



# OpenMP API User's Guide

---

Sun™ ONE Studio 8

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054 U.S.A.  
650-960-1300

Part No. 817-0933-10  
May 2003, Revision A

Send comments about this document to: [docfeedback@sun.com](mailto:docfeedback@sun.com)

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers. Portions of this product are derived in part from Cray90, a product of Cray Research, Inc.

libdwarf and libredblack are Copyright 2000 Silicon Graphics, Inc. and are available under the GNU Lesser General Public License from <http://www.sgi.com>.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Sun ONE Studio, the Solaris logo and the Sun ONE logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Netscape and Netscape Navigator are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits reserves.

Droits du gouvernement americain, utilisateurs gouvernementaux logiciel commercial. Les utilisateurs gouvernementaux sont soumis au contrat de licence standard de Sun Microsystems, Inc., ainsi qu aux dispositions en vigueur de la FAR (Federal Acquisition Regulations) et des supplements a celles-ci.

Distribue par des licences qui en restreignent l'utilisation.

Cette distribution peut comprendre des composants developpes par des tierces parties.

Des parties de ce produit pourront etre derivees Cray CF90, un produit de Cray Inc.

Des parties de ce produit pourront etre derivees des systemes Berkeley BSD licencies par l'Universite de Californie. UNIX est une marque deposee aux Etats-Unis et dans d'autres pays et licenciee exclusivement par X/Open Company, Ltd.

libdwarf et libredblack sont deposees 2000 Silicon Graphics, Inc. et sont disponible sous le GNU Moins Général Public Permis de <http://www.sgi.com>.

Sun, Sun Microsystems, le logo Sun, Java, Sun ONE Studio, le logo Solaris et le logo Sun ONE sont des marques de fabrique ou des marques deposees de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Netscape et Netscape Navigator sont des marques de fabrique ou des marques deposees de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisees sous licence et sont des marques de fabrique ou des marques deposees de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont bases sur une architecture developpee par Sun Microsystems, Inc.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont regis par la legislation americaine en matiere de controle des exportations et peuvent etre soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucleaires, des missiles, des armes biologiques et chimiques ou du nucleaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers des pays sous embargo des Etats-Unis, ou vers des entites figurant sur les listes d'exclusion d'exportation americaines, y compris, mais de maniere non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une facon directe ou indirecte, aux exportations des produits ou des services qui sont regis par la legislation americaine en matiere de controle des exportations et la liste de ressortissants specifiquement designes, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



# Contents

---

## **Before You Begin ix**

Typographic Conventions ix

Shell Prompts x

Accessing Compiler Collection Tools and Man Pages xi

Accessing Compiler Collection Documentation xiii

Accessing Related Solaris Documentation xv

Resources for Developers xvi

Contacting Sun Technical Support xvi

Sun Welcomes Your Comments xvi

## **1. OpenMP API Summary 1-1**

1.1 Where to Find the OpenMP Specifications 1-1

1.2 Special Conventions Used Here 1-2

1.3 Directive Formats 1-2

1.4 Conditional Compilation 1-3

1.5 **PARALLEL** - Parallel Region Construct 1-4

1.6 Work-Sharing Constructs 1-5

1.6.1 **DO** and **for** Constructs 1-5

1.6.2 **SECTIONS** Construct 1-6

1.6.3 **SINGLE** Construct 1-7

- 1.6.4 Fortran **WORKSHARE** Construct 1-8
- 1.7 Combined Parallel Work-sharing Constructs 1-8
  - 1.7.1 **PARALLEL DO** and **parallel for** Constructs 1-9
  - 1.7.2 **PARALLEL SECTIONS** Construct 1-9
  - 1.7.3 **PARALLEL WORKSHARE** Construct 1-10
- 1.8 Synchronization Constructs 1-10
  - 1.8.1 **MASTER** Construct 1-11
  - 1.8.2 **CRITICAL** Construct 1-11
  - 1.8.3 **BARRIER** Construct 1-12
  - 1.8.4 **ATOMIC** Construct 1-12
  - 1.8.5 **FLUSH** Construct 1-13
  - 1.8.6 **ORDERED** Construct 1-14
- 1.9 Data Environment Directives 1-15
  - 1.9.1 **THREADPRIVATE** Directive 1-15
- 1.10 OpenMP Directive Clauses 1-16
  - 1.10.1 Data Scoping Clauses 1-16
    - 1.10.1.1 **PRIVATE** Clause 1-16
    - 1.10.1.2 **SHARED** Clause 1-16
    - 1.10.1.3 **DEFAULT** Clause 1-16
    - 1.10.1.4 **FIRSTPRIVATE** Clause 1-17
    - 1.10.1.5 **LASTPRIVATE** Clause 1-17
    - 1.10.1.6 **COPYIN** Clause 1-17
    - 1.10.1.7 **COPYPRIVATE** Clause 1-18
    - 1.10.1.8 **REDUCTION** Clause 1-18
  - 1.10.2 Scheduling Clauses 1-19
    - 1.10.2.1 **STATIC** Scheduling 1-19
    - 1.10.2.2 **DYNAMIC** Scheduling 1-19
    - 1.10.2.3 **GUIDED** Scheduling 1-19

- 1.10.2.4 **RUNTIME** Scheduling 1-20
  - 1.10.3 **NUM\_THREADS** Clause 1-20
  - 1.10.4 Placement of Clauses on Directives 1-20
- 1.11 OpenMP Runtime Library Routines 1-22
  - 1.11.1 Fortran OpenMP Routines 1-22
  - 1.11.2 C/C++ OpenMP Routines 1-22
  - 1.11.3 Run-time Thread Management Routines 1-23
    - 1.11.3.1 **OMP\_SET\_NUM\_THREADS** Routine 1-23
    - 1.11.3.2 **OMP\_GET\_NUM\_THREADS** Routine 1-23
    - 1.11.3.3 **OMP\_GET\_MAX\_THREADS** Routine 1-23
    - 1.11.3.4 **OMP\_GET\_THREAD\_NUM** Routine 1-24
    - 1.11.3.5 **OMP\_GET\_NUM\_PROCS** Routine 1-24
    - 1.11.3.6 **OMP\_IN\_PARALLEL** Routine 1-24
    - 1.11.3.7 **OMP\_SET\_DYNAMIC** Routine 1-25
    - 1.11.3.8 **OMP\_GET\_DYNAMIC** Routine 1-25
    - 1.11.3.9 **OMP\_SET\_NESTED** Routine 1-25
    - 1.11.3.10 **OMP\_GET\_NESTED** Routine 1-26
  - 1.11.4 Routines That Manage Synchronization Locks 1-26
    - 1.11.4.1 **OMP\_INIT\_LOCK** and **OMP\_INIT\_NEST\_LOCK** Routines 1-27
    - 1.11.4.2 **OMP\_DESTROY\_LOCK** and **OMP\_DESTROY\_NEST\_LOCK** Routines 1-27
    - 1.11.4.3 **OMP\_SET\_LOCK** and **OMP\_SET\_NEST\_LOCK** Routines 1-27
    - 1.11.4.4 **OMP\_UNSET\_LOCK** and **OMP\_UNSET\_NEST\_LOCK** Routines 1-28
    - 1.11.4.5 **OMP\_TEST\_LOCK** and **OMP\_TEST\_NEST\_LOCK** Routines 1-28
  - 1.11.5 Timing Routines 1-28
    - 1.11.5.1 **OMP\_GET\_WTIME** Routine 1-29

- 2. Implementation-Dependent Issues 2-1**
- 3. Compiling for OpenMP 3-1**
  - 3.1 Compiler Options To Use 3-1
  - 3.2 Fortran 95 OpenMP Validation 3-3
  - 3.3 OpenMP Environment Variables 3-4
  - 3.4 Stacks and Stack Sizes 3-5
- 4. Converting to OpenMP 4-1**
  - 4.1 Converting Legacy Fortran Directives 4-1
    - 4.1.1 Converting Sun-Style Fortran Directives 4-2
      - 4.1.1.1 Issues Between Sun-Style Fortran Directives and OpenMP 4-3
    - 4.1.2 Converting Cray-Style Fortran Directives 4-4
      - 4.1.2.1 Issues Between Cray-Style Fortran Directives and OpenMP Directives 4-4
  - 4.2 Converting Legacy C Pragmas 4-4
    - 4.2.1 Issues Between Legacy C Pragmas and OpenMP 4-6

# Tables

---

TABLE 3-1	OpenMP Environment Variables	3-4
TABLE 3-2	Multiprocessing Environment Variables	3-5
TABLE 4-1	Converting Sun Parallelization Directives to OpenMP	4-2
TABLE 4-2	DOALL Qualifier Clauses and OpenMP Equivalent Clauses	4-2
TABLE 4-3	SCHEDTYPE Scheduling and OpenMP <code>schedule</code> Equivalents	4-3
TABLE 4-4	OpenMP Equivalents for Cray-Style DOALL Qualifier Clauses	4-4
TABLE 4-5	Converting Legacy C Parallelization Pragmas to OpenMP	4-5
TABLE 4-6	<code>taskloop</code> Optional Clauses and OpenMP Equivalents	4-5
TABLE 4-7	SCHEDTYPE Scheduling and OpenMP <code>schedule</code> Equivalents	4-6





# Before You Begin

---

The *OpenMP API User's Guide* summarizes the OpenMP Fortran 95, C, and C++ application program interface (API) for building multiprocessing applications. Sun™ ONE Studio compilers support the OpenMP API.

This guide is intended for scientists, engineers, and programmers who have a working knowledge of the Fortran, C, or C++ languages, and the OpenMP parallel programming model. Familiarity with the Solaris™ operating environment or UNIX® in general is also assumed.

---

## Typographic Conventions

TABLE P-1 Typeface Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
<b>AaBbCc123</b>	What you type, when contrasted with on-screen computer output	% <b>su</b> Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
<code>AaBbCc123</code>	Command-line placeholder text; replace with a real name or value	To delete a file, type <code>rm filename</code> .

**TABLE P-2** Code Conventions

Code Symbol	Meaning	Notation	Code Example
[ ]	Brackets contain arguments that are optional.	$O[n]$	<code>O4, O</code>
{ }	Braces contain a set of choices for a required option.	$d\{y n\}$	<code>dy</code>
	The “pipe” or “bar” symbol separates arguments, only one of which may be chosen.	$B\{dynamic static\}$	<code>Bstatic</code>
:	The colon, like the comma, is sometimes used to separate arguments.	$Rdir[:dir]$	<code>R/local/libs:/U/a</code>
...	The ellipsis indicates omission in a series.	$xinline=fl[,...fn]$	<code>xinline=alpha,dos</code>

## Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Superuser for Bourne shell and Korn shell	#

---

# Accessing Compiler Collection Tools and Man Pages

The compiler collection components and man pages are not installed into the standard `/usr/bin/` and `/usr/share/man` directories. To access the compilers and tools, you must have the compiler collection component directory in your `PATH` environment variable. To access the man pages, you must have the compiler collection man page directory in your `MANPATH` environment variable.

For more information about the `PATH` variable, see the `csh(1)`, `sh(1)`, and `ksh(1)` man pages. For more information about the `MANPATH` variable, see the `man(1)` man page. For more information about setting your `PATH` variable and `MANPATH` variables to access this release, see the installation guide or your system administrator.

---

**Note** – The information in this section assumes that your Sun ONE Studio compiler collection components are installed in the `/opt` directory. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

## Accessing the Compilers and Tools

Use the steps below to determine whether you need to change your `PATH` variable to access the compilers and tools.

### ▼ To Determine Whether You Need to Set Your `PATH` Environment Variable

1. Display the current value of the `PATH` variable by typing the following at a command prompt.

```
% echo $PATH
```

2. Review the output to find a string of paths that contain `/opt/SUNWspro/bin/`.

If you find the path, your `PATH` variable is already set to access the compilers and tools. If you do not find the path, set your `PATH` environment variable by following the instructions in the next procedure.

## ▼ To Set Your `PATH` Environment Variable to Enable Access to the Compilers and Tools

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `PATH` environment variable.

```
/opt/SUNWspro/bin
```

## Accessing the Man Pages

Use the following steps to determine whether you need to change your `MANPATH` variable to access the man pages.

## ▼ To Determine Whether You Need to Set Your `MANPATH` Environment Variable

1. Request the `dbx` man page by typing the following at a command prompt.

```
% man dbx
```

2. Review the output, if any.

If the `dbx(1)` man page cannot be found or if the man page displayed is not for the current version of the software installed, follow the instructions in the next procedure for setting your `MANPATH` environment variable.

## ▼ To Set Your `MANPATH` Environment Variable to Enable Access to the Man Pages

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `MANPATH` environment variable.

```
/opt/SUNWspro/man
```

---

# Accessing Compiler Collection Documentation

You can access the documentation at the following locations:

- The documentation is available from the documentation index that is installed with the software on your local system or network at `file:/opt/SUNWspro/docs/index.html`.

If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

- Most manuals are available from the `docs.sun.comsm` web site. The following titles are available through your installed software only:
  - *Standard C++ Library Class Reference*
  - *Standard C++ Library User's Guide*
  - *Tools.h++ Class Library Reference*
  - *Tools.h++ User's Guide*
- The release notes are available from the `docs.sun.com` web site.

The `docs.sun.com` web site (<http://docs.sun.com>) enables you to read, print, and buy Sun Microsystems manuals through the Internet. If you cannot find a manual, see the documentation index that is installed with the software on your local system or network.

---

**Note** – Sun is not responsible for the availability of third-party web sites mentioned in this document and does not endorse and is not responsible or liable for any content, advertising, products, or other materials on or available from such sites or resources. Sun will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services available on or through any such sites or resources.

---

# Documentation in Accessible Formats

The documentation is provided in accessible formats that are readable by assistive technologies for users with disabilities. You can find accessible versions of documentation as described in the following table. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

Type of Documentation	Format and Location of Accessible Version
Manuals (except third-party manuals)	HTML at <a href="http://docs.sun.com">http://docs.sun.com</a>
Third-party manuals: <ul style="list-style-type: none"><li>• <i>Standard C++ Library Class Reference</i></li><li>• <i>Standard C++ Library User's Guide</i></li><li>• <i>Tools.h++ Class Library Reference</i></li><li>• <i>Tools.h++ User's Guide</i></li></ul>	HTML in the installed software through the documentation index at <code>file:/opt/SUNWspr/docs/index.html</code>
Readmes and man pages	HTML in the installed software through the documentation index at <code>file:/opt/SUNWspr/docs/index.html</code>
Release notes	HTML at <a href="http://docs.sun.com/">http://docs.sun.com/</a>

---

## Related Compiler Collection Documentation

The following table describes related documentation that is available at `file:/opt/SUNWspr/docs/index.html` and <http://docs.sun.com>. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

Document Title	Description
<i>Fortran Programming Guide</i>	Describes how to write effective Fortran code on Solaris environments; input/output, libraries, performance, debugging, and parallel processing.
<i>Fortran Library Reference</i>	Details the Fortran library and intrinsic routines
<i>Fortran User's Guide</i>	Describes the compile-time environment and command-line options for the f95 compiler. Also includes guidelines for migrating legacy f77 programs to f95.

---

Document Title	Description
<i>C User's Guide</i>	Describes the compile-time environment and command-line options for the <code>cc</code> compiler.
<i>C++ User's Guide</i>	Describes the compile-time environment and command-line options for the <code>CC</code> compiler.
<i>Numerical Computation Guide</i>	Describes issues regarding the numerical accuracy of floating-point computations.

## Accessing Related Solaris Documentation

The following table describes related documentation that is available through the `docs.sun.com` web site.

Document Collection	Document Title	Description
Solaris Reference Manual Collection	See the titles of man page sections.	Provides information about the Solaris operating environment.
Solaris Software Developer Collection	<i>Linker and Libraries Guide</i>	Describes the operations of the Solaris link-editor and runtime linker.
Solaris Software Developer Collection	<i>Multithreaded Programming Guide</i>	Covers the POSIX and Solaris threads APIs, programming with synchronization objects, compiling multithreaded programs, and finding tools for multithreaded programs.

---

## Resources for Developers

Visit <http://www.sun.com/developers/studio> and click the Compiler Collection link to find these frequently updated resources:

- Articles on programming techniques and best practices
- A knowledge base of short programming tips
- Documentation of compiler collection components, as well as corrections to the documentation that is installed with your software
- Information on support levels
- User forums
- Downloadable code samples
- New technology previews

You can find additional resources for developers at <http://www.sun.com/developers/>.

---

## Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

---

## Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. Email your comments to Sun at this address:

[docfeedback@sun.com](mailto:docfeedback@sun.com)

Please include the part number (817-0933-10) of the document in the subject line of your email.



# OpenMP API Summary

---

The OpenMP™ Application Program Interface is a portable, parallel programming model for shared memory multiprocessor architectures, developed in collaboration with a number of computer vendors. The specifications were created and are published by the OpenMP Architecture Review Board. For more information on the OpenMP developer community, including tutorials and other resources, see their web site at:

<http://www.openmp.org/>.

The OpenMP API is the recommended parallel programming model for all Sun ONE Studio compilers on SPARC® and UltraSPARC® platforms. See Chapter 4 for guidelines on converting legacy Fortran and C parallelization directives to OpenMP.

This chapter summarizes the directives, run-time library routines, and environment variables comprising the OpenMP Version 2.0 Application Program Interfaces, as implemented by the Sun ONE Studio Fortran 95, C and C++ compilers.

---

## 1.1 Where to Find the OpenMP Specifications

The material presented in this chapter *is only a summary* with many details left out intentionally for the sake of brevity. In all cases, refer to the OpenMP specification documents for complete details.

The Fortran and C/C++ OpenMP 2.0 specifications can be found on the official OpenMP website, <http://www.openmp.org/>.

---

## 1.2 Special Conventions Used Here

In the tables and examples that follow, Fortran directives and source code are shown in upper case, but are case-insensitive.

The term *structured-block* refers to a block of Fortran or C/C++ statements having no transfers into or out of the block.

Constructs within square brackets, [...], are optional.

Throughout this manual, “Fortran” refers to the Fortran 95 language and compiler, **f95**.

The terms “directive” and “pragma” are used interchangeably in this manual.

---

## 1.3 Directive Formats

Only one *directive-name* can be specified on a directive line, and applies to the succeeding program statement.

### **Fortran:**

Fortran fixed format accepts three directive “sentinels”, free format only one. In the Fortran examples that follow, free format will be used.

### **C/C++:**

C and C++ use the standard preprocessing directive starting with **#pragma omp**.

---

### **OpenMP 2.0 Fortran**

---

#### **Fixed Format:**

C\$OMP *directive-name optional\_clauses...*

!\$OMP *directive-name optional\_clauses...*

\*\$OMP *directive-name optional\_clauses...*

The sentinel must start in column one; continuation lines must have a non-blank or non-zero character in column 6.

Comments may appear after column 6 on the directive line, initiated by an exclamation point (!). The rest of the line after the ! is ignored.

---

---

**OpenMP 2.0 Fortran**

---

**Free Format:**

!\$OMP *directive-name optional\_clauses...*

May appear anywhere on a line, preceded only by whitespace; an ampersand (&) at the end of the line identifies a continued line.

Comments may appear on the directive line, initiated by an exclamation point (!). The rest of the line is ignored.

---

---

**OpenMP 2.0 C/C++**

---

#pragma omp *directive-name optional\_clauses...*

Each pragma must end with a new-line character, and follows the conventions of standard C and C++ for compiler pragmas.

Pragmas are case sensitive. The order in which clauses appear is not significant. White space can appear after and before the # and between words.

The directive applies to the succeeding statement, which must be a structured block.

---

---

## 1.4 Conditional Compilation

The OpenMP API defines the preprocessor symbol `_OPENMP` to be used for conditional compilation. In addition, OpenMP Fortran API accepts a conditional compilation sentinel.

---

**OpenMP 2.0 Fortran**

---

**Fixed Format:**

!\$ *fortran\_95\_statement*

C\$ *fortran\_95\_statement*

\*\$ *fortran\_95\_statement*

c\$ *fortran\_95\_statement*

The sentinel must start in column 1 and have no intervening blanks. With OpenMP compilation enabled, the sentinel is replaced by two blanks. The rest of the line must conform to standard Fortran fixed format conventions. Example:

C23456789

!\$ 10 iam = OMP\_GET\_THREAD\_NUM() +

!\$ 1        index

---

---

**OpenMP 2.0 Fortran**

---

**Free Format:**

```
!$ fortran_95_statement
```

This sentinel can appear in any column, preceded only by white space, and must appear as a single word. Fortran free format conventions apply to the rest of the line.

Example:

```
C23456789
!$ iam = OMP_GET_THREAD_NUM() +      &
!$&          index
```

**Fortran Preprocessor:**

Compiling with OpenMP enabled defines the preprocessor symbol `_OPENMP`.

```
#ifdef _OPENMP
    iam = OMP_GET_THREAD_NUM()+index
#endif
```

---

---

**OpenMP 2.0 C/C++**

---

**C/C++ Preprocessor:**

Compiling with OpenMP enabled defines the macro `_OPENMP`.

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

---

---

## 1.5 **PARALLEL** - Parallel Region Construct

The **PARALLEL** directive defines a parallel region, which is a region of the program that must be executed by multiple threads in parallel.

---

**OpenMP 2.0 Fortran**

---

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

---

---

**OpenMP 2.0 C/C++**

---

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

---

There are many special conditions and restrictions. Programmers are urged to refer to the appropriate OpenMP specification document for the details.

TABLE 1-1 identifies the clauses that can appear with this construct.

---

## 1.6 Work-Sharing Constructs

Work-sharing constructs divide the execution of the enclosed code region among the members of the team of threads that encounter it. Work sharing constructs must be enclosed within a parallel region for the construct to execute in parallel.

There are many special conditions and restrictions on these directives and the code they apply to. Programmers are urged to refer to the appropriate OpenMP specification document for the details.

### 1.6.1 DO and **for** Constructs

Specifies that the iterations of the **DO** or **for** loop that follows must be executed in parallel.

---

**OpenMP 2.0 Fortran**

---

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

The **DO** directive specifies that the iterations of the **DO** loop that immediately follows should be executed in parallel. This directive must appear within a parallel region to be effective.

---

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

The `for` pragma specifies that the iterations of the *for-loop* that immediately follows should be executed in parallel. This pragma must appear within a parallel region to be effective. The `for` pragma places restrictions on the structure of the corresponding `for` loop, and it must have *canonical shape*:

```
for (initexpr; var logicop b; increxpr)
```

where:

- *initexpr* is one of the following:
  - var* = *lb*
  - integer\_type var* = *lb*
- *increxpr* is one of the following expression forms:
  - ++var*
  - var++*
  - var*
  - var--*
  - var += incr*
  - var -= incr*
  - var = var + incr*
  - var = incr + var*
  - var = var - incr*
- *var* is a signed integer variable, made implicitly private for the range of the `for`. *var* must not be modified within the body of the `for` statement. Its value is indeterminate after the loop, unless specified `lastprivate`.
- *logicop* is one of the following logical operators:
  - < <= > >=
- *lb*, *b*, and *incr* are loop invariant integer expressions.

There are further restrictions on the use of < or <= and > or >= as *logicalop* in the `for` statement. See the OpenMP C/C++ specifications for details.

---

## 1.6.2 SECTIONS Construct

The `SECTIONS` construct encloses a non-iterative block of code to be divided among threads in the team. Each block is executed once by a thread in the team.

Each section is preceded by a **SECTION** directive, which is optional for the first section.

---

**OpenMP 2.0 Fortran**

---

```
!$OMP SECTIONS [clause[,] clause...]  
[!$OMP SECTION  
  structured-block  
!$OMP SECTION  
  structured-block ]  
...  
!$OMP END SECTIONS [NOWAIT]
```

---

---

**OpenMP 2.0 C/C++**

---

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

---

TABLE 1-1 identifies the clauses that can appear with this construct.

## 1.6.3 **SINGLE** Construct

The structured block enclosed by **SINGLE** is executed by only one thread in the team. Threads in the team that are not executing the **SINGLE** block wait at the end of the block unless **NOWAIT** is specified.

---

**OpenMP 2.0 Fortran**

---

```
!$OMP SINGLE [clause[,] clause...]  
  structured-block  
!$OMP END SINGLE [end-modifier]
```

---

---

**OpenMP 2.0 C/C++**

---

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

---

TABLE 1-1 identifies the clauses that can appear with this construct.

## 1.6.4 Fortran **WORKSHARE** Construct

The **WORKSHARE** construct divides the work of executing the enclosed code block into separate units of work, and causes the threads of the team to share the work such that each unit is executed only once.

---

**OpenMP 2.0 Fortran**

---

```
!$OMP WORKSHARE  
structured-block  
!$OMP END WORKSHARE [NOWAIT]
```

---

There is no C/C++ equivalent to the Fortran **WORKSHARE** construct.

---

## 1.7 Combined Parallel Work-sharing Constructs

The combined parallel work-sharing constructs are shortcuts for specifying a parallel region that contains one work-sharing construct.

There are many special conditions and restrictions on these directives and the code they apply to. Refer to the appropriate OpenMP specification document for the complete details. The description that follows is intended only as a summary and is not complete.

TABLE 1-1 identifies the clauses that can appear with these constructs.



## 1.7.1 PARALLEL DO and `parallel for` Constructs

Shortcut for specifying a parallel region that contains a single `DO` or `for` loop. Equivalent to a `PARALLEL` directive followed immediately by a `DO` or `for` directive. *clause* can be any of the clauses accepted by the `PARALLEL` and `DO/for` directives, except the `NOWAIT` modifier.

---

### OpenMP 2.0 Fortran

---

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

---

---

### OpenMP 2.0 C/C++

---

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

---

## 1.7.2 PARALLEL SECTIONS Construct

Shortcut for specifying a parallel region that contains a single `SECTIONS` directive. Equivalent to a `PARALLEL` directive followed by a `SECTIONS` directive. *clause* can be any of the clauses accepted by the `PARALLEL` and `SECTIONS` directives, except the `NOWAIT` modifier.

---

### OpenMP 2.0 Fortran

---

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

---

---

**OpenMP 2.0 C/C++**

---

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

---

### 1.7.3 PARALLEL WORKSHARE Construct

The Fortran **PARALLEL WORKSHARE** construct provides a shortcut for specifying a parallel region that contains a single **WORKSHARE** directive. *clause* can be one of the clauses accepted by either the **PARALLEL** or **WORKSHARE** directive.

---

**OpenMP 2.0 Fortran**

---

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

---

There is no C/C++ equivalent.

---

## 1.8 Synchronization Constructs

The following constructs specify thread synchronization. There are many special conditions and restrictions regarding these constructs that are too numerous to summarize here. Programmers are urged to refer to the appropriate OpenMP specification document for the complete details.

## 1.8.1 MASTER Construct

Only the master thread of the team executes the block enclosed by this directive. The other threads skip this block and continue. There is no implied barrier on entry to or exit from the master construct.

---

### OpenMP 2.0 Fortran

---

```
!$OMP MASTER
  structured-block
!$OMP END MASTER
```

---

---

### OpenMP 2.0 C/C++

---

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

---

## 1.8.2 CRITICAL Construct

Restrict access to the structured block to only one thread at a time. The optional *name* argument identifies the critical region. All unnamed **CRITICAL** directives map to the same name. Critical section names are global entities of the program and must be unique. For Fortran, if *name* appears on the **CRITICAL** directive, it must also appear on the **END CRITICAL** directive. For C/C++, the identifier used to name a critical region has external linkage and is in a name space which is separate from the name spaces used by labels, tags, members, and ordinary identifiers.

---

### OpenMP 2.0 Fortran

---

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

---

---

### OpenMP 2.0 C/C++

---

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

---

## 1.8.3 **BARRIER** Construct

Synchronizes all the threads in a team. Each thread waits until all the others in the team have reached this point.

---

### OpenMP 2.0 Fortran

---

```
!$OMP BARRIER
```

---

---

### OpenMP 2.0 C/C++

---

```
#pragma omp barrier
```

---

After all threads in the team have encountered the barrier, each thread in the team begins executing the statements after the **BARRIER** directive in parallel.

Note that because the **barrier** pragma does not have a C/C++ statement as part of its syntax, there are restrictions on its placement within a program. See the C/C++ OpenMP specifications for details.

## 1.8.4 **ATOMIC** Construct

Ensures that a specific memory location is to be updated atomically, rather than exposing it to the possibility of multiple, simultaneous writing threads.

*This implementation replaces all **ATOMIC** directives by enclosing the expression-statement in a critical section.*

---

**OpenMP 2.0 Fortran**

---

```
!$OMP ATOMIC  
  expression-statement
```

The directive applies only to the immediately following statement, which must be in one of these forms:

```
x = x operator expression  
x = expression operator x  
x = intrinsic(x, expr-list)  
x = intrinsic(expr-list, x)
```

where:

- *x* is a scalar of intrinsic type
  - *expression* is a scalar expression that does not reference *x*
  - *expr-list* is a non-empty, comma-separated list of scalar expressions that do not reference *x* (see the OpenMP 2.0 Fortran specifications for details)
  - *intrinsic* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
  - *operator* is one of **+** **-** **\*** **/** **.AND.** **.OR.** **.EQV.** **.NEQV.**
- 

---

**OpenMP 2.0 C/C++**

---

```
#pragma omp atomic  
  expression-statement
```

The pragma applies only to the immediately following statement, which must be in one of these forms:

```
x binop = expr  
x++  
++x  
x--  
--x
```

where:

- *x* in an lvalue expression with scalar type.
  - *expr* is an expression with scalar type that does not reference *x*.
  - *binop* is not an overloaded operator and one of: **+**, **\***, **-**, **/**, **&**, **^**, **|**, **<<**, or **>>**.
- 

## 1.8.5 FLUSH Construct

Thread-visible Fortran variables or C objects are written back to memory at the point at which this directive appears. The **FLUSH** directive only provides consistency between operations within the executing thread and global memory. The optional

*variable-list* consists of a comma-separated list of variables or objects that need to be flushed. A **FLUSH** directive without a *variable-list* synchronizes all thread-visible shared variables or objects.

---

**OpenMP 2.0 Fortran**

---

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

---

---

**OpenMP 2.0 C/C++**

---

```
#pragma omp flush [(variable-list)]
```

---

Note that because the `flush` pragma does not have a C/C++ statement as part of its syntax, there are restrictions on its placement within a program. See the C/C++ OpenMP specifications for details.

## 1.8.6 ORDERED Construct

The enclosed block is executed in the order that iterations would be executed in a sequential execution of the loop.

---

**OpenMP 2.0 Fortran**

---

```
!$OMP ORDERED  
  structured-block  
!$OMP END ORDERED
```

The enclosed block is executed in the order that iterations would be executed in a sequential execution of the loop. It can appear only in the dynamic extent of a **DO** or **PARALLEL DO** directive. The **ORDERED** clause must be specified on the closest **DO** directive enclosing the block.

A loop to which a **DO** directive applies must not execute the same **ordered** directive more than once per iteration, and it must not execute more than one **ordered** directive.

---

---

**OpenMP 2.0 C/C++**

---

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

The enclosed block is executed in the order that iterations would be executed in a sequential execution of the loop. It can appear only in the dynamic extent of a **for** or **parallel for** directive with the **ordered** clause specified.

A loop with a **for** construct must not execute the same **ordered** directive more than once per iteration, and it must not execute more than one **ordered** directive.

---

---

## 1.9 Data Environment Directives

The following directives control the data environment during execution of parallel constructs.

### 1.9.1 **THREADPRIVATE** Directive

Makes the *list* of objects (Fortran common blocks and named variables, C and C++ named variables) private to a thread but global within the thread.

*See the OpenMP specifications (section 2.6.1 in the Fortran specifications, section 2.7.1 in the C/C++ specifications) for the complete details and restrictions.*

---

**OpenMP 2.0 Fortran**

---

```
!$OMP THREADPRIVATE (list)
```

Common block names must appear between slashes. To make a common block **THREADPRIVATE**, this directive must appear after every **COMMON** declaration of that block.

---

---

**OpenMP 2.0 C/C++**

---

```
#pragma omp threadprivate (list)
```

Each variable in *list* at file, namespace, or block scope must refer to a variable declaration at file, namespace, or block scope that lexically precedes the pragma.

---

---

## 1.10 OpenMP Directive Clauses

This section summarizes the data scoping and scheduling clauses that can appear on OpenMP directives.

### 1.10.1 Data Scoping Clauses

Several directives accept clauses that allow a user to control the scope attributes of variables within the extent of the construct. If no data scope clause is specified for a directive, the default scope for variables affected by the directive is **SHARED**.

**Fortran:** *list* is a comma-separated list of named variables or common blocks that are accessible in the scoping unit. Common block names must appear within slashes (for example, `/ABLOCK/`).

*There are important restrictions on the use of these scoping clauses. Refer to section 2.6.2 in the Fortran specifications, and section 2.7.2 in the C/C++ specifications for complete details.*

TABLE 1-1 identifies the directives on which these clauses can appear.

#### 1.10.1.1 PRIVATE Clause

`private(list)`

Declares the variables in the optional comma-separated *list* to be private to each thread in a team.

#### 1.10.1.2 SHARED Clause

`shared(list)`

All the threads in the team share the variables that appear in *list*, and access the same storage area.

#### 1.10.1.3 DEFAULT Clause

**Fortran**

`DEFAULT(PRIVATE | SHARED | NONE)`



C/C++

`default(shared | none)`

Specify scoping attribute for all variables within a parallel region. **THREADPRIVATE** variables are not affected by this clause. If not specified, **DEFAULT(SHARED)** is assumed. A variable's default data-sharing attribute can be overridden by using the **private**, **firstprivate**, **lastprivate**, **reduction**, and **shared** clauses.

#### 1.10.1.4 **FIRSTPRIVATE** Clause

`firstprivate(list)`

The variables in *list* are **PRIVATE**. In addition, private copies of the variables are initialized from the original object existing before the construct.

#### 1.10.1.5 **LASTPRIVATE** Clause

`lastprivate(list)`

The variables in the *list* are **PRIVATE**. In addition, when the **LASTPRIVATE** clause appears on a **DO** or **for** directive, the thread that executes the sequentially last iteration updates the version of the variable before the construct. On a **SECTIONS** directive, the thread that executes the lexically last **SECTION** updates the version of the object it had before the construct.

#### 1.10.1.6 **COPYIN** Clause

**Fortran**

`COPYIN(list)`

The **COPYIN** clause applies only to variables, common blocks, and variables in common blocks that are declared as **THREADPRIVATE**. In a parallel region, **COPYIN** specifies that the data in the master thread of the team be copied to the threadprivate copies of the common block at the beginning of the parallel region.

C/C++

`copyin(list)`

The **COPYIN** clause applies only to variables that are declared as **THREADPRIVATE**. In a parallel region, **COPYIN** specifies that the data in the master thread of the team be copied to the threadprivate copies at the beginning of the parallel region.

## 1.10.1.7 COPYPRIVATE Clause

### Fortran

`COPYPRIVATE`(*list*)

Uses a private variable to broadcast a value, or a pointer to a shared object, from one member of a team to the other members. `COPYPRIVATE` clause can only appear on the `END SINGLE` directive. The broadcast occurs after the execution of the structured block associated with the `single` construct, and before any threads in the team have left the barrier at the end of the construct. The variables in *list* must not appear in a `PRIVATE` or `FIRSTPRIVATE` clause of the `SINGLE` construct specifying `COPYPRIVATE`.

### C/C++

`copyprivate`(*list*)

Uses a private variable to broadcast a value from one member of a team to the other members. The `copyprivate` clause can only appear on the `single` directive. The broadcast occurs after the execution of the structured block associated with the `single` construct, and before any threads in the team have left the barrier at the end of the construct. The variables in *list* must not appear in a `private` or `firstprivate` clause for the same `single` directive.

## 1.10.1.8 REDUCTION Clause

### Fortran

`REDUCTION`(*operator* | *intrinsic*:*list*)

*operator* is one of: `+`, `*`, `-`, `.AND.`, `.OR.`, `.EQV.`, `.NEQV.`

*intrinsic* is one of: `MAX`, `MIN`, `IAND`, `IOR`, `IEOR`

Variables in *list* must be named variables of intrinsic type.

### C/C++

`reduction`(*operator*:*list*)

*operator* is one of: `+`, `*`, `-`, `&`, `^`, `|`, `&&`, `||`

The `REDUCTION` clause is intended to be used on a region in which the reduction variable is used only in reduction statements. The variables in *list* must be `SHARED` in the enclosing context. A private copy of each variable is created for each thread as if it were `PRIVATE`. At the end of the reduction, the shared variable is updated by combining the original value with the final value of each of the private copies.

See section 2.6.2.6 in the Fortran OpenMP specifications, and section 2.7.2.6 in the C/C++ specifications for complete details and restrictions on `REDUCTION` clauses and constructs.

## 1.10.2 Scheduling Clauses

The **SCHEDULE** clause specifies how iterations in a Fortran **DO** loop or C/C++ **for** loop are divided among the threads in a team. TABLE 1-1 shows which directives allow the **SCHEDULE** clause.

There are important restrictions on the use of these scheduling clauses. Refer to section 2.3.1 in the Fortran specification, and section 2.4.1 in the C/C++ specification for complete details.

```
schedule( type [, chunk])
```

Specifies how iterations of the **DO** or **for** loop are divided among the threads of the team. *type* can be one of **STATIC**, **DYNAMIC**, **GUIDED**, or **RUNTIME**. In the absence of a **SCHEDULE** clause, Sun ONE Studio compilers use **STATIC** scheduling. *chunk* must be an integer expression.

### 1.10.2.1 **STATIC** Scheduling

```
schedule(static[, chunk])
```

Iterations are divided into pieces of a size specified by *chunk*. The pieces are statically assigned to threads in the team in a round-robin fashion in the order of the thread number. If not specified, *chunk* is chosen so that the iterations divide into contiguous chunks nearly equal in size with one chunk assigned to each thread.

### 1.10.2.2 **DYNAMIC** Scheduling

```
schedule(dynamic[, chunk])
```

Iterations are divided into pieces of a size specified by *chunk*, and assigned to a waiting thread. As each thread finishes its piece of the iteration space, it dynamically obtains the next set of iterations. When no *chunk* is specified, it defaults to 1.

### 1.10.2.3 **GUIDED** Scheduling

```
schedule(guided[, chunk])
```

With **GUIDED**, the chunk size is reduced in an exponentially decreasing manner with each dispatched piece of the iterations. *chunk* specifies the minimum number of iterations to dispatch each time. (The size of the initial chunk of the iterations is implementation dependent; see Chapter 2.). When no *chunk* is specified, it defaults to 1.

## 1.10.2.4 **RUNTIME** Scheduling

`schedule(runtime)`

Scheduling is deferred until runtime. Schedule *type* and *chunk* size will be determined from the value of the `OMP_SCHEDULE` environment variable. (Default is `SCHEDULE(STATIC)`).

## 1.10.3 **NUM\_THREADS** Clause

The OpenMP API provides a `NUM_THREADS` clause on the `PARALLEL`, `PARALLEL SECTIONS`, `PARALLEL DO`, `PARALLEL for`, and `PARALLEL WORKSHARE` directives.

`num_threads(scalar_integer_expression)`

Specifies the number of threads in the team created when a thread enters a parallel region. *scalar\_integer\_expression* is the number of threads requested, and supersedes the number of threads defined by a prior call to the `OMP_SET_NUM_THREADS` library function, or the value of the `OMP_NUM_THREADS` environment variable. If dynamic thread management is enabled, the request is the *maximum* number of threads to use.

Note that `num_threads` does not apply to subsequent regions.

## 1.10.4 Placement of Clauses on Directives

TABLE 1-1 shows the clauses that can appear on these directives and pragmas:

- `PARALLEL`
- `DO`
- `for`
- `SECTIONS`
- `SINGLE`
- `PARALLEL DO`
- `parallel for`
- `PARALLEL SECTIONS`
- `PARALLEL WORKSHARE`

TABLE 1-1 Pragma Where Clauses Can Appear

Clause/Pragma	PARALLEL	DO/for	SECTIONS	SINGLE	PARALLEL DO/for	PARALLEL SECTIONS	PARALLEL WORKSHARE <sup>3</sup>
IF	•				•	•	•
PRIVATE	•	•	•	•	•	•	•
SHARED	•				•	•	•
FIRSTPRIVATE	•	•	•	•	•	•	•
LASTPRIVATE		•	•		•	•	
DEFAULT	•				•	•	•
REDUCTION	•	•	•		•	•	•
COPYIN	•				•	•	•
COPYPRIVATE				• <sup>1</sup>			
ORDERED		•			•		
SCHEDULE		•			•		
NOWAIT		• <sup>2</sup>	• <sup>2</sup>	• <sup>2</sup>			
NUM_THREADS	•				•	•	•

1. Fortran only: **COPYPRIVATE** can appear on the **END SINGLE** directive.
2. For Fortran, a **NOWAIT** modifier can appear on the **END DO**, **END SECTIONS**, **END SINGLE**, or **END WORKSHARE** directives.
3. Only Fortran supports **WORKSHARE** and **PARALLEL WORKSHARE**.

---

## 1.11 OpenMP Runtime Library Routines

OpenMP provides a set of callable library routines to control and query the parallel execution environment, a set of general purpose lock routines, and two portable timer routines. Full details appear in the Fortran and C/C++ OpenMP specifications.

### 1.11.1 Fortran OpenMP Routines

The Fortran run-time library routines are external procedures. In the following summary, *int\_expr* is a scalar integer expression, and *logical\_expr* is a scalar logical expression.

The `OMP_` functions returning `INTEGER(4)` and `LOGICAL(4)` are not intrinsic and must be declared properly, otherwise the compiler will assume `REAL`. Interface declarations for the OpenMP Fortran runtime library routines summarized below are provided by the Fortran include file `omp_lib.h` and a Fortran `MODULE omp_lib`, as described in the Fortran OpenMP specifications.

Supply an `INCLUDE 'omp_lib.h'` statement or `#include "omp_lib.h"` preprocessor directive, or a `USE omp_lib` statement in every program unit that references these library routines.

Compiling with `-xlist` will report any type mismatches.

The integer parameter `omp_lock_kind` defines the `KIND` type parameters used for simple lock variables in the `OMP_*_LOCK` routines.

The integer parameter `omp_nest_lock_kind` defines the `KIND` type parameters used for the nestable lock variables in the `OMP_*_NEST_LOCK` routines.

The integer parameter `openmp_version` is defined as a preprocessor macro `_OPENMP` having the form `YYYYMM` where `YYYY` and `MM` are the year and month designations of the version of the OpenMP Fortran API.

### 1.11.2 C/C++ OpenMP Routines

The C/C++ run-time library functions are external functions.

The header `<omp.h>` declares two types, several functions that can be used to control and query the parallel execution environment, and lock functions that can be used to synchronize access to data.

The type `omp_lock_t` is an object type capable of representing that a lock is available, or that a thread owns a lock. These locks are referred to as simple locks.

The type `omp_nest_lock_t` is an object type capable of representing that a lock is available, or that a thread owns a lock. These locks are referred to as nestable locks.

## 1.11.3 Run-time Thread Management Routines

For details, refer to the appropriate OpenMP specifications.

### 1.11.3.1 `OMP_SET_NUM_THREADS` Routine

Sets the number of threads to use for subsequent parallel regions

**Fortran**

```
SUBROUTINE OMP_SET_NUM_THREADS(int_expr)
```

**C/C++**

```
#include <omp.h>  
void omp_set_num_threads(int num_threads);
```

### 1.11.3.2 `OMP_GET_NUM_THREADS` Routine

Returns the number of threads currently in the team executing the parallel region from which it is called.

**Fortran**

```
INTEGER(4) FUNCTION OMP_GET_NUM_THREADS()
```

**C/C++**

```
#include <omp.h>  
int omp_get_num_threads(void);
```

### 1.11.3.3 `OMP_GET_MAX_THREADS` Routine

Returns maximum value that can be returned by calls to the `OMP_GET_NUM_THREADS` function.

**Fortran**

```
INTEGER(4) FUNCTION OMP_GET_MAX_THREADS()
```

**C/C++**

```
#include <omp.h>
int omp_get_max_threads(void);
```

### 1.11.3.4 **OMP\_GET\_THREAD\_NUM** Routine

Returns the thread number, within its team, of the thread executing the call to this function. This number lies between 0 and `OMP_GET_NUM_THREADS()-1`, with 0 being the master thread.

**Fortran**

```
INTEGER(4) FUNCTION OMP_GET_THREAD_NUM()
```

**C/C++**

```
#include <omp.h>
int omp_get_thread_num(void);
```

### 1.11.3.5 **OMP\_GET\_NUM\_PROCS** Routine

Return the number of processors available to the program.

**Fortran**

```
INTEGER(4) FUNCTION OMP_GET_NUM_PROCS()
```

**C/C++**

```
#include <omp.h>
int omp_get_num_procs(void);
```

### 1.11.3.6 **OMP\_IN\_PARALLEL** Routine

Determine whether or not thread is executing within the dynamic extent of a parallel region.

**Fortran**

```
LOGICAL(4) FUNCTION OMP_IN_PARALLEL()
```

Returns `.TRUE.` if called within the dynamic extent of a parallel region, `.FALSE.` otherwise.

**C/C++**

```
#include <omp.h>
int omp_in_parallel(void);
```



Returns nonzero if called within the dynamic extent of a parallel region, zero otherwise.

### 1.11.3.7 **OMP\_SET\_DYNAMIC** Routine

Enables or disables dynamic adjustment of the number of available threads. (Dynamic adjustment is enabled by default.)

#### **Fortran**

```
SUBROUTINE OMP_SET_DYNAMIC(logical_expr)
```

Dynamic adjustment is enabled when *logical\_expr* evaluates to `.TRUE.`, and is disabled otherwise.

#### **C/C++**

```
#include <omp.h>  
void omp_set_dynamic(int dynamic);
```

If *dynamic* evaluates as nonzero, dynamic adjustment is enabled; otherwise it is disabled.

### 1.11.3.8 **OMP\_GET\_DYNAMIC** Routine

Determine whether or not dynamic thread adjustment is enabled.

#### **Fortran**

```
LOGICAL(4) FUNCTION OMP_GET_DYNAMIC()
```

Returns `.TRUE.` if dynamic thread adjustment is enabled, `.FALSE.` otherwise.

#### **C/C++**

```
#include <omp.h>  
int omp_get_dynamic(void);
```

Returns nonzero if dynamic thread adjustment is enabled, zero otherwise.

### 1.11.3.9 **OMP\_SET\_NESTED** Routine

Enables or disables nested parallelism. (*Nested parallelism is not supported, and is disabled by default.*)

#### **Fortran**

```
SUBROUTINE OMP_SET_NESTED(logical_expr)
```

#### **C/C++**

```
#include <omp.h>
void omp_set_nested(int nested);
```

### 1.11.3.10 OMP\_GET\_NESTED Routine

Determine whether or not nested parallelism is enabled. (*Nested parallelism is not supported, and is disabled by default.*)

#### Fortran

```
LOGICAL(4) FUNCTION OMP_GET_NESTED()
```

Returns **.FALSE.** *Nested parallelism is not supported.*

#### C/C++

```
#include <omp.h>
int omp_get_nested(void);
```

Returns zero. *Nested parallelism is not supported.*

## 1.11.4 Routines That Manage Synchronization Locks

Two types of locks are supported: simple locks and nestable locks. Nestable locks may be locked multiple times by the same thread before being unlocked; simple locks may not be locked if they are already in a locked state. Simple lock variables may only be passed to simple lock routines, and nested lock variables only to nested lock routines.

#### Fortran:

The lock variable *var* must be accessed only through these routines. Use the parameters `OMP_LOCK_KIND` and `OMP_NEST_LOCK_KIND` (defined in `omp_lib.h INCLUDE` file and the `omp_lib MODULE`) for this purpose. For example,

```
INTEGER(KIND=OMP_LOCK_KIND)    :: var
INTEGER(KIND=OMP_NEST_LOCK_KIND) :: nvar
```

#### C/C++:

Simple lock variables must have type `omp_lock_t` and must be accessed only through these functions. All simple lock functions require an argument that points to `omp_lock_t` type.

Nested lock variables must have type `omp_nest_lock_t`, and similarly all nested lock functions require an argument that points to `omp_nest_lock_t` type.

### 1.11.4.1 **OMP\_INIT\_LOCK** and **OMP\_INIT\_NEST\_LOCK** Routines

Initialize a lock variable for subsequent calls.

#### **Fortran**

```
SUBROUTINE OMP_INIT_LOCK(var)  
SUBROUTINE OMP_INIT_NEST_LOCK(nvar)
```

#### **C/C++**

```
#include <omp.h>  
void omp_init_lock(omp_lock_t *lock);  
void omp_init_nest_lock(omp_nest_lock_t *lock);
```

### 1.11.4.2 **OMP\_DESTROY\_LOCK** and **OMP\_DESTROY\_NEST\_LOCK** Routines

Removes a lock variable.

#### **Fortran**

```
SUBROUTINE OMP_DESTROY_LOCK(var)  
SUBROUTINE OMP_DESTROY_NEST_LOCK(nvar)
```

#### **C/C++**

```
#include <omp.h>  
void omp_destroy_lock(omp_lock_t *lock);  
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

### 1.11.4.3 **OMP\_SET\_LOCK** and **OMP\_SET\_NEST\_LOCK** Routines

Forces the executing thread to wait until the specified lock is available. The thread is granted ownership of the lock when it is available.

#### **Fortran**

```
SUBROUTINE OMP_SET_LOCK(var)  
SUBROUTINE OMP_SET_NEST_LOCK(nvar)
```

#### **C/C++**

```
#include <omp.h>  
void omp_set_lock(omp_lock_t *lock);  
void omp_set_nest_lock(omp_nest_lock_t *lock);
```

#### 1.11.4.4 **OMP\_UNSET\_LOCK** and **OMP\_UNSET\_NEST\_LOCK** Routines

Releases the executing thread from ownership of the lock. Behavior is undefined if the thread does not own that lock.

##### **Fortran**

```
SUBROUTINE OMP_UNSET_LOCK(var)
SUBROUTINE OMP_UNSET_NEST_LOCK(nvar)
```

##### **C/C++**

```
#include <omp.h>
void omp_unset_lock(omp_lock_t *lock);
void omp_unset_nest_lock(omp_nest_lock_t *lock);
```

#### 1.11.4.5 **OMP\_TEST\_LOCK** and **OMP\_TEST\_NEST\_LOCK** Routines

**OMP\_TEST\_LOCK** attempts to set the lock associated with lock variable. Call does not block execution of the thread.

**OMP\_TEST\_NEST\_LOCK** returns the new nesting count if the lock was set successfully, otherwise it returns 0. Call does not block execution of the thread.

##### **Fortran**

```
LOGICAL(4) FUNCTION OMP_TEST_LOCK(var)
Returns .TRUE. if the lock was set, .FALSE. otherwise.
INTEGER(4) FUNCTION OMP_TEST_NEST_LOCK(nvar)
Returns nesting count if lock was set successfully, zero otherwise.
```

##### **C/C++**

```
#include <omp.h>
int omp_test_lock(omp_lock_t *lock);
Returns a nonzero value if lock was set successfully, zero otherwise.

int omp_test_nest_lock(omp_nest_lock_t *lock);
Returns lock nest count if lock was set successfully, zero otherwise.
```

### 1.11.5 **Timing Routines**

Two functions support a portable wall clock timer.

### 1.11.5.1 **OMP\_GET\_WTIME** Routine

Returns the elapsed wall clock time in seconds “since some arbitrary time in the past”.

#### **Fortran**

```
REAL(8) FUNCTION OMP_GET_WTIME()
```

#### **C/C++**

```
#include <omp.h>  
double omp_get_wtime(void);
```

### 1.11.5.2 **OMP\_GET\_WTICK** Routine

Returns the number of seconds between successive clock ticks.

#### **Fortran**

```
REAL(8) FUNCTION OMP_GET_WTICK()
```

#### **C/C++**

```
#include <omp.h>  
double omp_get_wtick(void);
```



# Implementation-Dependent Issues

---

This chapter notes specific issues in the OpenMP 2.0 Fortran and C/C++ specifications that are implementation dependent. For last-minute information regarding the latest compiler releases, see the C, C++, and Fortran readme files.

## Scheduling

- The default, in the absence of an explicit `OMP_SCHEDULE` environment variable, or an explicit `SCHEDULE` clause, is `static` scheduling.

## Number of Threads

- Without an explicit `num_threads()` clause, call to the `omp_set_num_threads()` function, or an explicit definition of the `OMP_NUM_THREADS` environment variable, the default number of threads in a team is 1.

## Dynamic Threads

- Without an explicit call to the `omp_set_dynamic()` function, or an explicit definition of the `OMP_DYNAMIC` environment variable, the default is to enable dynamic thread adjustment. When dynamic thread adjustment is enabled, the number of threads is limited to the number of available processors.

## Nested Parallelism

- Nested parallelism is not supported in this implementation, and is disabled by default. Nested parallel regions are executed by a single thread only.

## ATOMIC Directive

- This implementation replaces all `ATOMIC` directives and pragmas by enclosing the target statement in a critical region.

### **GUIDED Initial and Minimum Chunk Size**

- The default minimum chunk size with `SCHEDULE(GUIDED, chunk)` is 1. The default initial chunk size is the number of iterations in the loop divided by the number of threads executing the loop.

### **Explicitly Threaded Programs**

- Programs using threads can call routines that contain OpenMP directives.



# Compiling for OpenMP

---

This chapter describes how to compile programs that utilize the OpenMP API.

To run a parallelized program in a multithreaded environment, you must set the `OMP_NUM_THREADS` environment variable prior to program execution. This tells the runtime system the maximum number of threads the program can create. The default is 1. In general, set `OMP_NUM_THREADS` to the available number of processors on the target platform.

The compiler readme files contain information about limitations and known deficiencies regarding their OpenMP implementation. Readme files are viewable directly by invoking the compiler with the `-xhelp=readme` flag, or by pointing an HTML browser to the documentation index for the installed software at

```
file:/opt/SUNWspro/docs/index.html
```

---

## 3.1 Compiler Options To Use

To enable explicit parallelization with OpenMP directives, compile your program with the `cc`, `CC`, or `£95` option flag `-xopenmp`. This flag can take an optional keyword argument. (The `£95` compiler accepts both `-xopenmp` and `-openmp` as synonyms.)

The `-xopenmp` flag accepts the following keyword sub-options.

---

<code>-xopenmp=parallel</code>	Enables recognition of OpenMP pragmas. The minimum optimization level for <code>-xopenmp=parallel</code> is <code>-xO3</code> . The compiler changes the optimization from a lower level to <code>-xO3</code> if necessary, and issues a warning.
<code>-xopenmp=noopt</code>	Enables recognition of OpenMP pragmas. The compiler does not raise the level if it is lower than <code>-xO3</code> . If you explicitly set the optimization lower than <code>-xO3</code> , as in <code>-xO2 -openmp=noopt</code> the compiler will issue an error. If you do not specify an optimization level with <code>-openmp=noopt</code> , the OpenMP pragmas are recognized, the program is parallelized accordingly, but no optimization is done. (This sub-option applies to <code>cc</code> and <code>f95</code> only; <code>CC</code> issues a warning if specified, and no OpenMP parallelization is done.)
<code>-xopenmp=stubs</code>	Disables recognition of OpenMP pragmas, links to stub library routines, and does not change the optimization level. Use this option if your application makes explicit calls to the OpenMP runtime library routines and you want to compile it to execute serially.
<code>-xopenmp=none</code>	Disables recognition of OpenMP pragmas and does not change the optimization level. (Default)

---

If you do not specify `-xopenmp` on the command line, the compiler assumes `-xopenmp=none` (disabling recognition of OpenMP pragmas).

If you specify `-xopenmp` but without a keyword sub-option, the compiler assumes `-xopenmp=parallel`.

Do not specify `-xopenmp` together with `-xparallel` or `-xexplicitpar` on the command line.

Specifying `-xopenmp=` with `parallel`, `noopt`, or `stubs` will define the `_OPENMP` preprocessor token to be `YYYYMM` (specifically `200203L` for C/C++ and `200011` for Fortran 95).

When debugging OpenMP programs with `dbx`, compile with `-xopenmp=noopt -g`

The default optimization level for `-xopenmp` might change in future releases. Warning messages can be avoided by specifying an appropriate optimization level explicitly.

With Fortran 95, `-xopenmp`, `-xopenmp=parallel`, `-xopenmp=noopt` will add `-stackvar` automatically.

---

## 3.2 Fortran 95 OpenMP Validation

You can obtain a static, interprocedural validation of a Fortran 95 program's OpenMP directives by using the `f95` compiler's global program checking feature. Enable OpenMP checking by compiling with the `-xlistMP` flag. (Diagnostic messages from `-xlistMP` appear in a separate file created with the name of the source file and a `.lst` extension). The compiler will diagnose the following violations and parallelization inhibitors:

- Violations in the specifications of parallel directives, including improper nesting.
- Parallelization inhibitors due to data usage, detected by interprocedural dependence analysis.
- Parallelization inhibitors detected by interprocedural pointer analysis.

For example, compiling a source file `ord.f` with `-xlistMP` produces a diagnostic file `ord.lst`:

```
FILE "ord.f"
 1  !$OMP PARALLEL
 2  !$OMP DO ORDERED
 3              do i=1,100
 4                  call work(i)
 5              end do
 6  !$OMP END DO
 7  !$OMP END PARALLEL
 8
 9  !$OMP PARALLEL
10  !$OMP DO
11              do i=1,100
12                  call work(i)
13              end do
14  !$OMP END DO
15  !$OMP END PARALLEL
16              end
17              subroutine work(k)
18  !$OMP ORDERED
19              ^
          write(*,*) k
20  !$OMP END ORDERED
21              return
22              end

**** ERR-OMP: It is illegal for an ORDERED directive to bind to a
directive (ord.f, line 10, column 2) that does not have the
ORDERED clause specified.
```

In this example, the `ORDERED` directive in subroutine `WORK` receives a diagnostic that refers to the second `DO` directive because it lacks an `ORDERED` clause.

## 3.3 OpenMP Environment Variables

The OpenMP specifications define four environment variables that control the execution of OpenMP programs. These are summarized in the following table.

TABLE 3-1 OpenMP Environment Variables

Environment Variable	Function
<code>OMP_SCHEDULE</code>	Sets schedule type for <code>DO</code> , <code>PARALLEL DO</code> , <code>parallel for</code> , <code>for</code> , directives/pragmas with schedule type <code>RUNTIME</code> specified. If not defined, a default value of <code>STATIC</code> is used. <i>value</i> is “ <i>type[,chunk]</i> ” Example: <code>setenv OMP_SCHEDULE "GUIDED,4"</code>
<code>OMP_NUM_THREADS</code> or <code>PARALLEL</code>	Sets the number of threads to use during execution, unless set by a <code>NUM_THREADS</code> clause, or a call to <code>OMP_SET_NUM_THREADS()</code> . If not set, a default of 1 is used. <i>value</i> is a positive integer. ( <i>Current maximum is 128</i> ). For compatibility with legacy programs, setting the <code>PARALLEL</code> environment variable has the same effect as setting <code>OMP_NUM_THREADS</code> . However, if they are both set to different values, the runtime library will issue an error message. Example: <code>setenv OMP_NUM_THREADS 16</code>
<code>OMP_DYNAMIC</code>	Enables or disables dynamic adjustment of the number of threads available for execution of parallel regions. If not set, a default value of <code>TRUE</code> is used. <i>value</i> is either <code>TRUE</code> or <code>FALSE</code> . Example: <code>setenv OMP_DYNAMIC FALSE</code>
<code>OMP_NESTED</code>	Enables or disables nested parallelism. ( <i>Nested parallelism is not supported</i> ). <i>value</i> is either <code>TRUE</code> or <code>FALSE</code> . ( <i>This variable has no effect.</i> ) Example: <code>setenv OMP_NESTED FALSE</code>

Additional multiprocessing environment variables affect execution of OpenMP programs and are not part of the OpenMP specifications. These are summarized in the following table.

TABLE 3-2 Multiprocessing Environment Variables

Environment Variable	Function
<code>SUNW_MP_WARN</code>	<p>Controls warning messages issued by the OpenMP runtime library. If set <b>TRUE</b> the runtime library issues warning messages to <code>stderr</code>; <b>FALSE</b> disables warning messages. The default is <b>FALSE</b>.</p> <p>Example:</p> <pre>setenv SUNW_MP_WARN FALSE</pre>
<code>SUNW_MP_THR_IDLE</code>	<p>Controls the end-of-task status of each helper thread executing the parallel part of a program. You can set the value to <b>spin</b>, <b>sleep ns</b>, or <b>sleep nms</b>. The default is <b>SPIN</b> — the thread spins (or busy-waits) after completing a parallel task until a new parallel task arrives.</p> <p>Choosing <b>SLEEP time</b> specifies the amount of time a helper thread should spin-wait after completing a parallel task. If, while a thread is spinning, a new task arrives for the thread, the thread executes the new task immediately. Otherwise, the thread goes to sleep and is awakened when a new task arrives. <i>time</i> may be specified in seconds, (<i>ns</i>) or just (<i>n</i>), or milliseconds, (<i>nms</i>).</p> <p><b>SLEEP</b> with no argument puts the thread to sleep immediately after completing a parallel task. <b>SLEEP</b>, <b>SLEEP (0)</b>, <b>SLEEP (0s)</b>, and <b>SLEEP (0ms)</b> are all equivalent.</p> <p>Example: <pre>setenv SUNW_MP_THR_IDLE SLEEP(50ms)</pre></p>
<code>STACKSIZE</code>	<p>Sets the stack size for each thread. The value is in kilobytes. The default thread stack sizes are 4 Mb on 32-bit SPARC V8 platforms, and 8 Mb on 64-bit SPARC V9 platforms.</p> <p>Example:</p> <pre>setenv STACKSIZE 8192 sets the thread stack size to 8 Mb</pre>

## 3.4 Stacks and Stack Sizes

The executing program maintains a main memory stack for the initial thread executing the program, as well as distinct stacks for each helper thread. Stacks are temporary memory address spaces used to hold arguments and automatic variables over subprogram or function references.

In general, the default main stack size is about 8 megabytes. Compiling Fortran programs with the `f95 -stackvar` option forces the allocation of local variables and arrays on the stack as if they were automatic variables. Use of `-stackvar` with OpenMP programs is implied with explicitly parallelized programs because it

improves the optimizer's ability to parallelize calls in loops. (See the *Fortran User's Guide* for a discussion of the `-stackvar` flag.) However, this may lead to stack overflow if not enough memory is allocated for the stack.

Use the `limit` C-shell command, or the `ulimit` ksh/sh command, to display or set the size of the main stack.

Each helper thread of a multithreaded program has its own thread stack. This stack mimics the initial (or main) thread stack but is unique to the thread. The thread's **PRIVATE** arrays and variables (local to the thread) are allocated on the thread stack. The default size is 4 megabytes on 32-bit systems and 8 megabytes on 64-bit systems. The size of the helper thread stack is set with the **STACKSIZE** environment variable.

```
demo% setenv STACKSIZE 16384 <-Set thread stack size to 16 Mb (C shell)

demo% STACKSIZE=16384 <-Same, using Bourne/Korn shell
demo% export STACKSIZE
```

Finding the best stack size might have to be determined by trial and error. If the stack size is too small for a thread to run it may cause silent data corruption in neighboring threads, or segmentation faults. If you are unsure about stack overflows, compile your Fortran or C programs with the `-xcheck=stkovf` flag to force a segmentation fault on stack overflow. This stops the program before any data corruption can occur.

# Converting to OpenMP

---

This chapter gives guidelines for converting legacy programs using Sun or Cray directives and pragmas to OpenMP.

---

## 4.1 Converting Legacy Fortran Directives

Legacy Fortran programs use either Sun or Cray style parallelization directives. A description of these directives can be found in the chapter *Parallelization* in the *Fortran Programming Guide*.

## 4.1.1 Converting Sun-Style Fortran Directives

The following tables give OpenMP near equivalents to Sun parallelization directives and their subclauses. These are only suggestions.

**TABLE 4-1** Converting Sun Parallelization Directives to OpenMP

Sun Directive	Equivalent OpenMP Directive
C\$PAR DOALL [ <i>qualifiers</i> ]	!\$omp parallel do [ <i>qualifiers</i> ]
C\$PAR DOSERIAL	No exact equivalent. You can use: !\$omp master <i>loop</i> !\$omp end master
C\$PAR DOSERIAL*	No exact equivalent. You can use: !\$omp master <i>looptest</i> !\$omp end master
C\$PAR TASKCOMMON <i>block</i> [...]	!\$omp threadprivate (/block/[...])

The DOALL directive can take the following optional qualifier clauses.

**TABLE 4-2** DOALL Qualifier Clauses and OpenMP Equivalent Clauses

Sun DOALL Clause	OpenMP PARALLEL DO Equivalent Clauses
PRIVATE( <i>v1,v2,...</i> )	private( <i>v1,v2,...</i> )
SHARED( <i>v1,v2,...</i> )	shared( <i>v1,v2,...</i> )
MAXCPUS( <i>n</i> )	num_threads( <i>n</i> ). No exact equivalent.
READONLY( <i>v1,v2,...</i> )	No exact equivalent. You can achieve the same effect by using firstprivate( <i>v1,v2,...</i> ).
STOREBACK( <i>v1,v2,...</i> )	No exact equivalent. You can achieve the same effect by using lastprivate( <i>v1,v2,...</i> ).
SAVELAST	No exact equivalent. You can achieve the same effect by using lastprivate( <i>v1,v2,...</i> ).
REDUCTION( <i>v1,v2,...</i> )	reduction(operator: <i>v1,v2,...</i> ) Must supply the reduction operator as well as the list of variables.
SCHEDTYPE( <i>spec</i> )	schedule( <i>spec</i> ) (See TABLE 4-3)



The SCHEDTYPE(*spec*) clause accepts the following scheduling specifications.

**TABLE 4-3** SCHEDTYPE Scheduling and OpenMP schedule Equivalents

SCHEDTYPE( <i>spec</i> )	OpenMP schedule( <i>spec</i> ) Clause Equivalent
SCHEDTYPE(STATIC)	schedule(static)
SCHEDTYPE(SELF( <i>chunksize</i> ))	schedule(dynamic, <i>chunksize</i> ) Default <i>chunksize</i> is 1.
SCHEDTYPE(FACTORING( <i>m</i> ))	No OpenMP equivalent.
SCHEDTYPE(GSS( <i>m</i> ))	schedule(guided, <i>m</i> ) Default <i>m</i> is 1.

### 4.1.1.1 Issues Between Sun-Style Fortran Directives and OpenMP

- Scoping of variables (shared or private) must be declared explicitly with OpenMP. With Sun directives, the compiler uses its own default scoping rules for variables not explicitly scoped in a PRIVATE or SHARED clause: all scalars are treated as PRIVATE, and all array references are SHARED. With OpenMP, the default data scope is SHARED unless a DEFAULT(PRIVATE) clause appears on the PARALLEL DO directive. A DEFAULT(NONE) clause causes the compiler to flag variables not scoped explicitly.
- Since there is no DOSERIAL directive, mixing automatic and explicit OpenMP parallelization may have different effects: some loops may be automatically parallelized that would not have been with Sun directives.
- OpenMP provides a richer parallelism model by providing parallel regions and parallel sections. It could be possible to get better performance by redesigning the parallelism strategies of a program using Sun directives to take advantage of these features of OpenMP.

## 4.1.2 Converting Cray-Style Fortran Directives

Cray-style Fortran parallelization directives are identical to Sun-style except that the sentinel that identifies these directives is `!MIC$`. Also, the set of qualifier clauses on the `!MIC$ DOALL` is different.

**TABLE 4-4** OpenMP Equivalents for Cray-Style `DOALL` Qualifier Clauses

Cray <code>DOALL</code> Clause	OpenMP <code>PARALLEL DO</code> Equivalent Clauses
<code>SHARED(v1,v2,...)</code>	<code>SHARED(v1,v2,...)</code>
<code>PRIVATE(v1,v2,...)</code>	<code>PRIVATE(v1,v2,...)</code>
<code>AUTOSCOPE</code>	No equivalent. Scoping must be explicit, or with the <code>DEFAULT</code> clause.
<code>SAVELAST</code>	No exact equivalent. You can achieve the same effect by using <code>lastprivate</code> .
<code>MAXCPUS(n)</code>	<code>num_threads(n)</code> . No exact equivalent.
<code>GUIDED</code>	<code>schedule(guided, m)</code> Default <code>m</code> is 1.
<code>SINGLE</code>	<code>schedule(dynamic, 1)</code>
<code>CHUNKSIZE(n)</code>	<code>schedule(dynamic, n)</code>
<code>NUMCHUNKS(m)</code>	<code>schedule(dynamic, n/m)</code> where <code>n</code> is the number of iterations

### 4.1.2.1 Issues Between Cray-Style Fortran Directives and OpenMP Directives

The differences are the same as for Sun-style directives, except that there is no equivalent for the Cray `AUTOSCOPE`.

---

## 4.2 Converting Legacy C Pragmas

The C compiler accepts legacy pragmas for explicit parallelization. These are described in the *C User's Guide*. As with the Fortran directives, these are only suggestions.

The legacy parallelization pragmas are:

**TABLE 4-5** Converting Legacy C Parallelization Pragmas to OpenMP

Legacy C Pragma	Equivalent OpenMP Pragma
<code>#pragma MP taskloop [clauses]</code>	<code>#pragma omp parallel for [clauses]</code>
<code>#pragma MP serial_loop</code>	No exact equivalent. You can use <code>#pragma omp master</code> <i>loop</i>
<code>#pragma MP serial_loop_nested</code>	No exact equivalent. You can use <code>#pragma omp master</code> <i>loopnest</i>

The `taskloop` pragma can take on one or more of the following optional clauses.

**TABLE 4-6** `taskloop` Optional Clauses and OpenMP Equivalents

<code>taskloop</code> Clause	OpenMP <code>parallel for</code> Equivalent Clause
<code>maxcpus(n)</code>	No equivalent. Use <code>num_threads(n)</code>
<code>private(v1,v2,...)</code>	<code>private(v1,v2,...)</code>
<code>shared(v1,v2,...)</code>	<code>shared(v1,v2,...)</code>
<code>readonly(v1,v2,...)</code>	No exact equivalent. You can achieve the same effect by using <code>firstprivate(v1,v2,...)</code> .
<code>storeback(v1,v2,...)</code>	No exact equivalent. You can achieve the same effect by using <code>lastprivate(v1,v2,...)</code> .
<code>savelast(v1,v2,...)</code>	No exact equivalent. You can achieve the same effect by using <code>lastprivate(v1,v2,...)</code> .
<code>reduction(v1,v2,...)</code>	<code>reduction(operator:v1,v2,...)</code> Must supply the reduction operator as well as the list of variables.
<code>schedtype(spec)</code>	<code>schedule(spec)</code> (See TABLE 4-7)

The `schedtype(spec)` clause accepts the following scheduling specifications.

**TABLE 4-7** SCHEDTYPE Scheduling and OpenMP schedule Equivalents

<b>schedtype(<i>spec</i>)</b>	<b>OpenMP <code>schedule(<i>spec</i>)</code> Clause Equivalent</b>
SCHEDTYPE (STATIC)	<code>schedule(static)</code>
SCHEDTYPE (SELF ( <i>chunksize</i> ))	<code>schedule(dynamic, <i>chunksize</i>)</code> Note: Default <i>chunksize</i> is 1.
SCHEDTYPE (FACTORING ( <i>m</i> ))	No OpenMP equivalent.
SCHEDTYPE (GSS ( <i>m</i> ))	<code>schedule(guided, <i>m</i>)</code> Default <i>m</i> is 1.

## 4.2.1 Issues Between Legacy C Pragmas and OpenMP

- Variables declared within a parallel construct are scoped `private`. A `default(none)` clause on a `#pragma omp parallel for` directive causes the compiler to flag variables not scoped explicitly.
- Since there is no `serial_loop` directive, mixing automatic and explicit OpenMP parallelization may have different effects: some loops may be automatically parallelized that would not have been with legacy C directives.
- Because OpenMP provides a richer parallelism model, it might often be possible to get better performance by redesigning the parallelism strategies of a program that uses legacy C directives to take advantage of these features.

# Index

---

## A

accessible documentation, 1-xiv

## B

barrier, 1-10

## C

C, 3-1

common blocks

in data scoping clauses, 1-16

compilers, accessing, 1-xi

compiling for OpenMP, 3-1

conditional compilation, 1-3

converting to OpenMP

Cray-style Fortran directives, 4-4

legacy C pragmas, 4-4

Sun-style Fortran directives, 4-2

critical region, 1-10

## D

data scoping clauses

**COPYIN**, 1-17

**COPYPRIVATE**, 1-18

**DEFAULT**, 1-17

**FIRSTPRIVATE**, 1-17

**LASTPRIVATE**, 1-17

**PRIVATE**, 1-16

**REDUCTION**, 1-18

**SHARED**, 1-16

directive

formats, 1-2

See pragma

directive clauses

data scoping, 1-16

scheduling, 1-19

directives

**ATOMIC**, 1-12, 2-1

**BARRIER**, 1-11

**CRITICAL**, 1-11

**DO**, 1-5

**FLUSH**, 1-13

**for**, 1-6

**MASTER**, 1-11

**ORDERED**, 1-14

**PARALLEL**, 1-3, 1-4

**PARALLEL DO**, 1-9

**parallel for**, 1-9

**PARALLEL SECTIONS**, 1-9

**PARALLEL WORKSHARE**, 1-10

**SECTION**, 1-6

**SECTIONS**, 1-6

**SINGLE**, 1-7

**THREADPRIVATE**, 1-15

validation (Fortran 95), 3-3

**WORKSHARE**, 1-8

documentation index, 1-xiii

documentation, accessing, 1-xiii to 1-xiv

dynamic thread adjustment, 3-4

dynamic threads, 2-1

## E

environment variables, 3-4  
explicitly threaded programs, 2-2

## H

header files  
    omp.h, 1-22  
    omp\_lib.h, 1-22

## I

idle threads, 3-5  
implementation, 2-1

## M

man pages, accessing, 1-xi  
MANPATH environment variable, setting, 1-xii  
master thread, 1-10

## N

nested parallelism, 2-1, 3-4  
NUM\_THREADS, 1-20  
number of threads, 1-20, 2-1  
    OMP\_NUM\_THREADS, 3-4

## O

omp.h, 1-22  
OMP\_DESTROY\_LOCK(), 1-27  
OMP\_DESTROY\_NEST\_LOCK(), 1-27  
OMP\_DYNAMIC, 3-4  
OMP\_GET\_DYNAMIC(), 1-25  
OMP\_GET\_MAX\_THREADS(), 1-23  
OMP\_GET\_NESTED(), 1-26  
OMP\_GET\_NUM\_PROCS(), 1-24  
OMP\_GET\_NUM\_THREADS(), 1-23  
OMP\_GET\_THREAD\_NUM(), 1-24  
OMP\_GET\_WTICK(), 1-29  
OMP\_GET\_WTIME(), 1-29  
OMP\_IN\_PARALLEL(), 1-24  
OMP\_INIT\_LOCK(), 1-27  
OMP\_INIT\_NEST\_LOCK(), 1-27

omp\_lib.h, 1-22  
OMP\_NESTED, 3-4  
OMP\_NUM\_THREADS, 3-4  
OMP\_SCHEDULE, 3-4  
OMP\_SET\_DYNAMIC(), 1-25  
OMP\_SET\_LOCK(), 1-27  
OMP\_SET\_NEST\_LOCK(), 1-27  
OMP\_SET\_NESTED(), 1-25  
OMP\_SET\_NUM\_THREADS(), 1-23  
OMP\_TEST\_LOCK(), 1-28  
OMP\_TEST\_NEST\_LOCK(), 1-28  
OMP\_UNSET\_LOCK(), 1-28  
OMP\_UNSET\_NEST\_LOCK(), 1-28  
OpenMP 2.0 specifications, 1-1  
ordered region, 1-14

## P

parallel region, 1-3, 1-4  
PATH environment variable, setting, 1-xii  
pragma  
    See directive

## R

run-time  
    C/C++, 1-22  
    Fortran, 1-22

## S

scheduling, 2-1, 2-2  
    OMP\_SCHEDULE, 3-4  
scheduling clauses  
    SCHEDULE, 1-19, 2-1, 2-2  
shell prompts, 1-x  
SLEEP, 3-5  
SPIN, 3-5  
stack size, 3-5  
STACKSIZE, 3-5  
SUNW\_MP\_THR\_IDLE, 3-5  
SUNW\_MP\_WARN, 3-5  
synchronization, 1-10  
synchronization locks, 1-26

## **T**

thread stack size, 3-5  
timing routines, 1-28  
typographic conventions, 1-ix

## **V**

validation of directives (Fortran 95), 3-3

## **W**

warning messages, 3-5  
work-sharing, 1-5  
    combined directives, 1-8

## **X**

-xlistMP, 3-3  
-xopenmp, 3-1

