads receive varying computational resources, which has much the same effect as varying amounts of work for each iteration. Similarly, the `DYNAMIC` schedule can also be useful if

the threads arrive at the DO construct at varying times, though in some of these cases the GUIDED schedule may be preferable.

The GUIDED schedule is appropriate for the case in which the threads may arrive at varying times at a DO construct with each iteration requiring about the same amount of work. This can happen if, for example, the DO construct is preceded by one or more SECTIONS or DO constructs with NOWAIT clauses.

```
!$OMP PARALLEL
!$OMP SECTIONS
.....
!$OMP END SECTIONS NOWAIT

!$OMP DO SCHEDULE(GUIDED)
DO I=1,N
    CALL INVARIANT_AMOUNT_OF_WORK(I)
ENDDO
```

Like DYNAMIC, the GUIDED schedule guarantees that no thread waits at the barrier longer than it takes another thread to execute its final iteration, or final  $K$  iterations if a chunk size of  $K$  is specified. Among such schedules, the GUIDED schedule is characterized by the property that it requires the fewest synchronizations. For chunk size  $K$ , a typical implementation will assign  $Q = \text{CEILING}(N/P)$  iterations to the first available thread, set  $N$  to the larger of  $N-Q$  and  $P*K$ , and repeat until all iterations are assigned.

When the choice of the optimum schedule is not as clear as it is for these examples, the RUNTIME schedule is convenient for experimenting with different schedules and chunk sizes without having to modify and recompile the program. It can also be useful when the optimum schedule depends (in some predictable way) on the input data to which the program is applied.

To see an example of the trade-offs between different schedules, consider sharing 1000 iterations among 8 threads. Suppose there is an invariant amount of work in each iteration, and use that as the unit of time.

If all threads start at the same time, the STATIC schedule will cause the construct to execute in 125 units, with no synchronization. But suppose that one thread is 100 units late in arriving. Then the remaining seven threads wait for 100 units at the barrier, and the execution time for the whole construct increases to 225.

Because both the DYNAMIC and GUIDED schedules ensure that no thread waits for more than one unit at the barrier, the delayed thread causes their execution times for the construct to increase only to 138 units, possibly increased by

delays from synchronization. If such delays are not negligible, it becomes important that the number of synchronizations is 1000 for DYNAMIC but only 41 for GUIDED, assuming the default chunk size of one. With a chunk size of 25, DYNAMIC and GUIDED both finish in 150 units, plus any delays from the required synchronizations, which now number only 40 and 20, respectively.