OpenMP Application Programming Interface

Examples

Version 5.0.1 – June 2020

Source codes for OpenMP 5.0.1 Examples can be downloaded from github.

This page intentionally left blank

Contents

Foreword

The OpenMP Examples document has been updated with new features found in the OpenMP 5.0 Specification. The additional examples and updates are referenced in the Document Revision History of the Appendix on page 424.

Text describing an example with a 5.0 feature specifically states that the feature support begins in the OpenMP 5.0 Specification. Also, an `omp_5.0` keyword has been added to metadata in the source code. These distinctions are presented to remind readers that a 5.0 compliant OpenMP implementation is necessary to use these features in codes.

Examples for most of the 5.0 features are included in this document, and incremental releases will become available as more feature examples and updates are submitted, and approved by the OpenMP Examples Subcommittee.

Examples Subcommitee Co-chairs:

Henry Jin (NASA Ames Research Center)
Kent Milfeld (TACC, Texas Advanced Research Center)

Introduction

This collection of programming examples supplements the OpenMP API for Shared Memory Parallelization specifications, and is not part of the formal specifications. It assumes familiarity with the OpenMP specifications, and shares the typographical conventions used in that document.

The OpenMP API specification provides a model for parallel programming that is portable across shared memory architectures from different vendors. Compilers from numerous vendors support the OpenMP API.

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

The latest source codes for OpenMP Examples can be downloaded from the **`sources`** directory at https://github.com/OpenMP/Examples. The codes for this OpenMP 5.0.1 Examples document have the tag *v5.0.1*.

Complete information about the OpenMP API and a list of the compilers that support the OpenMP API can be found at the OpenMP.org web site

`http://www.openmp.org`

Examples

The following are examples of the OpenMP API directives, constructs, and routines.

C / C++

A statement following a directive is compound only when necessary, and a non-compound statement is indented with respect to a directive preceding it.

C / C++

Each example is labeled as *ename.seqno.ext*, where *ename* is the example name, *seqno* is the sequence number in a section, and *ext* is the source file extension to indicate the code type and source form. *ext* is one of the following:

- *c* – C code,
- *cpp* – C++ code,
- *f* – Fortran code in fixed form, and
- *f90* – Fortran code in free form.

Some of the example labels may include version information (**omp**_*verno*) to indicate features that are illustrated by an example for a specific OpenMP version, such as "*scan.1.c* (**omp_5.0**)."

CHAPTER 1

Parallel Execution

A single thread, the *initial thread*, begins sequential execution of an OpenMP enabled program, as if the whole program is in an implicit parallel region consisting of an implicit task executed by the *initial thread*.

A **parallel** construct encloses code, forming a parallel region. An *initial thread* encountering a **parallel** region forks (creates) a team of threads at the beginning of the **parallel** region, and joins them (removes from execution) at the end of the region. The initial thread becomes the master thread of the team in a **parallel** region with a *thread* number equal to zero, the other threads are numbered from 1 to number of threads minus 1. A team may be comprised of just a single thread.

Each thread of a team is assigned an implicit task consisting of code within the parallel region. The task that creates a parallel region is suspended while the tasks of the team are executed. A thread is tied to its task; that is, only the thread assigned to the task can execute that task. After completion of the **parallel** region, the master thread resumes execution of the generating task.

Any task within a **parallel** region is allowed to encounter another **parallel** region to form a nested **parallel** region. The parallelism of a nested **parallel** region (whether it forks additional threads, or is executed serially by the encountering task) can be controlled by the **OMP_NESTED** environment variable or the **omp_set_nested()** API routine with arguments indicating true or false.

The number of threads of a **parallel** region can be set by the **OMP_NUM_THREADS** environment variable, the **omp_set_num_threads()** routine, or on the **parallel** directive with the **num_threads** clause. The routine overrides the environment variable, and the clause overrides all. Use the **OMP_DYNAMIC** or the **omp_set_dynamic()** function to specify that the OpenMP implementation dynamically adjust the number of threads for **parallel** regions. The default setting for dynamic adjustment is implementation defined. When dynamic adjustment is on and the number of threads is specified, the number of threads becomes an upper limit for the number of threads to be provided by the OpenMP runtime.

WORKSHARING CONSTRUCTS

A worksharing construct distributes the execution of the associated region among the members of the team that encounter it. There is an implied barrier at the end of the worksharing region (there is no barrier at the beginning). The worksharing constructs are:

- loop constructs: **`for`** and **`do`**
- **`sections`**
- **`single`**
- **`workshare`**

The **`for`** and **`do`** constructs (loop constructs) create a region consisting of a loop. A loop controlled by a loop construct is called an *associated* loop. Nested loops can form a single region when the **`collapse`** clause (with an integer argument) designates the number of *associated* loops to be executed in parallel, by forming a "single iteration space" for the specified number of nested loops. The **`ordered`** clause can also control multiple associated loops.

An associated loop must adhere to a "canonical form" (specified in the *Canonical Loop Form* of the OpenMP Specifications document) which allows the iteration count (of all associated loops) to be computed before the (outermost) loop is executed. Most common loops comply with the canonical form, including C++ iterators.

A **`single`** construct forms a region in which only one thread (any one of the team) executes the region. The other threads wait at the implied barrier at the end, unless the **`nowait`** clause is specified.

The **`sections`** construct forms a region that contains one or more structured blocks. Each block of a **`sections`** directive is constructed with a **`section`** construct, and executed once by one of the threads (any one) in the team. (If only one block is formed in the region, the **`section`** construct, which is used to separate blocks, is not required.) The other threads wait at the implied barrier at the end, unless the **`nowait`** clause is specified.

The **`workshare`** construct is a Fortran feature that consists of a region with a single structure block (section of code). Statements in the **`workshare`** region are divided into units of work, and executed (once) by threads of the team.

MASTER CONSTRUCT

The **`master`** construct is not a worksharing construct. The master region is is executed only by the master thread. There is no implicit barrier (and flush) at the end of the **`master`** region; hence the other threads of the team continue execution beyond code statements beyond the **`master`** region.

1.1 A Simple Parallel Loop

The following example demonstrates how to parallelize a simple loop using the parallel loop construct. The loop iteration variable is private by default, so it is not necessary to specify it explicitly in a **`private`** clause.

C / C++

Example ploop.1.c

```
S-1  void simple(int n, float *a, float *b)
S-2  {
S-3      int i;
S-4
S-5  #pragma omp parallel for
S-6      for (i=1; i<n; i++) /* i is private by default */
S-7          b[i] = (a[i] + a[i-1]) / 2.0;
S-8  }
```

C / C++

Fortran

Example ploop.1.f

```
S-1        SUBROUTINE SIMPLE(N, A, B)
S-2
S-3        INTEGER I, N
S-4        REAL B(N), A(N)
S-5
S-6  !$OMP PARALLEL DO  !I is private by default
S-7        DO I=2,N
S-8            B(I) = (A(I) + A(I-1)) / 2.0
S-9        ENDDO
S-10 !$OMP END PARALLEL DO
S-11
S-12       END SUBROUTINE SIMPLE
```

Fortran

1.2 The `parallel` Construct

The **parallel** construct 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:

C / C++

Example parallel.1.c

```
S-1   #include <omp.h>
S-2
S-3   void subdomain(float *x, int istart, int ipoints)
S-4   {
S-5     int i;
S-6
S-7     for (i = 0; i < ipoints; i++)
S-8         x[istart+i] = 123.456;
S-9   }
S-10
S-11  void sub(float *x, int npoints)
S-12  {
S-13      int iam, nt, ipoints, istart;
S-14
S-15  #pragma omp parallel default(shared) private(iam,nt,ipoints,istart)
S-16      {
S-17          iam = omp_get_thread_num();
S-18          nt =  omp_get_num_threads();
S-19          ipoints = npoints / nt;     /* size of partition */
S-20          istart = iam * ipoints;  /* starting array index */
S-21          if (iam == nt-1)      /* last thread may do more */
S-22            ipoints = npoints - istart;
S-23          subdomain(x, istart, ipoints);
S-24      }
S-25  }
S-26
S-27  int main()
S-28  {
S-29      float array[10000];
S-30
S-31      sub(array, 10000);
S-32
S-33      return 0;
S-34  }
```

C / C++

Fortran

1 *Example parallel.1.f*

```
S-1         SUBROUTINE SUBDOMAIN(X, ISTART, IPOINTS)
S-2             INTEGER ISTART, IPOINTS
S-3             REAL X(*)
S-4
S-5             INTEGER I
S-6
S-7             DO 100 I=1,IPOINTS
S-8                X(ISTART+I) = 123.456
S-9      100    CONTINUE
S-10
S-11        END SUBROUTINE SUBDOMAIN
S-12
S-13        SUBROUTINE SUB(X, NPOINTS)
S-14            INCLUDE "omp_lib.h"     ! or USE OMP_LIB
S-15
S-16            REAL X(*)
S-17            INTEGER NPOINTS
S-18            INTEGER IAM, NT, IPOINTS, ISTART
S-19
S-20    !$OMP PARALLEL DEFAULT(PRIVATE) SHARED(X,NPOINTS)
S-21
S-22            IAM = OMP_GET_THREAD_NUM()
S-23            NT =  OMP_GET_NUM_THREADS()
S-24            IPOINTS = NPOINTS/NT
S-25            ISTART = IAM * IPOINTS
S-26            IF (IAM .EQ. NT-1) THEN
S-27                IPOINTS = NPOINTS - ISTART
S-28            ENDIF
S-29            CALL SUBDOMAIN(X,ISTART,IPOINTS)
S-30
S-31    !$OMP END PARALLEL
S-32        END SUBROUTINE SUB
S-33
S-34        PROGRAM PAREXAMPLE
S-35            REAL ARRAY(10000)
S-36            CALL SUB(ARRAY, 10000)
S-37        END PROGRAM PAREXAMPLE
```

Fortran

1.3 `teams` Construct on Host

Originally the **`teams`** construct was created for devices (such as GPUs) for independent executions of a structured block by teams within a league (on SMs). It was only available through offloading with the **`target`** construct, and the execution of a **`teams`** region could only be directed to host execution by various means such as **`if`** and **`device`** clauses, and the **`OMP_TARGET_OFFLOAD`** environment variable.

In OpenMP 5.0 the **`teams`** construct was extended to enable the host to execute a **`teams`** region (without an associated **`target`** construct), with anticipation of further affinity and threading controls in future OpenMP releases.

In the example below the **`teams`** construct is used to create two teams, one to execute single precision code, and the other to execute double precision code. Two teams are required, and the thread limit for each team is set to 1/2 of the number of available processors.

C / C++

Example host_teams.1.c (**`omp_5.0`**)

```
S-1   #include <stdio.h>
S-2   #include <stdlib.h>
S-3   #include  <math.h>
S-4   #include   <omp.h>
S-5   #define    N 1000
S-6
S-7   int main(){
S-8      int     nteams_required=2, max_thrds, tm_id;
S-9      float   sp_x[N], sp_y[N], sp_a=0.0001e0;
S-10     double  dp_x[N], dp_y[N], dp_a=0.0001e0;
S-11
S-12     // Create 2 teams, each team works in a different precision
S-13     #pragma omp teams num_teams(nteams_required) \
S-14                       thread_limit(max_thrds)  private(tm_id)
S-15     {
S-16        tm_id = omp_get_team_num();
S-17
S-18        if( omp_get_num_teams() != 2 )   //if only getting 1, quit
S-19        { printf("error: Insufficient teams on host, 2 required\n");
S-20          exit(0);
S-21        }
S-22
S-23        if(tm_id == 0)    // Do Single Precision Work (SAXPY) with this team
S-24        {
S-25           #pragma omp parallel
S-26           {
S-27              #pragma omp for                              //init
S-28              for(int i=0; i<N; i++){sp_x[i] = i*0.0001;  sp_y[i]=i; }
```

```
S-29
S-30                #pragma omp for simd simdlen(8)
S-31                for(int i=0; i<N; i++){sp_x[i] = sp_a*sp_x[i] + sp_y[i];}
S-32             }
S-33          }
S-34
S-35          if(tm_id == 1)    // Do Double Precision Work (DAXPY) with this team
S-36          {
S-37             #pragma omp parallel
S-38             {
S-39                #pragma omp for                                  //init
S-40                for(int i=0; i<N; i++){dp_x[i] = i*0.0001;  dp_y[i]=i; }
S-41
S-42                #pragma omp for simd simdlen(4)
S-43                for(int i=0; i<N; i++){dp_x[i] = dp_a*dp_x[i] + dp_y[i];}
S-44             }
S-45          }
S-46       }
S-47
S-48       printf("i=%d  sp|dp  %f %f \n",N-1, sp_x[N-1], dp_x[N-1]);
S-49       printf("i=%d  sp|dp  %f %f \n",N/2, sp_x[N/2], dp_x[N/2]);
S-50     //OUTPUT1:i=999  sp|dp  999.000000 999.000010
S-51     //OUTPUT2:i=500  sp|dp  500.000000 500.000005
S-52
S-53       return 0;
S-54    }
```

C / C++

Fortran

1 *Example host_teams.1.f90* (`omp_5.0`)

```
S-1
S-2     program main
S-3        use omp_lib
S-4        integer           :: nteams_required=2, max_thrds, tm_id
S-5        integer,parameter ::  N=1000
S-6        real              :: sp_x(N), sp_y(N), sp_a=0.0001e0
S-7        double precision  :: dp_x(N), dp_y(N), dp_a=0.0001d0
S-8
S-9        max_thrds  = omp_get_num_procs()/nteams_required
S-10
S-11       !! Create 2 teams, each team works in a different precision
S-12       !$omp teams num_teams(nteams_required) thread_limit(max_thrds) private(tm_id)
S-13
S-14          tm_id = omp_get_team_num()
S-15
S-16          if( omp_get_num_teams() /= 2 ) then   !! if only getting 1, quit
```

```
S-17              stop "error: Insufficient teams on host, 2 required."
S-18           endif
S-19
S-20           if(tm_id == 0) then    !! Do Single Precision Work (SAXPY) with this team
S-21
S-22              !$omp parallel
S-23                 !$omp do         !! init
S-24                 do i = 1,N
S-25                    sp_x(i) = i*0.0001e0
S-26                    sp_y(i) = i
S-27                 end do
S-28
S-29                 !$omp do simd simdlen(8)
S-30                 do i = 1,N
S-31                    sp_x(i) = sp_a*sp_x(i) + sp_y(i)
S-32                 end do
S-33              !$omp end parallel
S-34
S-35           endif
S-36
S-37           if(tm_id == 1) then    !! Do Double Precision Work (DAXPY) with this team
S-38
S-39              !$omp parallel
S-40                 !$omp do         !! init
S-41                 do i = 1,N
S-42                    dp_x(i) = i*0.0001d0
S-43                    dp_y(i) = i
S-44                 end do
S-45
S-46                 !$omp do simd simdlen(4)
S-47                 do i = 1,N
S-48                    dp_x(i) = dp_a*dp_x(i) + dp_y(i)
S-49                 end do
S-50              !$omp end parallel
S-51
S-52           endif
S-53        !$omp end teams
S-54
S-55        write(*,'( "i=",i4," sp|dp= ", e15.7, d25.16  )') N, sp_x(N), dp_x(N)
S-56        write(*,'( "i=",i4," sp|dp= ", e15.7, d25.16  )') N/2, sp_x(N/2), dp_x(N/2)
S-57                 !! i=1000 sp|dp=   0.1000000E+04   0.10000000100000000D+04
S-58                 !! i= 500 sp|dp=   0.5000000E+03   0.50000000500000000D+03
S-59     end program
```

Fortran

1.4 Controlling the Number of Threads on Multiple Nesting Levels

The following examples demonstrate how to use the **OMP_NUM_THREADS** environment variable to control the number of threads on multiple nesting levels:

C / C++

Example nthrs_nesting.1.c

```
S-1   #include <stdio.h>
S-2   #include <omp.h>
S-3   int main (void)
S-4   {
S-5      omp_set_nested(1);
S-6      omp_set_dynamic(0);
S-7      #pragma omp parallel
S-8      {
S-9         #pragma omp parallel
S-10        {
S-11           #pragma omp single
S-12           {
S-13           /*
S-14           * If OMP_NUM_THREADS=2,3 was set, the following should print:
S-15           * Inner: num_thds=3
S-16           * Inner: num_thds=3
S-17           *
S-18           * If nesting is not supported, the following should print:
S-19           * Inner: num_thds=1
S-20           * Inner: num_thds=1
S-21           */
S-22              printf ("Inner: num_thds=%d\n", omp_get_num_threads());
S-23           }
S-24        }
S-25        #pragma omp barrier
S-26        omp_set_nested(0);
S-27        #pragma omp parallel
S-28        {
S-29           #pragma omp single
S-30           {
S-31           /*
S-32           * Even if OMP_NUM_THREADS=2,3 was set, the following should
S-33           * print, because nesting is disabled:
S-34           * Inner: num_thds=1
S-35           * Inner: num_thds=1
S-36           */
S-37              printf ("Inner: num_thds=%d\n", omp_get_num_threads());
S-38           }
```

```
S-39        }
S-40        #pragma omp barrier
S-41        #pragma omp single
S-42        {
S-43           /*
S-44           * If OMP_NUM_THREADS=2,3 was set, the following should print:
S-45           * Outer: num_thds=2
S-46           */
S-47           printf ("Outer: num_thds=%d\n", omp_get_num_threads());
S-48        }
S-49     }
S-50     return 0;
S-51  }
```

C / C++

Fortran

1 *Example nthrs_nesting.1.f*

```
S-1           program icv
S-2           use omp_lib
S-3           call omp_set_nested(.true.)
S-4           call omp_set_dynamic(.false.)
S-5     !$omp parallel
S-6     !$omp parallel
S-7     !$omp single
S-8           ! If OMP_NUM_THREADS=2,3 was set, the following should print:
S-9           ! Inner: num_thds= 3
S-10          ! Inner: num_thds= 3
S-11          ! If nesting is not supported, the following should print:
S-12          ! Inner: num_thds= 1
S-13          ! Inner: num_thds= 1
S-14          print *, "Inner: num_thds=", omp_get_num_threads()
S-15    !$omp end single
S-16    !$omp end parallel
S-17    !$omp barrier
S-18          call omp_set_nested(.false.)
S-19    !$omp parallel
S-20    !$omp single
S-21          ! Even if OMP_NUM_THREADS=2,3 was set, the following should print,
S-22          ! because nesting is disabled:
S-23          ! Inner: num_thds= 1
S-24          ! Inner: num_thds= 1
S-25          print *, "Inner: num_thds=", omp_get_num_threads()
S-26    !$omp end single
S-27    !$omp end parallel
S-28    !$omp barrier
S-29    !$omp single
```

```
S-30            ! If OMP_NUM_THREADS=2,3 was set, the following should print:
S-31            ! Outer: num_thds= 2
S-32            print *, "Outer: num_thds=", omp_get_num_threads()
S-33    !$omp end single
S-34    !$omp end parallel
S-35            end
```

Fortran

1.5 Interaction Between the `num_threads` Clause and `omp_set_dynamic`

The following example demonstrates the **`num_threads`** clause and the effect of the **`omp_set_dynamic`** routine on it.

The call to the **`omp_set_dynamic`** routine with argument **`0`** in C/C++, or **`.FALSE.`** in Fortran, disables the dynamic adjustment of the number of threads in OpenMP implementations that support it. In this case, 10 threads are provided. Note that in case of an error the OpenMP implementation is free to abort the program or to supply any number of threads available.

C / C++

Example nthrs_dynamic.1.c

```
S-1   #include <omp.h>
S-2   int main()
S-3   {
S-4     omp_set_dynamic(0);
S-5     #pragma omp parallel num_threads(10)
S-6     {
S-7       /* do work here */
S-8     }
S-9     return 0;
S-10  }
```

C / C++

Fortran

Example nthrs_dynamic.1.f

```
S-1         PROGRAM EXAMPLE
S-2           INCLUDE "omp_lib.h"      ! or USE OMP_LIB
S-3           CALL OMP_SET_DYNAMIC(.FALSE.)
S-4   !$OMP     PARALLEL NUM_THREADS(10)
S-5               ! do work here
S-6   !$OMP     END PARALLEL
S-7         END PROGRAM EXAMPLE
```

Fortran

The call to the **omp_set_dynamic** routine with a non-zero argument in C/C++, or **.TRUE.** in Fortran, allows the OpenMP implementation to choose any number of threads between 1 and 10.

C / C++

Example nthrs_dynamic.2.c

```
S-1   #include <omp.h>
S-2   int main()
S-3   {
S-4     omp_set_dynamic(1);
S-5     #pragma omp parallel num_threads(10)
S-6     {
S-7       /* do work here */
S-8     }
S-9     return 0;
S-10  }
```

C / C++

Fortran

Example nthrs_dynamic.2.f

```
S-1         PROGRAM EXAMPLE
S-2           INCLUDE "omp_lib.h"      ! or USE OMP_LIB
S-3           CALL OMP_SET_DYNAMIC(.TRUE.)
S-4   !$OMP     PARALLEL NUM_THREADS(10)
S-5             ! do work here
S-6   !$OMP     END PARALLEL
S-7         END PROGRAM EXAMPLE
```

Fortran

It is good practice to set the *dyn-var* ICV explicitly by calling the **omp_set_dynamic** routine, as its default setting is implementation defined.

1.6 Fortran Restrictions on the do Construct

Fortran

If an **end do** directive follows a *do-construct* in which several **DO** statements share a **DO** termination statement, then a **do** directive can only be specified for the outermost of these **DO** statements. The following example contains correct usages of loop constructs:

Example fort_do.1.f

```
S-1         SUBROUTINE WORK(I, J)
S-2         INTEGER I,J
S-3         END SUBROUTINE WORK
S-4
S-5         SUBROUTINE DO_GOOD()
S-6           INTEGER I, J
S-7           REAL A(1000)
S-8
S-9           DO 100 I = 1,10
S-10    !$OMP   DO
S-11            DO 100 J = 1,10
S-12              CALL WORK(I,J)
S-13    100     CONTINUE      !  !$OMP ENDDO implied here
S-14
S-15    !$OMP   DO
S-16            DO 200 J = 1,10
S-17    200       A(I) = I + 1
S-18    !$OMP   ENDDO
S-19
S-20    !$OMP   DO
S-21            DO 300 I = 1,10
S-22              DO 300 J = 1,10
S-23                CALL WORK(I,J)
S-24    300     CONTINUE
S-25    !$OMP   ENDDO
S-26        END SUBROUTINE DO_GOOD
```

The following example is non-conforming because the matching **do** directive for the **end do** does not precede the outermost loop:

Example fort_do.2.f

```
S-1         SUBROUTINE WORK(I, J)
S-2         INTEGER I,J
S-3         END SUBROUTINE WORK
S-4
S-5         SUBROUTINE DO_WRONG
S-6           INTEGER I, J
S-7
```

```
 S-8            DO 100 I = 1,10
 S-9     !$OMP    DO
S-10              DO 100 J = 1,10
S-11                CALL WORK(I,J)
S-12     100      CONTINUE
S-13     !$OMP    ENDDO
S-14           END SUBROUTINE DO_WRONG
```

Fortran

1.7 The `nowait` Clause

If there are multiple independent loops within a **parallel** region, you can use the **nowait** clause to avoid the implied barrier at the end of the loop construct, as follows:

C / C++

Example nowait.1.c

```
S-1   #include <math.h>
S-2
S-3   void nowait_example(int n, int m, float *a, float *b, float *y, float *z)
S-4   {
S-5     int i;
S-6     #pragma omp parallel
S-7     {
S-8       #pragma omp for nowait
S-9         for (i=1; i<n; i++)
S-10          b[i] = (a[i] + a[i-1]) / 2.0;
S-11
S-12      #pragma omp for nowait
S-13        for (i=0; i<m; i++)
S-14          y[i] = sqrt(z[i]);
S-15    }
S-16  }
```

C / C++

Fortran

Example nowait.1.f

```
S-1           SUBROUTINE NOWAIT_EXAMPLE(N, M, A, B, Y, Z)
S-2
S-3           INTEGER N, M
S-4           REAL A(*), B(*), Y(*), Z(*)
S-5
S-6           INTEGER I
S-7
S-8   !$OMP PARALLEL
S-9
S-10  !$OMP DO
S-11          DO I=2,N
S-12            B(I) = (A(I) + A(I-1)) / 2.0
S-13          ENDDO
S-14  !$OMP END DO NOWAIT
S-15
S-16  !$OMP DO
S-17          DO I=1,M
S-18            Y(I) = SQRT(Z(I))
```

```
S-19            ENDDO
S-20    !$OMP END DO NOWAIT
S-21
S-22    !$OMP END PARALLEL
S-23
S-24            END SUBROUTINE NOWAIT_EXAMPLE
```

Fortran

In the following example, static scheduling distributes the same logical iteration numbers to the threads that execute the three loop regions. This allows the **nowait** clause to be used, even though there is a data dependence between the loops. The dependence is satisfied as long the same thread executes the same logical iteration numbers in each loop.

Note that the iteration count of the loops must be the same. The example satisfies this requirement, since the iteration space of the first two loops is from **0** to **n−1** (from **1** to **N** in the Fortran version), while the iteration space of the last loop is from **1** to **n** (**2** to **N+1** in the Fortran version).

C / C++

Example nowait.2.c

```
S-1
S-2     #include <math.h>
S-3     void nowait_example2(int n, float *a, float *b, float *c, float *y, float
S-4     *z)
S-5     {
S-6        int i;
S-7     #pragma omp parallel
S-8        {
S-9     #pragma omp for schedule(static) nowait
S-10       for (i=0; i<n; i++)
S-11          c[i] = (a[i] + b[i]) / 2.0f;
S-12    #pragma omp for schedule(static) nowait
S-13       for (i=0; i<n; i++)
S-14          z[i] = sqrtf(c[i]);
S-15    #pragma omp for schedule(static) nowait
S-16       for (i=1; i<=n; i++)
S-17          y[i] = z[i-1] + a[i];
S-18       }
S-19    }
```

C / C++

Fortran

Example nowait.2.f90

```
S-1       SUBROUTINE NOWAIT_EXAMPLE2(N, A, B, C, Y, Z)
S-2       INTEGER N
S-3       REAL A(*), B(*), C(*), Y(*), Z(*)
S-4       INTEGER I
S-5    !$OMP PARALLEL
S-6    !$OMP DO SCHEDULE(STATIC)
S-7       DO I=1,N
S-8          C(I) = (A(I) + B(I)) / 2.0
S-9       ENDDO
S-10   !$OMP END DO NOWAIT
S-11   !$OMP DO SCHEDULE(STATIC)
S-12      DO I=1,N
S-13         Z(I) = SQRT(C(I))
S-14      ENDDO
S-15   !$OMP END DO NOWAIT
S-16   !$OMP DO SCHEDULE(STATIC)
S-17      DO I=2,N+1
S-18         Y(I) = Z(I-1) + A(I)
S-19      ENDDO
S-20   !$OMP END DO NOWAIT
S-21   !$OMP END PARALLEL
S-22      END SUBROUTINE NOWAIT_EXAMPLE2
```

Fortran

1.8 The `collapse` Clause

In the following example, the **k** and **j** loops are associated with the loop construct. So the iterations of the **k** and **j** loops are collapsed into one loop with a larger iteration space, and that loop is then divided among the threads in the current team. Since the **i** loop is not associated with the loop construct, it is not collapsed, and the **i** loop is executed sequentially in its entirety in every iteration of the collapsed **k** and **j** loop.

The variable **j** can be omitted from the **private** clause when the **collapse** clause is used since it is implicitly private. However, if the **collapse** clause is omitted then **j** will be shared if it is omitted from the **private** clause. In either case, **k** is implicitly private and could be omitted from the **private** clause.

C / C++

Example collapse.1.c (**omp_3.0**)

```
S-1
S-2   void bar(float *a, int i, int j, int k);
S-3
S-4   int kl, ku, ks, jl, ju, js, il, iu,is;
S-5
S-6   void sub(float *a)
S-7   {
S-8       int i, j, k;
S-9
S-10      #pragma omp for collapse(2) private(i, k, j)
S-11      for (k=kl; k<=ku; k+=ks)
S-12         for (j=jl; j<=ju; j+=js)
S-13            for (i=il; i<=iu; i+=is)
S-14               bar(a,i,j,k);
S-15  }
```

C / C++

Fortran

Example collapse.1.f (**omp_3.0**)

```
S-1
S-2         subroutine sub(a)
S-3
S-4         real a(*)
S-5         integer kl, ku, ks, jl, ju, js, il, iu, is
S-6         common /csub/ kl, ku, ks, jl, ju, js, il, iu, is
S-7         integer i, j, k
S-8
S-9   !$omp do collapse(2) private(i,j,k)
S-10         do k = kl, ku, ks
S-11           do j = jl, ju, js
```

```
S-12              do i = il, iu, is
S-13                call bar(a,i,j,k)
S-14              enddo
S-15            enddo
S-16          enddo
S-17    !$omp end do
S-18
S-19          end subroutine
```

Fortran

In the next example, the **k** and **j** loops are associated with the loop construct. So the iterations of the **k** and **j** loops are collapsed into one loop with a larger iteration space, and that loop is then divided among the threads in the current team.

The sequential execution of the iterations in the **k** and **j** loops determines the order of the iterations in the collapsed iteration space. This implies that in the sequentially last iteration of the collapsed iteration space, **k** will have the value **2** and **j** will have the value **3**. Since **klast** and **jlast** are **lastprivate**, their values are assigned by the sequentially last iteration of the collapsed **k** and **j** loop. This example prints: **2 3**.

C / C++

Example collapse.2.c (**omp_3.0**)

```
S-1
S-2     #include <stdio.h>
S-3     void test()
S-4     {
S-5        int j, k, jlast, klast;
S-6        #pragma omp parallel
S-7        {
S-8           #pragma omp for collapse(2) lastprivate(jlast, klast)
S-9           for (k=1; k<=2; k++)
S-10             for (j=1; j<=3; j++)
S-11             {
S-12                jlast=j;
S-13                klast=k;
S-14             }
S-15          #pragma omp single
S-16          printf("%d %d\n", klast, jlast);
S-17       }
S-18    }
```

C / C++

Fortran

Example collapse.2.f (**omp_3.0**)

```
S-1
S-2         program test
S-3   !$omp parallel
S-4   !$omp do private(j,k) collapse(2) lastprivate(jlast, klast)
S-5         do k = 1,2
S-6           do j = 1,3
S-7             jlast=j
S-8             klast=k
S-9           enddo
S-10        enddo
S-11  !$omp end do
S-12  !$omp single
S-13        print *, klast, jlast
S-14  !$omp end single
S-15  !$omp end parallel
S-16        end program test
```

Fortran

The next example illustrates the interaction of the **collapse** and **ordered** clauses.

In the example, the loop construct has both a **collapse** clause and an **ordered** clause. The **collapse** clause causes the iterations of the **k** and **j** loops to be collapsed into one loop with a larger iteration space, and that loop is divided among the threads in the current team. An **ordered** clause is added to the loop construct, because an ordered region binds to the loop region arising from the loop construct.

According to Section 2.12.8 of the OpenMP 4.0 specification, a thread must not execute more than one ordered region that binds to the same loop region. So the **collapse** clause is required for the example to be conforming. With the **collapse** clause, the iterations of the **k** and **j** loops are collapsed into one loop, and therefore only one ordered region will bind to the collapsed **k** and **j** loop. Without the **collapse** clause, there would be two ordered regions that bind to each iteration of the **k** loop (one arising from the first iteration of the **j** loop, and the other arising from the second iteration of the **j** loop).

The code prints

```
0 1 1
0 1 2
0 2 1
1 2 2
1 3 1
1 3 2
```

C / C++

Example collapse.3.c (`omp_3.0`)

```
S-1   #include <omp.h>
S-2   #include <stdio.h>
S-3   void work(int a, int j, int k);
S-4   void sub()
S-5   {
S-6      int j, k, a;
S-7      #pragma omp parallel num_threads(2)
S-8      {
S-9         #pragma omp for collapse(2) ordered private(j,k) schedule(static,3)
S-10        for (k=1; k<=3; k++)
S-11           for (j=1; j<=2; j++)
S-12           {
S-13              #pragma omp ordered
S-14              printf("%d %d %d\n", omp_get_thread_num(), k, j);
S-15              /* end ordered */
S-16              work(a,j,k);
S-17           }
S-18     }
S-19  }
```

C / C++

Fortran

Example collapse.3.f (`omp_3.0`)

```
S-1         program test
S-2         include 'omp_lib.h'
S-3   !$omp parallel num_threads(2)
S-4   !$omp do collapse(2) ordered private(j,k) schedule(static,3)
S-5         do k = 1,3
S-6           do j = 1,2
S-7   !$omp ordered
S-8             print *, omp_get_thread_num(), k, j
S-9   !$omp end ordered
S-10            call work(a,j,k)
S-11          enddo
S-12        enddo
S-13  !$omp end do
S-14  !$omp end parallel
S-15        end program test
```

Fortran

1.9 `linear` Clause in Loop Constructs

The following example shows the use of the **`linear`** clause in a loop construct to allow the proper parallelization of a loop that contains an induction variable (*j*). At the end of the execution of the loop construct, the original variable *j* is updated with the value *N/2* from the last iteration of the loop.

C / C++

Example linear_in_loop.1.c (**`omp_4.5`**)

```
S-1   #include <stdio.h>
S-2
S-3   #define N 100
S-4   int main(void)
S-5   {
S-6      float a[N], b[N/2];
S-7      int i, j;
S-8
S-9      for ( i = 0; i < N; i++ )
S-10        a[i] = i + 1;
S-11
S-12     j = 0;
S-13     #pragma omp parallel
S-14     #pragma omp for linear(j:1)
S-15     for ( i = 0; i < N; i += 2 ) {
S-16        b[j] = a[i] * 2.0f;
S-17        j++;
S-18     }
S-19
S-20     printf( "%d %f %f\n", j, b[0], b[j-1] );
S-21     /* print out: 50 2.0 198.0 */
S-22
S-23     return 0;
S-24  }
```

C / C++

Fortran

Example linear_in_loop.1.f90 (**omp_4.5**)

```
S-1   program linear_loop
S-2      implicit none
S-3      integer, parameter :: N = 100
S-4      real :: a(N), b(N/2)
S-5      integer :: i, j
S-6
S-7      do i = 1, N
S-8         a(i) = i
S-9      end do
S-10
S-11     j = 0
S-12     !$omp parallel
S-13     !$omp do linear(j:1)
S-14     do i = 1, N, 2
S-15        j = j + 1
S-16        b(j) = a(i) * 2.0
S-17     end do
S-18     !$omp end parallel
S-19
S-20     print *, j, b(1), b(j)
S-21     ! print out: 50 2.0 198.0
S-22
S-23  end program
```

Fortran

1.10 The `parallel sections` Construct

In the following example routines **XAXIS**, **YAXIS**, and **ZAXIS** can be executed concurrently. The first **section** directive is optional. Note that all **section** directives need to appear in the **parallel sections** construct.

C / C++

Example psections.1.c

```
S-1   void XAXIS();
S-2   void YAXIS();
S-3   void ZAXIS();
S-4
S-5   void sect_example()
S-6   {
S-7     #pragma omp parallel sections
S-8     {
S-9       #pragma omp section
S-10        XAXIS();
S-11
S-12      #pragma omp section
S-13        YAXIS();
S-14
S-15      #pragma omp section
S-16        ZAXIS();
S-17    }
S-18  }
```

C / C++

Fortran

Example psections.1.f

```
S-1         SUBROUTINE SECT_EXAMPLE()
S-2   !$OMP PARALLEL SECTIONS
S-3   !$OMP SECTION
S-4           CALL XAXIS()
S-5   !$OMP SECTION
S-6           CALL YAXIS()
S-7
S-8   !$OMP SECTION
S-9           CALL ZAXIS()
S-10
S-11  !$OMP END PARALLEL SECTIONS
S-12        END SUBROUTINE SECT_EXAMPLE
```

Fortran

1.11 The `firstprivate` Clause and the `sections` Construct

In the following example of the **`sections`** construct the **`firstprivate`** clause is used to initialize the private copy of **`section_count`** of each thread. The problem is that the **`section`** constructs modify **`section_count`**, which breaks the independence of the **`section`** constructs. When different threads execute each section, both sections will print the value 1. When the same thread executes the two sections, one section will print the value 1 and the other will print the value 2. Since the order of execution of the two sections in this case is unspecified, it is unspecified which section prints which value.

C / C++

Example fpriv_sections.1.c

```
S-1   #include <omp.h>
S-2   #include <stdio.h>
S-3   #define NT 4
S-4   int main( ) {
S-5       int section_count = 0;
S-6       omp_set_dynamic(0);
S-7       omp_set_num_threads(NT);
S-8   #pragma omp parallel
S-9   #pragma omp sections firstprivate( section_count )
S-10  {
S-11  #pragma omp section
S-12      {
S-13          section_count++;
S-14          /* may print the number one or two */
S-15          printf( "section_count %d\n", section_count );
S-16      }
S-17  #pragma omp section
S-18      {
S-19          section_count++;
S-20          /* may print the number one or two */
S-21          printf( "section_count %d\n", section_count );
S-22      }
S-23  }
S-24      return 0;
S-25  }
```

C / C++

Fortran

Example fpriv_sections.1.f90

```
S-1   program section
S-2       use omp_lib
S-3       integer :: section_count = 0
S-4       integer, parameter :: NT = 4
S-5       call omp_set_dynamic(.false.)
S-6       call omp_set_num_threads(NT)
S-7   !$omp parallel
S-8   !$omp sections firstprivate ( section_count )
S-9   !$omp section
S-10      section_count = section_count + 1
S-11  ! may print the number one or two
S-12      print *, 'section_count', section_count
S-13  !$omp section
S-14      section_count = section_count + 1
S-15  ! may print the number one or two
S-16      print *, 'section_count', section_count
S-17  !$omp end sections
S-18  !$omp end parallel
S-19  end program section
```

Fortran

1.12 The `single` Construct

The following example demonstrates the **`single`** construct. In the example, only one thread prints each of the progress messages. All other threads will skip the **`single`** region and stop at the barrier at the end of the **`single`** construct until all threads in the team have reached the barrier. If other threads can proceed without waiting for the thread executing the **`single`** region, a **`nowait`** clause can be specified, as is done in the third **`single`** construct in this example. The user must not make any assumptions as to which thread will execute a **`single`** region.

C / C++

Example single.1.c

```
S-1   #include <stdio.h>
S-2
S-3   void work1() {}
S-4   void work2() {}
S-5
S-6   void single_example()
S-7   {
S-8     #pragma omp parallel
S-9     {
S-10      #pragma omp single
S-11        printf("Beginning work1.\n");
S-12
S-13      work1();
S-14
S-15      #pragma omp single
S-16        printf("Finishing work1.\n");
S-17
S-18      #pragma omp single nowait
S-19        printf("Finished work1 and beginning work2.\n");
S-20
S-21      work2();
S-22    }
S-23  }
```

C / C++

Fortran

Example single.1.f

```
S-1         SUBROUTINE WORK1()
S-2         END SUBROUTINE WORK1
S-3
S-4         SUBROUTINE WORK2()
S-5         END SUBROUTINE WORK2
S-6
S-7         PROGRAM SINGLE_EXAMPLE
S-8   !$OMP PARALLEL
S-9
S-10  !$OMP SINGLE
S-11          print *, "Beginning work1."
S-12  !$OMP END SINGLE
S-13
S-14          CALL WORK1()
S-15
S-16  !$OMP SINGLE
S-17          print *, "Finishing work1."
S-18  !$OMP END SINGLE
S-19
S-20  !$OMP SINGLE
S-21          print *, "Finished work1 and beginning work2."
S-22  !$OMP END SINGLE NOWAIT
S-23
S-24          CALL WORK2()
S-25
S-26  !$OMP END PARALLEL
S-27
S-28        END PROGRAM SINGLE_EXAMPLE
```

Fortran

1.13 The `workshare` Construct

Fortran

The following are examples of the **workshare** construct.

In the following example, **workshare** spreads work across the threads executing the **parallel** region, and there is a barrier after the last statement. Implementations must enforce Fortran execution rules inside of the **workshare** block.

Example workshare.1.f

```
S-1         SUBROUTINE WSHARE1(AA, BB, CC, DD, EE, FF, N)
S-2         INTEGER N
S-3         REAL AA(N,N), BB(N,N), CC(N,N), DD(N,N), EE(N,N), FF(N,N)
S-4
S-5   !$OMP    PARALLEL
S-6   !$OMP     WORKSHARE
S-7               AA = BB
S-8               CC = DD
S-9               EE = FF
S-10  !$OMP     END WORKSHARE
S-11  !$OMP   END PARALLEL
S-12
S-13        END SUBROUTINE WSHARE1
```

In the following example, the barrier at the end of the first **workshare** region is eliminated with a **nowait** clause. Threads doing **CC = DD** immediately begin work on **EE = FF** when they are done with **CC = DD**.

Example workshare.2.f

```
S-1         SUBROUTINE WSHARE2(AA, BB, CC, DD, EE, FF, N)
S-2         INTEGER N
S-3         REAL AA(N,N), BB(N,N), CC(N,N)
S-4         REAL DD(N,N), EE(N,N), FF(N,N)
S-5
S-6   !$OMP   PARALLEL
S-7   !$OMP     WORKSHARE
S-8               AA = BB
S-9               CC = DD
S-10  !$OMP     END WORKSHARE NOWAIT
S-11  !$OMP     WORKSHARE
S-12              EE = FF
S-13  !$OMP     END WORKSHARE
S-14  !$OMP   END PARALLEL
S-15        END SUBROUTINE WSHARE2
```

The following example shows the use of an **atomic** directive inside a **workshare** construct. The computation of **SUM(AA)** is workshared, but the update to **R** is atomic.

Example workshare.3.f

```
S-1         SUBROUTINE WSHARE3(AA, BB, CC, DD, N)
S-2         INTEGER N
S-3         REAL AA(N,N), BB(N,N), CC(N,N), DD(N,N)
S-4         REAL R
S-5           R=0
S-6   !$OMP   PARALLEL
S-7   !$OMP     WORKSHARE
S-8               AA = BB
S-9   !$OMP       ATOMIC UPDATE
S-10                R = R + SUM(AA)
S-11              CC = DD
S-12  !$OMP     END WORKSHARE
S-13  !$OMP   END PARALLEL
S-14        END SUBROUTINE WSHARE3
```

Fortran **WHERE** and **FORALL** statements are *compound statements*, made up of a *control* part and a *statement* part. When **workshare** is applied to one of these compound statements, both the control and the statement parts are workshared. The following example shows the use of a **WHERE** statement in a **workshare** construct.

Each task gets worked on in order by the threads:

AA = BB then
CC = DD then
EE .ne. 0 then
FF = 1 / EE then
GG = HH

Example workshare.4.f

```
S-1         SUBROUTINE WSHARE4(AA, BB, CC, DD, EE, FF, GG, HH, N)
S-2         INTEGER N
S-3         REAL AA(N,N), BB(N,N), CC(N,N)
S-4         REAL DD(N,N), EE(N,N), FF(N,N)
S-5         REAL GG(N,N), HH(N,N)
S-6
S-7   !$OMP   PARALLEL
S-8   !$OMP     WORKSHARE
S-9               AA = BB
S-10              CC = DD
S-11              WHERE (EE .ne. 0) FF = 1 / EE
```

```
S-12              GG = HH
S-13     !$OMP    END WORKSHARE
S-14     !$OMP  END PARALLEL
S-15
S-16           END SUBROUTINE WSHARE4
```

In the following example, an assignment to a shared scalar variable is performed by one thread in a **workshare** while all other threads in the team wait.

Example workshare.5.f

```
S-1            SUBROUTINE WSHARE5(AA, BB, CC, DD, N)
S-2            INTEGER N
S-3            REAL AA(N,N), BB(N,N), CC(N,N), DD(N,N)
S-4
S-5              INTEGER SHR
S-6
S-7      !$OMP   PARALLEL SHARED(SHR)
S-8      !$OMP     WORKSHARE
S-9                AA = BB
S-10               SHR = 1
S-11               CC = DD * SHR
S-12     !$OMP     END WORKSHARE
S-13     !$OMP   END PARALLEL
S-14
S-15           END SUBROUTINE WSHARE5
```

The following example contains an assignment to a private scalar variable, which is performed by one thread in a **workshare** while all other threads wait. It is non-conforming because the private scalar variable is undefined after the assignment statement.

Example workshare.6.f

```
S-1            SUBROUTINE WSHARE6_WRONG(AA, BB, CC, DD, N)
S-2            INTEGER N
S-3            REAL AA(N,N), BB(N,N), CC(N,N), DD(N,N)
S-4
S-5              INTEGER PRI
S-6
S-7      !$OMP   PARALLEL PRIVATE(PRI)
S-8      !$OMP     WORKSHARE
S-9                AA = BB
S-10               PRI = 1
S-11               CC = DD * PRI
S-12     !$OMP     END WORKSHARE
S-13     !$OMP   END PARALLEL
```

```
      END SUBROUTINE WSHARE6_WRONG
```

Fortran execution rules must be enforced inside a **workshare** construct. In the following example, the same result is produced in the following program fragment regardless of whether the code is executed sequentially or inside an OpenMP program with multiple threads:

Example workshare.7.f

```
      SUBROUTINE WSHARE7(AA, BB, CC, N)
      INTEGER N
      REAL AA(N), BB(N), CC(N)

!$OMP   PARALLEL
!$OMP     WORKSHARE
            AA(1:50)  = BB(11:60)
            CC(11:20) = AA(1:10)
!$OMP     END WORKSHARE
!$OMP   END PARALLEL

      END SUBROUTINE WSHARE7
```

Fortran

1.14 The `master` Construct

The following example demonstrates the master construct . In the example, the master keeps track of how many iterations have been executed and prints out a progress report. The other threads skip the master region without waiting.

C / C++

Example master.1.c

```
S-1   #include <stdio.h>
S-2
S-3   extern float average(float,float,float);
S-4
S-5   void master_example( float* x, float* xold, int n, float tol )
S-6   {
S-7     int c, i, toobig;
S-8     float error, y;
S-9     c = 0;
S-10    #pragma omp parallel
S-11    {
S-12      do{
S-13        #pragma omp for private(i)
S-14        for( i = 1; i < n-1; ++i ){
S-15          xold[i] = x[i];
S-16        }
S-17        #pragma omp single
S-18        {
S-19          toobig = 0;
S-20        }
S-21        #pragma omp for private(i,y,error) reduction(+:toobig)
S-22        for( i = 1; i < n-1; ++i ){
S-23          y = x[i];
S-24          x[i] = average( xold[i-1], x[i], xold[i+1] );
S-25          error = y - x[i];
S-26          if( error > tol || error < -tol ) ++toobig;
S-27        }
S-28        #pragma omp master
S-29        {
S-30          ++c;
S-31          printf( "iteration %d, toobig=%d\n", c, toobig );
S-32        }
S-33      }while( toobig > 0 );
S-34    }
S-35  }
```

C / C++

Fortran

1 *Example master.1.f*

```
S-1         SUBROUTINE MASTER_EXAMPLE( X, XOLD, N, TOL )
S-2         REAL X(*), XOLD(*), TOL
S-3         INTEGER N
S-4         INTEGER C, I, TOOBIG
S-5         REAL ERROR, Y, AVERAGE
S-6         EXTERNAL AVERAGE
S-7         C = 0
S-8         TOOBIG = 1
S-9      !$OMP PARALLEL
S-10           DO WHILE( TOOBIG > 0 )
S-11     !$OMP     DO PRIVATE(I)
S-12                 DO I = 2, N-1
S-13                   XOLD(I) = X(I)
S-14                 ENDDO
S-15     !$OMP     SINGLE
S-16                 TOOBIG = 0
S-17     !$OMP     END SINGLE
S-18     !$OMP     DO PRIVATE(I,Y,ERROR), REDUCTION(+:TOOBIG)
S-19                 DO I = 2, N-1
S-20                   Y = X(I)
S-21                   X(I) = AVERAGE( XOLD(I-1), X(I), XOLD(I+1) )
S-22                   ERROR = Y-X(I)
S-23                   IF( ERROR > TOL .OR. ERROR < -TOL ) TOOBIG = TOOBIG+1
S-24                 ENDDO
S-25     !$OMP     MASTER
S-26                 C = C + 1
S-27                 PRINT *, 'Iteration ', C, 'TOOBIG=', TOOBIG
S-28     !$OMP     END MASTER
S-29           ENDDO
S-30     !$OMP END PARALLEL
S-31        END SUBROUTINE MASTER_EXAMPLE
```

Fortran

1.15 The `loop` Construct

The following example illustrates the use of the OpenMP 5.0 **`loop`** construct for the execution of a loop. The **`loop`** construct asserts to the compiler that the iterations of the loop are free of data dependencies and may be executed concurrently. It allows the compiler to use heuristics to select the parallelization scheme and compiler-level optimizations for the concurrency.

C / C++

Example loop.1.c (**`omp_5.0`**)

```
S-1   #include  <stdio.h>
S-2   #define N 100
S-3   int main()
S-4   {
S-5     float x[N], y[N];
S-6     float a = 2.0;
S-7     for(int i=0;i<N;i++){ x[i]=i; y[i]=0;}   // initialize
S-8
S-9     #pragma omp parallel
S-10    {
S-11      #pragma omp loop
S-12      for(int i = 0; i < N; ++i) y[i] = a*x[i] + y[i];
S-13    }
S-14    if(y[N-1] != (N-1)*2.0) printf("Error: 2*(N-1) != y[N-1]=%f",y[N-1]);
S-15  }
```

C / C++

Fortran

Example loop.1.f90 (**`omp_5.0`**)

```
S-1
S-2   program main
S-3     integer, parameter :: N=100
S-4     real :: x(N), y(N)
S-5     real :: a = 2.0e0
S-6
S-7     x=(/ (i,i=1,N) /); y=1.0e0                    !! initialize
S-8
S-9     !$omp parallel
S-10      !$omp loop
S-11         do i=1,N; y(i) = a*x(i) + y(i); enddo
S-12    !$omp end parallel
S-13
S-14    if(y(N) /= N*2.0e0) print*,"Error: 2*N /= y(N); y(N)=",y(N)
S-15  end program
```

Fortran

1.16 Parallel Random Access Iterator Loop

C++

The following example shows a parallel random access iterator loop.

Example pra_iterator.1.cpp (`omp_3.0`)

```
S-1   #include <vector>
S-2   void iterator_example()
S-3   {
S-4     std::vector<int> vec(23);
S-5     std::vector<int>::iterator it;
S-6   #pragma omp parallel for default(none) shared(vec)
S-7     for (it = vec.begin(); it < vec.end(); it++)
S-8     {
S-9       // do work with *it //
S-10    }
S-11  }
```

C++

1.17 The `omp_set_dynamic` and `omp_set_num_threads` Routines

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 defined, such programs can choose to turn off the dynamic threads capability and set the number of threads explicitly to ensure portability. The following example shows how to do this using **`omp_set_dynamic`**, and **`omp_set_num_threads`**.

In this example, the program executes correctly only if it is executed by 16 threads. If the implementation is not capable of supporting 16 threads, the behavior of this example is implementation defined. Note that the number of threads executing a **`parallel`** region remains constant during the 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.

C / C++

Example set_dynamic_nthrs.1.c

```
S-1   #include <omp.h>
S-2   #include <stdlib.h>
S-3
S-4   void do_by_16(float *x, int iam, int ipoints) {}
S-5
S-6   void dynthreads(float *x, int npoints)
S-7   {
S-8     int iam, ipoints;
S-9
S-10    omp_set_dynamic(0);
S-11    omp_set_num_threads(16);
S-12
S-13    #pragma omp parallel shared(x, npoints) private(iam, ipoints)
S-14    {
S-15      if (omp_get_num_threads() != 16)
S-16        abort();
S-17
S-18      iam = omp_get_thread_num();
S-19      ipoints = npoints/16;
S-20      do_by_16(x, iam, ipoints);
S-21    }
S-22  }
```

C / C++

Fortran

Example set_dynamic_nthrs.1.f

```
S-1         SUBROUTINE DO_BY_16(X, IAM, IPOINTS)
S-2           REAL X(*)
S-3           INTEGER IAM, IPOINTS
S-4         END SUBROUTINE DO_BY_16
S-5
S-6         SUBROUTINE DYNTHREADS(X, NPOINTS)
S-7
S-8           INCLUDE "omp_lib.h"      ! or USE OMP_LIB
S-9
S-10          INTEGER NPOINTS
S-11          REAL X(NPOINTS)
S-12
S-13          INTEGER IAM, IPOINTS
S-14
S-15          CALL OMP_SET_DYNAMIC(.FALSE.)
S-16          CALL OMP_SET_NUM_THREADS(16)
S-17
S-18    !$OMP   PARALLEL SHARED(X,NPOINTS) PRIVATE(IAM, IPOINTS)
S-19
S-20            IF (OMP_GET_NUM_THREADS() .NE. 16) THEN
S-21              STOP
S-22            ENDIF
S-23
S-24            IAM = OMP_GET_THREAD_NUM()
S-25            IPOINTS = NPOINTS/16
S-26            CALL DO_BY_16(X,IAM,IPOINTS)
S-27
S-28    !$OMP   END PARALLEL
S-29
S-30        END SUBROUTINE DYNTHREADS
```

Fortran

1.18 The `omp_get_num_threads` Routine

In the following example, the **omp_get_num_threads** call returns 1 in the sequential part of the code, so **np** will always be equal to 1. To determine the number of threads that will be deployed for the **parallel** region, the call should be inside the **parallel** region.

C / C++

Example get_nthrs.1.c

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

void incorrect() {
  int np, i;

  np = omp_get_num_threads();  /* misplaced */

  #pragma omp parallel for schedule(static)
  for (i=0; i < np; i++)
    work(i);
}
```

C / C++

Fortran

Example get_nthrs.1.f

```
      SUBROUTINE WORK(I)
      INTEGER I
        I = I + 1
      END SUBROUTINE WORK

      SUBROUTINE INCORRECT()
        INCLUDE "omp_lib.h"      ! or USE OMP_LIB
        INTEGER I, NP

        NP = OMP_GET_NUM_THREADS()   !misplaced: will return 1
!$OMP   PARALLEL DO SCHEDULE(STATIC)
          DO I = 0, NP-1
            CALL WORK(I)
          ENDDO
!$OMP   END PARALLEL DO
      END SUBROUTINE INCORRECT
```

Fortran

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

C / C++

Example get_nthrs.2.c

```
S-1   #include <omp.h>
S-2   void work(int i);
S-3
S-4   void correct()
S-5   {
S-6     int i;
S-7
S-8     #pragma omp parallel private(i)
S-9     {
S-10      i = omp_get_thread_num();
S-11      work(i);
S-12    }
S-13  }
```

C / C++

Fortran

Example get_nthrs.2.f

```
S-1         SUBROUTINE WORK(I)
S-2           INTEGER I
S-3
S-4           I = I + 1
S-5
S-6         END SUBROUTINE WORK
S-7
S-8         SUBROUTINE CORRECT()
S-9           INCLUDE "omp_lib.h"     ! or USE OMP_LIB
S-10          INTEGER I
S-11
S-12  !$OMP    PARALLEL PRIVATE(I)
S-13            I = OMP_GET_THREAD_NUM()
S-14            CALL WORK(I)
S-15  !$OMP   END PARALLEL
S-16
S-17        END SUBROUTINE CORRECT
```

Fortran

CHAPTER 2

OpenMP Affinity

OpenMP Affinity consists of a **proc_bind** policy (thread affinity policy) and a specification of places ("location units" or *processors* that may be cores, hardware threads, sockets, etc.). OpenMP Affinity enables users to bind computations on specific places. The placement will hold for the duration of the parallel region. However, the runtime is free to migrate the OpenMP threads to different cores (hardware threads, sockets, etc.) prescribed within a given place, if two or more cores (hardware threads, sockets, etc.) have been assigned to a given place.

Often the binding can be managed without resorting to explicitly setting places. Without the specification of places in the **OMP_PLACES** variable, the OpenMP runtime will distribute and bind threads using the entire range of processors for the OpenMP program, according to the **OMP_PROC_BIND** environment variable or the **proc_bind** clause. When places are specified, the OMP runtime binds threads to the places according to a default distribution policy, or those specified in the **OMP_PROC_BIND** environment variable or the **proc_bind** clause.

In the OpenMP Specifications document a processor refers to an execution unit that is enabled for an OpenMP thread to use. A processor is a core when there is no SMT (Simultaneous Multi-Threading) support or SMT is disabled. When SMT is enabled, a processor is a hardware thread (HW-thread). (This is the usual case; but actually, the execution unit is implementation defined.) Processor numbers are numbered sequentially from 0 to the number of cores less one (without SMT), or 0 to the number HW-threads less one (with SMT). OpenMP places use the processor number to designate binding locations (unless an "abstract name" is used.)

The processors available to a process may be a subset of the system's processors. This restriction may be the result of a wrapper process controlling the execution (such as **numactl** on Linux systems), compiler options, library-specific environment variables, or default kernel settings. For instance, the execution of multiple MPI processes, launched on a single compute node, will each have a subset of processors as determined by the MPI launcher or set by MPI affinity environment variables for the MPI library.

Threads of a team are positioned onto places in a compact manner, a scattered distribution, or onto the master's place, by setting the **OMP_PROC_BIND** environment variable or the **proc_bind**

clause to *close*, *spread*, or *master*, respectively. When **OMP_PROC_BIND** is set to FALSE no binding is enforced; and when the value is TRUE, the binding is implementation defined to a set of places in the **OMP_PLACES** variable or to places defined by the implementation if the **OMP_PLACES** variable is not set.

The **OMP_PLACES** variable can also be set to an abstract name (*threads*, *cores*, *sockets*) to specify that a place is either a single hardware thread, a core, or a socket, respectively. This description of the **OMP_PLACES** is most useful when the number of threads is equal to the number of hardware thread, cores or sockets. It can also be used with a *close* or *spread* distribution policy when the equality doesn't hold.

2.1 The `proc_bind` Clause

The following examples demonstrate how to use the **`proc_bind`** clause to control the thread binding for a team of threads in a **`parallel`** region. The machine architecture is depicted in the figure below. It consists of two sockets, each equipped with a quad-core processor and configured to execute two hardware threads simultaneously on each core. These examples assume a contiguous core numbering starting from 0, such that the hardware threads 0,1 form the first physical core.

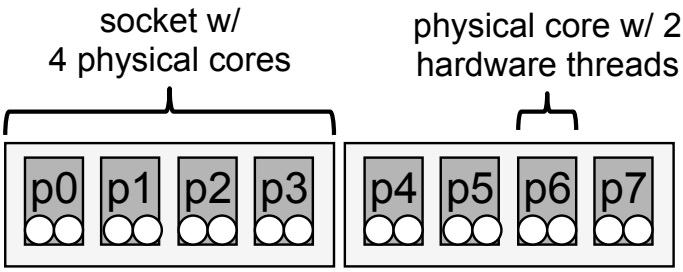

The following equivalent place list declarations consist of eight places (which we designate as p0 to p7):

```
OMP_PLACES="{0,1},{2,3},{4,5},{6,7},{8,9},{10,11},{12,13},{14,15}"
```

or

```
OMP_PLACES="{0:2}:8:2"
```

2.1.1 Spread Affinity Policy

The following example shows the result of the **`spread`** affinity policy on the partition list when the number of threads is less than or equal to the number of places in the parent's place partition, for the machine architecture depicted above. Note that the threads are bound to the first place of each subpartition.

C / C++

Example affinity.1.c (**`omp_4.0`**)

```
S-1
S-2   void work();
S-3
S-4   int main()
S-5   {
S-6
S-7   #pragma omp parallel proc_bind(spread) num_threads(4)
S-8      {
S-9         work();
```

```
S-10      }
S-11
S-12      return 0;
S-13
S-14   }
```

C / C++

Fortran

Example affinity.1.f (`omp_4.0`)

```
S-1         PROGRAM EXAMPLE
S-2   !$OMP PARALLEL PROC_BIND(SPREAD) NUM_THREADS(4)
S-3         CALL WORK()
S-4   !$OMP END PARALLEL
S-5         END PROGRAM EXAMPLE
```

Fortran

It is unspecified on which place the master thread is initially started. If the master thread is initially started on p0, the following placement of threads will be applied in the parallel region:

- thread 0 executes on p0 with the place partition p0,p1
- thread 1 executes on p2 with the place partition p2,p3
- thread 2 executes on p4 with the place partition p4,p5
- thread 3 executes on p6 with the place partition p6,p7

If the master thread would initially be started on p2, the placement of threads and distribution of the place partition would be as follows:

- thread 0 executes on p2 with the place partition p2,p3
- thread 1 executes on p4 with the place partition p4,p5
- thread 2 executes on p6 with the place partition p6,p7
- thread 3 executes on p0 with the place partition p0,p1

The following example illustrates the **`spread`** thread affinity policy when the number of threads is greater than the number of places in the parent's place partition.

Let *T* be the number of threads in the team, and *P* be the number of places in the parent's place partition. The first *T/P* threads of the team (including the master thread) execute on the parent's place. The next *T/P* threads execute on the next place in the place partition, and so on, with wrap around.

C / C++

Example affinity.2.c (`omp_4.0`)

```
S-1  void work();
S-2  void foo()
S-3  {
S-4    #pragma omp parallel num_threads(16) proc_bind(spread)
S-5    {
S-6      work();
S-7    }
S-8  }
```

C / C++

Fortran

Example affinity.2.f90 (`omp_4.0`)

```
S-1  subroutine foo
S-2  !$omp parallel num_threads(16) proc_bind(spread)
S-3        call work()
S-4  !$omp end parallel
S-5  end subroutine
```

Fortran

It is unspecified on which place the master thread is initially started. If the master thread is initially started on p0, the following placement of threads will be applied in the parallel region:

- threads 0,1 execute on p0 with the place partition p0
- threads 2,3 execute on p1 with the place partition p1
- threads 4,5 execute on p2 with the place partition p2
- threads 6,7 execute on p3 with the place partition p3
- threads 8,9 execute on p4 with the place partition p4
- threads 10,11 execute on p5 with the place partition p5
- threads 12,13 execute on p6 with the place partition p6
- threads 14,15 execute on p7 with the place partition p7

If the master thread would initially be started on p2, the placement of threads and distribution of the place partition would be as follows:

- threads 0,1 execute on p2 with the place partition p2
- threads 2,3 execute on p3 with the place partition p3
- threads 4,5 execute on p4 with the place partition p4
- threads 6,7 execute on p5 with the place partition p5
- threads 8,9 execute on p6 with the place partition p6
- threads 10,11 execute on p7 with the place partition p7
- threads 12,13 execute on p0 with the place partition p0
- threads 14,15 execute on p1 with the place partition p1

2.1.2 Close Affinity Policy

The following example shows the result of the **close** affinity policy on the partition list when the number of threads is less than or equal to the number of places in parent's place partition, for the machine architecture depicted above. The place partition is not changed by the **close** policy.

C / C++

Example affinity.3.c (**omp_4.0**)

```
S-1  void work();
S-2  int main()
S-3  {
S-4  #pragma omp parallel proc_bind(close) num_threads(4)
S-5     {
S-6        work();
S-7     }
S-8     return 0;
S-9  }
```

C / C++

Fortran

Example affinity.3.f (**omp_4.0**)

```
S-1        PROGRAM EXAMPLE
S-2  !$OMP PARALLEL PROC_BIND(CLOSE) NUM_THREADS(4)
S-3        CALL WORK()
S-4  !$OMP END PARALLEL
S-5        END PROGRAM EXAMPLE
```

Fortran

It is unspecified on which place the master thread is initially started. If the master thread is initially started on p0, the following placement of threads will be applied in the **parallel** region:

- thread 0 executes on p0 with the place partition p0-p7
- thread 1 executes on p1 with the place partition p0-p7
- thread 2 executes on p2 with the place partition p0-p7
- thread 3 executes on p3 with the place partition p0-p7

If the master thread would initially be started on p2, the placement of threads and distribution of the place partition would be as follows:

- thread 0 executes on p2 with the place partition p0-p7
- thread 1 executes on p3 with the place partition p0-p7
- thread 2 executes on p4 with the place partition p0-p7
- thread 3 executes on p5 with the place partition p0-p7

The following example illustrates the **close** thread affinity policy when the number of threads is greater than the number of places in the parent's place partition.

Let *T* be the number of threads in the team, and *P* be the number of places in the parent's place partition. The first *T/P* threads of the team (including the master thread) execute on the parent's place. The next *T/P* threads execute on the next place in the place partition, and so on, with wrap around. The place partition is not changed by the **close** policy.

C / C++

Example affinity.4.c (**omp_4.0**)

```
S-1  void work();
S-2  void foo()
S-3  {
S-4    #pragma omp parallel num_threads(16) proc_bind(close)
S-5    {
S-6      work();
S-7    }
S-8  }
```

C / C++

Fortran

Example affinity.4.f90 (**omp_4.0**)

```
S-1  subroutine foo
S-2  !$omp parallel num_threads(16) proc_bind(close)
S-3        call work()
S-4  !$omp end parallel
S-5  end subroutine
```

Fortran

It is unspecified on which place the master thread is initially started. If the master thread is initially running on p0, the following placement of threads will be applied in the parallel region:

- threads 0,1 execute on p0 with the place partition p0-p7
- threads 2,3 execute on p1 with the place partition p0-p7
- threads 4,5 execute on p2 with the place partition p0-p7
- threads 6,7 execute on p3 with the place partition p0-p7
- threads 8,9 execute on p4 with the place partition p0-p7
- threads 10,11 execute on p5 with the place partition p0-p7
- threads 12,13 execute on p6 with the place partition p0-p7
- threads 14,15 execute on p7 with the place partition p0-p7

If the master thread would initially be started on p2, the placement of threads and distribution of the place partition would be as follows:

- threads 0,1 execute on p2 with the place partition p0-p7

- threads 2,3 execute on p3 with the place partition p0-p7
- threads 4,5 execute on p4 with the place partition p0-p7
- threads 6,7 execute on p5 with the place partition p0-p7
- threads 8,9 execute on p6 with the place partition p0-p7
- threads 10,11 execute on p7 with the place partition p0-p7
- threads 12,13 execute on p0 with the place partition p0-p7
- threads 14,15 execute on p1 with the place partition p0-p7

2.1.3 Master Affinity Policy

The following example shows the result of the **master** affinity policy on the partition list for the machine architecture depicted above. The place partition is not changed by the master policy.

C / C++

Example affinity.5.c (**omp_4.0**)

```
S-1  void work();
S-2  int main()
S-3  {
S-4  #pragma omp parallel proc_bind(master) num_threads(4)
S-5     {
S-6        work();
S-7     }
S-8     return 0;
S-9  }
```

C / C++

Fortran

Example affinity.5.f (**omp_4.0**)

```
S-1        PROGRAM EXAMPLE
S-2  !$OMP PARALLEL PROC_BIND(MASTER) NUM_THREADS(4)
S-3        CALL WORK()
S-4  !$OMP END PARALLEL
S-5        END PROGRAM EXAMPLE
```

Fortran

It is unspecified on which place the master thread is initially started. If the master thread is initially running on p0, the following placement of threads will be applied in the parallel region:

- threads 0-3 execute on p0 with the place partition p0-p7

If the master thread would initially be started on p2, the placement of threads and distribution of the place partition would be as follows:

- threads 0-3 execute on p2 with the place partition p0-p7

2.2 Task Affinity

The next example illustrates the use of the **affinity** clause with a **task** construct. The variables in the **affinity** clause provide a hint to the runtime that the task should execute "close" to the physical storage location of the variables. For example, on a two-socket platform with a local memory component close to each processor socket, the runtime will attempt to schedule the task execution on the socket where the storage is located.

Because the C/C++ code employs a pointer, an array section is used in the **affinity** clause. Fortran code can use an array reference to specify the storage, as shown here.

Note, in the second task of the C/C++ code the *B* pointer is declared shared. Otherwise, by default, it would be firstprivate since it is a local variable, and would probably be saved for the second task before being assigned a storage address by the first task. Also, one might think it reasonable to use the **affinity** clause *affinity(B[:N])* on the second **task** construct. However, the storage behind *B* is created in the first task, and the array section reference may not be valid when the second task is generated. The use of the *A* array is sufficient for this case, because one would expect the storage for *A* and *B* would be physically "close" (as provided by the hint in the first task).

C / C++

Example affinity.6.c (**omp_5.0**)

```
S-1
S-2   double * alloc_init_B(double *A, int N);
S-3   void     compute_on_B(double *B, int N);
S-4
S-5   void task_affinity(double *A, int N)
S-6   {
S-7      double * B;
S-8      #pragma omp task depend(out:B) shared(B) affinity(A[0:N])
S-9      {
S-10       B = alloc_init_B(A,N);
S-11     }
S-12
S-13     #pragma omp task depend( in:B) shared(B) affinity(A[0:N])
S-14     {
S-15       compute_on_B(B,N);
S-16     }
S-17
S-18     #pragma omp taskwait
S-19  }
S-20
```

C / C++

Fortran

Example affinity.6.f90 (**omp_5.0**)

```
S-1
S-2   subroutine task_affinity(A, N)
S-3
S-4     external alloc_init_B
S-5     external compute_on_B
S-6     double precision, allocatable :: B(:)
S-7
S-8      !$omp task depend(out:B) shared(B) affinity(A)
S-9        call alloc_init_B(B,A)
S-10     !$omp end task
S-11
S-12     !$omp task depend(in:B) shared(B) affinity(A)
S-13       call compute_on_B(B)
S-14     !$omp end task
S-15
S-16     !$omp taskwait
S-17
S-18  end subroutine
S-19
```

Fortran

2.3 Affinity Display

The following examples illustrate ways to display thread affinity. Automatic display of affinity can be invoked by setting the **OMP_DISPLAY_AFFINITY** environment variable to **TRUE**. The format of the output can be customized by setting the **OMP_AFFINITY_FORMAT** environment variable to an appropriate string. Also, there are API calls for the user to display thread affinity at selected locations within code.

For the first example the environment variable **OMP_DISPLAY_AFFINITY** has been set to **TRUE**, and execution occurs on an 8-core system with **OMP_NUM_THREADS** set to 8.

The affinity for the master thread is reported through a call to the API **omp_display_affinity()** routine. For default affinity settings the report shows that the master thread can execute on any of the cores. In the following parallel region the affinity for each of the team threads is reported automatically since the **OMP_DISPLAY_AFFINITY** environment variable has been set to **TRUE**.

These two reports are often useful (as in hybrid codes using both MPI and OpenMP) to observe the affinity (for an MPI task) before the parallel region, and during an OpenMP parallel region. Note:

the next parallel region uses the same number of threads as in the previous parallel region and affinities are not changed, so affinity is NOT reported.

In the last parallel region, the thread affinities are reported because the thread affinity has changed.

C / C++

Example affinity_display.1.c (`omp_5.0`)

```
#include <stdio.h>
#include <omp.h>

int main(void){                        //MAX threads = 8, single socket system

   omp_display_affinity(NULL);  //API call-- Displays Affinity of Master Thread

// API CALL OUTPUT (default format):
//team_num= 0, nesting_level= 0, thread_num= 0, thread_affinity= 0,1,2,3,4,5,6,7


                    // OMP_DISPLAY_AFFINITY=TRUE, OMP_NUM_THREADS=8
   #pragma omp parallel num_threads(omp_get_num_procs())
   {
     if(omp_get_thread_num()==0)
        printf("1st Parallel Region -- Affinity Reported \n");

         // DISPLAY OUTPUT (default format) has been sorted:
         // team_num= 0, nesting_level= 1, thread_num= 0, thread_affinity= 0
         // team_num= 0, nesting_level= 1, thread_num= 1, thread_affinity= 1
         // ...
         // team_num= 0, nesting_level= 1, thread_num= 7, thread_affinity= 7

     // doing work here
   }

   #pragma omp parallel num_threads( omp_get_num_procs() )
   {
     if(omp_get_thread_num()==0)
        printf("%s%s\n","Same Affinity as in Previous Parallel Region",
                        " -- no Affinity Reported\n");

         // NO AFFINITY OUTPUT:
         //(output in 1st parallel region only for OMP_DISPLAY_AFFINITY=TRUE)

     // doing more work here
   }

                       // Report Affinity for 1/2 number of threads
   #pragma omp parallel num_threads( omp_get_num_procs()/2 )
```

```
S-41        {
S-42          if(omp_get_thread_num()==0)
S-43             printf("Report Affinity for using 1/2 of max threads.\n");
S-44
S-45              // DISPLAY OUTPUT (default format) has been sorted:
S-46              // team_num= 0, nesting_level= 1, thread_num= 0, thread_affinity= 0,1
S-47              // team_num= 0, nesting_level= 1, thread_num= 1, thread_affinity= 2,3
S-48              // team_num= 0, nesting_level= 1, thread_num= 2, thread_affinity= 4,5
S-49              // team_num= 0, nesting_level= 1, thread_num= 3, thread_affinity= 6,7
S-50
S-51          // do work
S-52        }
S-53
S-54        return 0;
S-55    }
```

C / C++

Fortran

Example affinity_display.1.f90 (`omp_5.0`)

```
S-1
S-2     program affinity_display        ! MAX threads = 8, single socket system
S-3
S-4        use omp_lib
S-5        implicit none
S-6        character(len=0) :: null
S-7
S-8        call omp_display_affinity(null) !API call- Displays Affinity of Master Thread
S-9
S-10    ! API CALL OUTPUT (default format):
S-11    ! team_num= 0, nesting_level= 0, thread_num= 0, thread_affinity= 0,1,2,3,4,5,6,7
S-12
S-13
S-14                            ! OMP_DISPLAY_AFFINITY=TRUE, OMP_NUM_THREADS=8
S-15
S-16       !$omp parallel num_threads(omp_get_num_procs())
S-17
S-18         if(omp_get_thread_num()==0) then
S-19            print*, "1st Parallel Region  -- Affinity Reported"
S-20         endif
S-21
S-22              ! DISPLAY OUTPUT (default format) has been sorted:
S-23              ! team_num= 0, nesting_level= 1, thread_num= 0, thread_affinity= 0
S-24              ! team_num= 0, nesting_level= 1, thread_num= 1, thread_affinity= 1
S-25              ! ...
S-26              ! team_num= 0, nesting_level= 1, thread_num= 7, thread_affinity= 7
S-27
```

```
S-28          ! doing work here
S-29
S-30        !$omp end parallel
S-31
S-32        !$omp parallel num_threads( omp_get_num_procs() )
S-33
S-34          if(omp_get_thread_num()==0) then
S-35             print*, "Same Affinity in Parallel Region -- no Affinity Reported"
S-36          endif
S-37
S-38               ! NO AFFINITY OUTPUT:
S-39               !(output in 1st parallel region only for OMP_DISPLAY_AFFINITY=TRUE)
S-40
S-41           ! doing more work here
S-42
S-43        !$omp end parallel
S-44
S-45                            ! Report Affinity for 1/2 number of threads
S-46        !$omp parallel num_threads( omp_get_num_procs()/2 )
S-47
S-48          if(omp_get_thread_num()==0) then
S-49             print*, "Different Affinity in Parallel Region -- Affinity Reported"
S-50          endif
S-51
S-52               ! DISPLAY OUTPUT (default format) has been sorted:
S-53               ! team_num= 0, nesting_level= 1, thread_num= 0, thread_affinity= 0,1
S-54               ! team_num= 0, nesting_level= 1, thread_num= 1, thread_affinity= 2,3
S-55               ! team_num= 0, nesting_level= 1, thread_num= 2, thread_affinity= 4,5
S-56               ! team_num= 0, nesting_level= 1, thread_num= 3, thread_affinity= 6,7
S-57
S-58           ! do work
S-59
S-60        !$omp end parallel
S-61
S-62    end program
```

Fortran

In the following example 2 threads are forked, and each executes on a socket. Next, a nested parallel region runs half of the available threads on each socket.

These OpenMP environment variables have been set:

- **OMP_PROC_BIND**="TRUE"
- **OMP_NUM_THREADS**="2,4"
- **OMP_PLACES**="{0,2,4,6},{1,3,5,7}"
- **OMP_AFFINITY_FORMAT**="nest_level= %L, parent_thrd_num= %a, thrd_num= %n, thrd_affinity= %A"

where the numbers correspond to core ids for the system. Note, **OMP_DISPLAY_AFFINITY** is not set and is **FALSE** by default. This example shows how to use API routines to perform affinity display operations.

For each of the two first-level threads the **OMP_PLACES** variable specifies a place with all the core-ids of the socket ({0,2,4,6} for one thread and {1,3,5,7} for the other). (As is sometimes the case in 2-socket systems, one socket may consist of the even id numbers, while the other may have the odd id numbers.) The affinities are printed according to the **OMP_AFFINITY_FORMAT** format: providing the parallel nesting level (%L), the ancestor thread number (%a), the thread number (%n) and the thread affinity (%A). In the nested parallel region within the *socket_work* routine the affinities for the threads on each socket are printed according to this format.

C / C++

Example affinity_display.2.c (**omp_5.0**)

```
#include <stdio.h>
#include <stdlib.h>
#include <omp.h>

void socket_work(int socket_num, int n_thrds);

int main(void)
{
 int n_sockets, socket_num, n_thrds_on_socket;

 omp_set_nested(1);             // or env var= OMP_NESTED=true
 omp_set_max_active_levels(2); // or env var= OMP_MAX_ACTIVE_LEVELS=2

  n_sockets         = omp_get_num_places();
  n_thrds_on_socket = omp_get_place_num_procs(0);

     // OMP_NUM_THREADS=2,4
     // OMP_PLACES="{0,2,4,6},{1,3,5,7}"  #2 sockets; even/odd proc-ids
     // OMP_AFFINITY_FORMAT=\
     // "nest_level= %L, parent_thrd_num= %a, thrd_num= %n, thrd_affinity= %A"

  #pragma omp parallel num_threads(n_sockets) private(socket_num)
  {
    socket_num = omp_get_place_num();

    if(socket_num==0)
       printf(" LEVEL 1 AFFINITIES 1 thread/socket, %d sockets:\n\n", n_sockets);

    omp_display_affinity(NULL);   // not needed if OMP_DISPLAY_AFFINITY=TRUE

     // OUTPUT:
     // LEVEL 1 AFFINITIES 1 thread/socket, 2 sockets:
```

```
S-33        // nest_level= 1, parent_thrd_num= 0, thrd_num= 0, thrd_affinity= 0,2,4,6
S-34        // nest_level= 1, parent_thrd_num= 0, thrd_num= 1, thrd_affinity= 1,3,5,7
S-35
S-36      socket_work(socket_num, n_thrds_on_socket);
S-37    }
S-38
S-39    return 0;
S-40  }
S-41
S-42  void socket_work(int socket_num, int n_thrds)
S-43  {
S-44    #pragma omp parallel num_threads(n_thrds)
S-45    {
S-46      if(omp_get_thread_num()==0)
S-47    printf(" LEVEL 2 AFFINITIES, %d threads on socket %d\n",n_thrds, socket_num);
S-48
S-49      omp_display_affinity(NULL);   // not needed if OMP_DISPLAY_AFFINITY=TRUE
S-50
S-51      // OUTPUT:
S-52      // LEVEL 2 AFFINITIES, 4 threads on socket 0
S-53      // nest_level= 2, parent_thrd_num= 0, thrd_num= 0, thrd_affinity= 0
S-54      // nest_level= 2, parent_thrd_num= 0, thrd_num= 1, thrd_affinity= 2
S-55      // nest_level= 2, parent_thrd_num= 0, thrd_num= 2, thrd_affinity= 4
S-56      // nest_level= 2, parent_thrd_num= 0, thrd_num= 3, thrd_affinity= 6
S-57
S-58      // LEVEL 2 AFFINITIES, 4 threads on socket 1
S-59      // nest_level= 2, parent_thrd_num= 1, thrd_num= 0, thrd_affinity= 1
S-60      // nest_level= 2, parent_thrd_num= 1, thrd_num= 1, thrd_affinity= 3
S-61      // nest_level= 2, parent_thrd_num= 1, thrd_num= 2, thrd_affinity= 5
S-62      // nest_level= 2, parent_thrd_num= 1, thrd_num= 3, thrd_affinity= 7
S-63
S-64      // ... Do Some work on Socket
S-65
S-66    }
S-67  }
```

C / C++

Fortran

Example affinity_display.2.f90 (`omp_5.0`)

```
S-1
S-2   program affinity_display
S-3
S-4      use omp_lib
S-5      implicit none
S-6      character(len=0) :: null
S-7      integer          :: n_sockets, socket_num, n_thrds_on_socket;
```

```
S-8
S-9       call omp_set_nested(.true.)        ! or env var= OMP_NESTED=true
S-10      call omp_set_max_active_levels(2)  ! or env var= OMP_MAX_ACTIVE_LEVELS=2
S-11
S-12      n_sockets         = omp_get_num_places()
S-13      n_thrds_on_socket = omp_get_place_num_procs(0)
S-14
S-15       ! OMP_NUM_THREADS=2,4
S-16       ! OMP_PLACES="{0,2,4,6},{1,3,5,7}"  #2 sockets; even/odd proc-ids
S-17       ! OMP_AFFINITY_FORMAT=\
S-18       !     "nest_level= %L, parent_thrd_num= %a, thrd_num= %n, thrd_affinity= %A"
S-19
S-20      !$omp parallel num_threads(n_sockets) private(socket_num)
S-21
S-22        socket_num = omp_get_place_num()
S-23
S-24        if(socket_num==0) then
S-25          write(*,'("LEVEL 1 AFFINITIES 1 thread/socket ",i0," sockets")')n_sockets
S-26        endif
S-27
S-28        call omp_display_affinity(null)  !not needed if OMP_DISPLAY_AFFINITY=TRUE
S-29
S-30          ! OUTPUT:
S-31          ! LEVEL 1 AFFINITIES 1 thread/socket, 2 sockets:
S-32          ! nest_level= 1, parent_thrd_num= 0, thrd_num= 0, thrd_affinity= 0,2,4,6
S-33          ! nest_level= 1, parent_thrd_num= 0, thrd_num= 1, thrd_affinity= 1,3,5,7
S-34
S-35        call socket_work(socket_num, n_thrds_on_socket)
S-36
S-37      !$omp end parallel
S-38
S-39    end program
S-40
S-41    subroutine socket_work(socket_num, n_thrds)
S-42       use omp_lib
S-43       implicit none
S-44       integer :: socket_num, n_thrds
S-45       character(len=0) :: null
S-46
S-47       !$omp parallel num_threads(n_thrds)
S-48
S-49          if(omp_get_thread_num()==0) then
S-50          write(*,'("LEVEL 2 AFFINITIES, ",i0," threads on socket ",i0)') &
S-51                n_thrds,socket_num
S-52          endif
S-53
S-54          call omp_display_affinity(null);  !not needed if OMP_DISPLAY_AFFINITY=TRUE
```

```
S-55
S-56          ! OUTPUT:
S-57          ! LEVEL 2 AFFINITIES, 4 threads on socket 0
S-58          ! nest_level= 2, parent_thrd_num= 0, thrd_num= 0, thrd_affinity= 0
S-59          ! nest_level= 2, parent_thrd_num= 0, thrd_num= 1, thrd_affinity= 2
S-60          ! nest_level= 2, parent_thrd_num= 0, thrd_num= 2, thrd_affinity= 4
S-61          ! nest_level= 2, parent_thrd_num= 0, thrd_num= 3, thrd_affinity= 6
S-62
S-63          ! LEVEL 2 AFFINITIES, 4 thrds on socket 1
S-64          ! nest_level= 2, parent_thrd_num= 1, thrd_num= 0, thrd_affinity= 1
S-65          ! nest_level= 2, parent_thrd_num= 1, thrd_num= 1, thrd_affinity= 3
S-66          ! nest_level= 2, parent_thrd_num= 1, thrd_num= 2, thrd_affinity= 5
S-67          ! nest_level= 2, parent_thrd_num= 1, thrd_num= 3, thrd_affinity= 7
S-68
S-69          ! ... Do Some work on Socket
S-70
S-71       !$omp end parallel
S-72
S-73    end subroutine
```

Fortran

The next example illustrates more details about affinity formatting. First, the **omp_get_affininity_format()** API routine is used to obtain the default format. The code checks to make sure the storage provides enough space to hold the format. Next, the **omp_set_affinity_format()** API routine sets a user-defined format: *host=%20H thrd_num=%0.4n binds_to=%A*.

The host, thread number and affinity fields are specified by *%20H*, *%0.4n* and *%A*: *H*, *n* and *A* are single character "short names" for the host, thread_num and thread_affinity data to be printed, with format sizes of *20*, *4*, and "size as needed". The period (.) indicates that the field is displayed right-justified (default is left-justified) and the "0" indicates that any unused space is to be prefixed with zeros (e.g. instead of "1", "0001" is displayed for the field size of 4).

Within the parallel region the affinity for each thread is captured by **omp_capture_affinity()** into a buffer array with elements indexed by the thread number (*thrd_num*). After the parallel region, the thread affinities are printed in thread-number order.

If the storage area in buffer is inadequate for holding the affinity data, the stored affinity data is truncated. The maximum value for the number of characters (*nchars*) returned by **omp_capture_affinity** is captured by the **reduction(max:max_req_store)** clause and the *if(nchars >= max_req_store) max_req_store=nchars* statement. It is used to report possible truncation (if *max_req_store* > *buffer_store*).

Example affinity_display.3.c (`omp_5.0`)

```
#include <stdio.h>
#include <stdlib.h>  // also null is in <stddef.h>
#include <stddef.h>
#include <string.h>
#include <omp.h>

#define FORMAT_STORE   80
#define BUFFER_STORE   80

int main(void){

   int i, n, thrd_num, max_req_store;
   size_t nchars;

   char default_format[FORMAT_STORE];
   char my_format[]  = "host=%20H thrd_num=%0.4n binds_to=%A";
   char **buffer;


// CODE SEGMENT 1         AFFINITY FORMAT

//                        Get and Display Default Affinity Format

   nchars = omp_get_affinity_format(default_format,(size_t)FORMAT_STORE);
   printf("Default Affinity Format is: %s\n",default_format);

   if(nchars >= FORMAT_STORE){
      printf("Caution: Reported Format is truncated.  Increase\n");
      printf("         FORMAT_STORE to %d.\n", nchars+1);
   }

//                        Set Affinity Format

   omp_set_affinity_format(my_format);
   printf("Affinity Format set to: %s\n",my_format);


// CODE SEGMENT 2         CAPTURE AFFINITY

//                        Set up buffer for affinity of n threads

   n = omp_get_num_procs();
   buffer = (char **)malloc( sizeof(char *) * n );
   for(i=0;i<n;i++){ buffer[i]=(char *)malloc( sizeof(char) * BUFFER_STORE); }
```

```
S-45
S-46    //                          Capture Affinity using Affinity Format set above.
S-47    //                          Use max reduction to check size of buffer areas
S-48       max_req_store = 0;
S-49       #pragma omp parallel private(thrd_num,nchars) reduction(max:max_req_store)
S-50       {
S-51          if(omp_get_num_threads()>n) exit(1); //safety: don't exceed # of buffers
S-52
S-53          thrd_num=omp_get_thread_num();
S-54          nchars=omp_capture_affinity(buffer[thrd_num],(size_t)BUFFER_STORE,NULL);
S-55          if(nchars > max_req_store) max_req_store=nchars;
S-56
S-57          // ...
S-58       }
S-59
S-60       for(i=0;i<n;i++){
S-61          printf("thrd_num= %d, affinity: %s\n", i,buffer[i]);
S-62       }
S-63             // For 4 threads with OMP_PLACES='{0,1},{2,3},{4,5},{6,7}'
S-64             // Format    host=%20H thrd_num=%0.4n binds_to=%A
S-65
S-66             // affinity: host=hpc.cn567            thrd_num=0000 binds_to=0,1
S-67             // affinity: host=hpc.cn567            thrd_num=0001 binds_to=2,3
S-68             // affinity: host=hpc.cn567            thrd_num=0002 binds_to=4,5
S-69             // affinity: host=hpc.cn567            thrd_num=0003 binds_to=6,7
S-70
S-71
S-72       if(max_req_store>=BUFFER_STORE){
S-73          printf("Caution: Affinity string truncated.  Increase\n");
S-74          printf("         BUFFER_STORE to %d\n",max_req_store+1);
S-75       }
S-76
S-77       for(i=0;i<n;i++) free(buffer[i]);
S-78       free (buffer);
S-79
S-80       return 0;
S-81    }
```

C / C++

Fortran

Example affinity_display.3.f90 (`omp_5.0`)

```
S-1
S-2     program affinity_display
S-3        use omp_lib
S-4        implicit none
S-5        integer, parameter :: FORMAT_STORE=80
S-6        integer, parameter :: BUFFER_STORE=80
S-7
S-8        integer            :: i, n, thrd_num, nchars, max_req_store
S-9
S-10       character(FORMAT_STORE)     :: default_format
S-11       character(*), parameter     :: my_format = &
S-12                                      "host=%20H thrd_num=%0.4n binds_to=%A"
S-13       character(:), allocatable   :: buffer(:)
S-14       character(len=0)            :: null
S-15
S-16
S-17    !  CODE SEGMENT 1         AFFINITY FORMAT
S-18
S-19    !                         Get and Display Default Affinity Format
S-20
S-21       nchars = omp_get_affinity_format(default_format)
S-22       print*,"Default Affinity Format: ", trim(default_format)
S-23
S-24       if( nchars > FORMAT_STORE) then
S-25          print*,"Caution: Reported Format is truncated.  Increase"
S-26          print*,"         FORMAT_STORE to ", nchars
S-27       endif
S-28
S-29    !                         Set Affinity Format
S-30
S-31       call omp_set_affinity_format(my_format)
S-32       print*,"Affinity Format set to: ", my_format
S-33
S-34
S-35    !  CODE SEGMENT 2         CAPTURE AFFINITY
S-36
S-37    !                         Set up buffer for affinity of n threads
S-38
S-39       n = omp_get_num_procs()
S-40       allocate( character(len=BUFFER_STORE)::buffer(0:n-1) )
S-41
S-42    !                         Capture Affinity using Affinity Format set above.
S-43    !                         Use max reduction to check size of buffer areas
S-44       max_req_store = 0
```

```
S-45        !$omp parallel private(thrd_num,nchars) reduction(max:max_req_store)
S-46
S-47           if(omp_get_num_threads()>n) stop "ERROR: increase buffer lines"
S-48
S-49           thrd_num=omp_get_thread_num()
S-50           nchars=omp_capture_affinity(buffer(thrd_num),null)
S-51           if(nchars>max_req_store) max_req_store=nchars
S-52           !  ...
S-53
S-54        !$omp end parallel
S-55
S-56        do i = 0, n-1
S-57           print*, "thrd_num= ",i,"   affinity:", trim(buffer(i))
S-58        end do
S-59              !  For 4 threads with OMP_PLACES='{0,1},{2,3},{4,5},{6,7}'
S-60              !  Format:   host=%20H thrd_num=%0.4n binds_to=%A
S-61
S-62              !  affinity: host=hpc.cn567            thrd_num=0000 binds_to=0,1
S-63              !  affinity: host=hpc.cn567            thrd_num=0001 binds_to=2,3
S-64              !  affinity: host=hpc.cn567            thrd_num=0002 binds_to=4,5
S-65              !  affinity: host=hpc.cn567            thrd_num=0003 binds_to=6,7
S-66
S-67        if(max_req_store > BUFFER_STORE) then
S-68           print*,  "Caution: Affinity string truncated.  Increase"
S-69           print*,  "         BUFFER_STORE to ",max_req_store
S-70        endif
S-71
S-72        deallocate(buffer)
S-73     end program
```

Fortran

2.4 Affinity Query Functions

In the example below a team of threads is generated on each socket of the system, using nested parallelism. Several query functions are used to gather information to support the creation of the teams and to obtain socket and thread numbers.

For proper execution of the code, the user must create a place partition, such that each place is a listing of the core numbers for a socket. For example, in a 2 socket system with 8 cores in each socket, and sequential numbering in the socket for the core numbers, the **OMP_PLACES** variable would be set to "{0:8},{8:8}", using the place syntax {*lower_bound*:*length*:*stride*}, and the default stride of 1.

The code determines the number of sockets (*n_sockets*) using the **omp_get_num_places()** query function. In this example each place is constructed with a list of each socket's core numbers, hence the number of places is equal to the number of sockets.

The outer parallel region forms a team of threads, and each thread executes on a socket (place) because the **proc_bind** clause uses **spread** in the outer **parallel** construct. Next, in the *socket_init* function, an inner parallel region creates a team of threads equal to the number of elements (core numbers) from the place of the parent thread. Because the outer **parallel** construct uses a **spread** affinity policy, each of its threads inherits a subpartition of the original partition. Hence, the **omp_get_place_num_procs** query function returns the number of elements (here procs = cores) in the subpartition of the thread. After each parent thread creates its nested parallel region on the section, the socket number and thread number are reported.

Note: Portable tools like hwloc (Portable HardWare LOCality package), which support many common operating systems, can be used to determine the configuration of a system. On some systems there are utilities, files or user guides that provide configuration information. For instance, the socket number and proc_id's for a socket can be found in the /proc/cpuinfo text file on Linux systems.

C / C++

Example affinity_query.1.c (**omp_4.5**)

```
#include <stdio.h>
#include <omp.h>

void socket_init(int socket_num)
{
   int n_procs;

   n_procs = omp_get_place_num_procs(socket_num);
   #pragma omp parallel num_threads(n_procs) proc_bind(close)
   {
      printf("Reporting in from socket num, thread num:  %d %d\n",
                                  socket_num,omp_get_thread_num() );
   }
}

int main()
{
   int n_sockets, socket_num;

   omp_set_nested(1);              // or export OMP_NESTED=true
   omp_set_max_active_levels(2);   // or export OMP_MAX_ACTIVE_LEVELS=2

   n_sockets = omp_get_num_places();
   #pragma omp parallel num_threads(n_sockets) private(socket_num) \
                        proc_bind(spread)
```

```
S-26        {
S-27           socket_num = omp_get_place_num();
S-28           socket_init(socket_num);
S-29        }
S-30
S-31        return 0;
S-32     }
```

C / C++

Fortran

1 *Example affinity_query.1.f90* (`omp_4.5`)

```
S-1
S-2      subroutine socket_init(socket_num)
S-3         use omp_lib
S-4         integer  :: socket_num, n_procs
S-5
S-6         n_procs = omp_get_place_num_procs(socket_num)
S-7         !$omp parallel num_threads(n_procs) proc_bind(close)
S-8
S-9            print*,"Reporting in from socket num, thread num: ",  &
S-10                                  socket_num,omp_get_thread_num()
S-11        !$omp end parallel
S-12     end subroutine
S-13
S-14     program numa_teams
S-15        use omp_lib
S-16        integer :: n_sockets, socket_num
S-17
S-18        call omp_set_nested(.true.)             ! or export OMP_NESTED=true
S-19        call omp_set_max_active_levels(2) ! or export OMP_MAX_ACTIVE_LEVELS=2
S-20
S-21        n_sockets = omp_get_num_places()
S-22        !$omp parallel num_threads(n_sockets) private(socket_num) &
S-23        !$omp&         proc_bind(spread)
S-24
S-25           socket_num = omp_get_place_num()
S-26           call socket_init(socket_num)
S-27
S-28        !$omp end parallel
S-29     end program
```

Fortran

CHAPTER 3

Tasking

Tasking constructs provide units of work to a thread for execution. Worksharing constructs do this, too (e.g. **for**, **do**, **sections**, and **singles** constructs); but the work units are tightly controlled by an iteration limit and limited scheduling, or a limited number of **sections** or **single** regions. Worksharing was designed with "data parallel" computing in mind. Tasking was designed for "task parallel" computing and often involves non-locality or irregularity in memory access.

The **task** construct can be used to execute work chunks: in a while loop; while traversing nodes in a list; at nodes in a tree graph; or in a normal loop (with a **taskloop** construct). Unlike the statically scheduled loop iterations of worksharing, a task is often enqueued, and then dequeued for execution by any of the threads of the team within a parallel region. The generation of tasks can be from a single generating thread (creating sibling tasks), or from multiple generators in a recursive graph tree traversals. A **taskloop** construct bundles iterations of an associated loop into tasks, and provides similar controls found in the **task** construct.

Sibling tasks are synchronized by the **taskwait** construct, and tasks and their descendent tasks can be synchronized by containing them in a **taskgroup** region. Ordered execution is accomplished by specifying dependences with a **depend** clause. Also, priorities can be specified as hints to the scheduler through a **priority** clause.

Various clauses can be used to manage and optimize task generation, as well as reduce the overhead of execution and to relinquish control of threads for work balance and forward progress.

Once a thread starts executing a task, it is the designated thread for executing the task to completion, even though it may leave the execution at a scheduling point and return later. The thread is tied to the task. Scheduling points can be introduced with the **taskyield** construct. With an **untied** clause any other thread is allowed to continue the task. An **if** clause with an expression that evaluates to *false* results in an *undeferred* task, which instructs the runtime to suspend the generating task until the undeferred task completes its execution. By including the data environment of the generating task into the generated task with the **mergeable** and **final** clauses, task generation overhead can be reduced.

A complete list of the tasking constructs and details of their clauses can be found in the *Tasking Constructs* chapter of the OpenMP Specifications, in the *OpenMP Application Programming Interface* section.

3.1 The `task` and `taskwait` Constructs

The following example shows how to traverse a tree-like structure using explicit tasks. Note that the **`traverse`** function should be called from within a parallel region for the different specified tasks to be executed in parallel. Also note that the tasks will be executed in no specified order because there are no synchronization directives. Thus, assuming that the traversal will be done in post order, as in the sequential code, is wrong.

C / C++

Example tasking.1.c (`omp_3.0`)

```
S-1
S-2   struct node {
S-3     struct node *left;
S-4     struct node *right;
S-5   };
S-6
S-7   extern void process(struct node *);
S-8
S-9   void traverse( struct node *p )
S-10  {
S-11    if (p->left)
S-12  #pragma omp task   // p is firstprivate by default
S-13        traverse(p->left);
S-14    if (p->right)
S-15  #pragma omp task    // p is firstprivate by default
S-16        traverse(p->right);
S-17    process(p);
S-18  }
```

C / C++

Fortran

Example tasking.1.f90 (`omp_3.0`)

```
S-1
S-2           RECURSIVE SUBROUTINE traverse ( P )
S-3              TYPE Node
S-4                 TYPE(Node), POINTER :: left, right
S-5              END TYPE Node
S-6              TYPE(Node) :: P
S-7
S-8              IF (associated(P%left)) THEN
S-9                 !$OMP TASK     ! P is firstprivate by default
S-10                    CALL traverse(P%left)
S-11                !$OMP END TASK
S-12             ENDIF
S-13             IF (associated(P%right)) THEN
```

```
S-14              !$OMP TASK     ! P is firstprivate by default
S-15                 CALL traverse(P%right)
S-16              !$OMP END TASK
S-17           ENDIF
S-18           CALL process ( P )
S-19
S-20        END SUBROUTINE
```

Fortran

In the next example, we force a postorder traversal of the tree by adding a **taskwait** directive. Now, we can safely assume that the left and right sons have been executed before we process the current node.

C / C++

Example tasking.2.c (**omp_3.0**)

```
S-1   struct node {
S-2     struct node *left;
S-3     struct node *right;
S-4   };
S-5   extern void process(struct node *);
S-6   void postorder_traverse( struct node *p ) {
S-7       if (p->left)
S-8          #pragma omp task    // p is firstprivate by default
S-9              postorder_traverse(p->left);
S-10      if (p->right)
S-11         #pragma omp task   // p is firstprivate by default
S-12             postorder_traverse(p->right);
S-13      #pragma omp taskwait
S-14      process(p);
S-15  }
```

C / C++

Fortran

Example tasking.2.f90 (`omp_3.0`)

```
      RECURSIVE SUBROUTINE traverse ( P )
         TYPE Node
             TYPE(Node), POINTER :: left, right
         END TYPE Node
         TYPE(Node) :: P
         IF (associated(P%left)) THEN
             !$OMP TASK    ! P is firstprivate by default
                 CALL traverse(P%left)
             !$OMP END TASK
         ENDIF
         IF (associated(P%right)) THEN
             !$OMP TASK    ! P is firstprivate by default
                 CALL traverse(P%right)
             !$OMP END TASK
         ENDIF
         !$OMP TASKWAIT
         CALL process ( P )
      END SUBROUTINE
```

Fortran

The following example demonstrates how to use the **`task`** construct to process elements of a linked list in parallel. The thread executing the **`single`** region generates all of the explicit tasks, which are then executed by the threads in the current team. The pointer *p* is **`firstprivate`** by default on the **`task`** construct so it is not necessary to specify it in a **`firstprivate`** clause.

C / C++

Example tasking.3.c (`omp_3.0`)

```

typedef struct node node;
struct node {
      int data;
      node * next;
};

void process(node * p)
{
    /* do work here */
}

void increment_list_items(node * head)
{
    #pragma omp parallel
    {
```

```
S-17            #pragma omp single
S-18                {
S-19                   node * p = head;
S-20                   while (p) {
S-21                        #pragma omp task
S-22                         // p is firstprivate by default
S-23                               process(p);
S-24                         p = p->next;
S-25                      }
S-26                }
S-27         }
S-28    }
```

C / C++

Fortran

1 *Example tasking.3.f90* (`omp_3.0`)

```
S-1
S-2         MODULE LIST
S-3             TYPE NODE
S-4                INTEGER :: PAYLOAD
S-5                TYPE (NODE), POINTER :: NEXT
S-6             END TYPE NODE
S-7         CONTAINS
S-8
S-9             SUBROUTINE PROCESS(p)
S-10               TYPE (NODE), POINTER :: P
S-11                   ! do work here
S-12            END SUBROUTINE
S-13
S-14            SUBROUTINE INCREMENT_LIST_ITEMS (HEAD)
S-15
S-16               TYPE (NODE), POINTER :: HEAD
S-17               TYPE (NODE), POINTER :: P
S-18               !$OMP PARALLEL PRIVATE(P)
S-19                   !$OMP SINGLE
S-20                       P => HEAD
S-21                       DO
S-22                          !$OMP TASK
S-23                              ! P is firstprivate by default
S-24                              CALL PROCESS(P)
S-25                          !$OMP END TASK
S-26                          P => P%NEXT
S-27                          IF ( .NOT. ASSOCIATED (P) ) EXIT
S-28                       END DO
S-29                   !$OMP END SINGLE
S-30               !$OMP END PARALLEL
```

```
S-31
S-32            END SUBROUTINE
S-33
S-34        END MODULE
```

Fortran

The **fib()** function should be called from within a **parallel** region for the different specified tasks to be executed in parallel. Also, only one thread of the **parallel** region should call **fib()** unless multiple concurrent Fibonacci computations are desired.

C / C++

Example tasking.4.c (**omp_3.0**)

```
S-1     int fib(int n) {
S-2        int i, j;
S-3        if (n<2)
S-4          return n;
S-5        else {
S-6           #pragma omp task shared(i)
S-7              i=fib(n-1);
S-8           #pragma omp task shared(j)
S-9              j=fib(n-2);
S-10          #pragma omp taskwait
S-11             return i+j;
S-12       }
S-13    }
```

C / C++

Fortran

Example tasking.4.f (**omp_3.0**)

```
S-1         RECURSIVE INTEGER FUNCTION fib(n) RESULT(res)
S-2         INTEGER n, i, j
S-3         IF ( n .LT. 2) THEN
S-4           res = n
S-5         ELSE
S-6     !$OMP TASK SHARED(i)
S-7           i = fib( n-1 )
S-8     !$OMP END TASK
S-9     !$OMP TASK SHARED(j)
S-10          j = fib( n-2 )
S-11    !$OMP END TASK
S-12    !$OMP TASKWAIT
S-13          res = i+j
S-14        END IF
S-15        END FUNCTION
```

Fortran

Note: There are more efficient algorithms for computing Fibonacci numbers. This classic recursion algorithm is for illustrative purposes.

The following example demonstrates a way to generate a large number of tasks with one thread and execute them with the threads in the team. While generating these tasks, the implementation may reach its limit on unassigned tasks. If it does, the implementation is allowed to cause the thread executing the task generating loop to suspend its task at the task scheduling point in the **task** directive, and start executing unassigned tasks. Once the number of unassigned tasks is sufficiently low, the thread may resume execution of the task generating loop.

C / C++

Example tasking.5.c (**omp_3.0**)

```
S-1   #define LARGE_NUMBER 10000000
S-2   double item[LARGE_NUMBER];
S-3   extern void process(double);
S-4
S-5   int main()
S-6   {
S-7   #pragma omp parallel
S-8     {
S-9       #pragma omp single
S-10      {
S-11        int i;
S-12        for (i=0; i<LARGE_NUMBER; i++)
S-13              #pragma omp task    // i is firstprivate, item is shared
S-14                   process(item[i]);
S-15      }
S-16    }
S-17  }
```

C / C++

Fortran

Example tasking.5.f (**omp_3.0**)

```
S-1          real*8 item(10000000)
S-2          integer i
S-3
S-4   !$omp parallel
S-5   !$omp single ! loop iteration variable i is private
S-6          do i=1,10000000
S-7   !$omp task
S-8             ! i is firstprivate, item is shared
S-9              call process(item(i))
S-10  !$omp end task
S-11         end do
S-12  !$omp end single
```

```
!$omp end parallel

      end
```

Fortran

The following example is the same as the previous one, except that the tasks are generated in an untied task. While generating the tasks, the implementation may reach its limit on unassigned tasks. If it does, the implementation is allowed to cause the thread executing the task generating loop to suspend its task at the task scheduling point in the **task** directive, and start executing unassigned tasks. If that thread begins execution of a task that takes a long time to complete, the other threads may complete all the other tasks before it is finished.

In this case, since the loop is in an untied task, any other thread is eligible to resume the task generating loop. In the previous examples, the other threads would be forced to idle until the generating thread finishes its long task, since the task generating loop was in a tied task.

C / C++

Example tasking.6.c (**omp_3.0**)

```
#define LARGE_NUMBER 10000000
double item[LARGE_NUMBER];
extern void process(double);
int main() {
#pragma omp parallel
  {
    #pragma omp single
    {
      int i;
      #pragma omp task untied
      // i is firstprivate, item is shared
      {
         for (i=0; i<LARGE_NUMBER; i++)
             #pragma omp task
                  process(item[i]);
      }
    }
  }
  return 0;
}
```

C / C++

Fortran

Example tasking.6.f (**omp_3.0**)

```
S-1         real*8 item(10000000)
S-2   !$omp parallel
S-3   !$omp single
S-4   !$omp task untied
S-5         ! loop iteration variable i is private
S-6         do i=1,10000000
S-7   !$omp task ! i is firstprivate, item is shared
S-8            call process(item(i))
S-9   !$omp end task
S-10        end do
S-11  !$omp end task
S-12  !$omp end single
S-13  !$omp end parallel
S-14        end
```

Fortran

The following two examples demonstrate how the scheduling rules illustrated in Section 2.11.3 of the OpenMP 4.0 specification affect the usage of **threadprivate** variables in tasks. A **threadprivate** variable can be modified by another task that is executed by the same thread. Thus, the value of a **threadprivate** variable cannot be assumed to be unchanged across a task scheduling point. In untied tasks, task scheduling points may be added in any place by the implementation.

A task switch may occur at a task scheduling point. A single thread may execute both of the task regions that modify **tp**. The parts of these task regions in which **tp** is modified may be executed in any order so the resulting value of **var** can be either 1 or 2.

C / C++

Example tasking.7.c (**omp_3.0**)

```
S-1
S-2   int tp;
S-3   #pragma omp threadprivate(tp)
S-4   int var;
S-5   void work()
S-6   {
S-7   #pragma omp task
S-8       {
S-9           /* do work here */
S-10  #pragma omp task
S-11          {
S-12              tp = 1;
S-13              /* do work here */
S-14  #pragma omp task
```

```
S-15                {
S-16                    /* no modification of tp */
S-17                }
S-18                var = tp; //value of tp can be 1 or 2
S-19            }
S-20            tp = 2;
S-21        }
S-22    }
```

C / C++

Fortran

Example tasking.7.f (`omp_3.0`)

```
S-1         module example
S-2         integer tp
S-3     !$omp threadprivate(tp)
S-4         integer var
S-5         contains
S-6         subroutine work
S-7     !$omp task
S-8             ! do work here
S-9     !$omp task
S-10            tp = 1
S-11            ! do work here
S-12    !$omp task
S-13              ! no modification of tp
S-14    !$omp end task
S-15            var = tp    ! value of var can be 1 or 2
S-16    !$omp end task
S-17          tp = 2
S-18    !$omp end task
S-19        end subroutine
S-20        end module
```

Fortran

In this example, scheduling constraints prohibit a thread in the team from executing a new task that modifies **tp** while another such task region tied to the same thread is suspended. Therefore, the value written will persist across the task scheduling point.

C / C++

Example tasking.8.c (`omp_3.0`)

```
S-1
S-2     int tp;
S-3     #pragma omp threadprivate(tp)
S-4     int var;
S-5     void work()
S-6     {
S-7     #pragma omp parallel
S-8         {
S-9             /* do work here */
S-10    #pragma omp task
S-11            {
S-12                tp++;
S-13                /* do work here */
S-14    #pragma omp task
S-15                {
S-16                    /* do work here but don't modify tp */
S-17                }
S-18                var = tp; //Value does not change after write above
S-19            }
S-20        }
S-21    }
```

C / C++

Fortran

Example tasking.8.f (`omp_3.0`)

```
S-1           module example
S-2           integer tp
S-3     !$omp threadprivate(tp)
S-4           integer var
S-5           contains
S-6           subroutine work
S-7     !$omp parallel
S-8              ! do work here
S-9     !$omp task
S-10             tp = tp + 1
S-11             ! do work here
S-12    !$omp task
S-13               ! do work here but don't modify tp
S-14    !$omp end task
S-15             var = tp    ! value does not change after write above
S-16    !$omp end task
S-17    !$omp end parallel
```

```
S-18        end subroutine
S-19        end module
```

Fortran

The following two examples demonstrate how the scheduling rules illustrated in Section 2.11.3 of the OpenMP 4.0 specification affect the usage of locks and critical sections in tasks. If a lock is held across a task scheduling point, no attempt should be made to acquire the same lock in any code that may be interleaved. Otherwise, a deadlock is possible.

In the example below, suppose the thread executing task 1 defers task 2. When it encounters the task scheduling point at task 3, it could suspend task 1 and begin task 2 which will result in a deadlock when it tries to enter critical region 1.

C / C++

Example tasking.9.c (**omp_3.0**)

```
S-1   void work()
S-2   {
S-3      #pragma omp task
S-4      { //Task 1
S-5          #pragma omp task
S-6          { //Task 2
S-7               #pragma omp critical //Critical region 1
S-8               {/*do work here */ }
S-9          }
S-10         #pragma omp critical //Critical Region 2
S-11         {
S-12             //Capture data for the following task
S-13             #pragma omp task
S-14             { /* do work here */ } //Task 3
S-15         }
S-16     }
S-17  }
```

C / C++

Fortran

Example tasking.9.f (`omp_3.0`)

```
       module example
       contains
       subroutine work
!$omp task
       ! Task 1
!$omp task
       ! Task 2
!$omp critical
       ! Critical region 1
       ! do work here
!$omp end critical
!$omp end task
!$omp critical
       ! Critical region 2
       ! Capture data for the following task
!$omp task
       !Task 3
       ! do work here
!$omp end task
!$omp end critical
!$omp end task
      end subroutine
      end module
```

Fortran

In the following example, **`lock`** is held across a task scheduling point. However, according to the scheduling restrictions, the executing thread can't begin executing one of the non-descendant tasks that also acquires **`lock`** before the task region is complete. Therefore, no deadlock is possible.

C / C++

Example tasking.10.c (`omp_3.0`)

```
#include <omp.h>
void work() {
    omp_lock_t lock;
    omp_init_lock(&lock);
#pragma omp parallel
    {
        int i;
#pragma omp for
        for (i = 0; i < 100; i++) {
#pragma omp task
            {
        // lock is shared by default in the task
```

```
S-13            omp_set_lock(&lock);
S-14                    // Capture data for the following task
S-15    #pragma omp task
S-16            // Task Scheduling Point 1
S-17                    { /* do work here */ }
S-18                    omp_unset_lock(&lock);
S-19                }
S-20            }
S-21        }
S-22        omp_destroy_lock(&lock);
S-23    }
```

C / C++

Fortran

Example tasking.10.f90 (omp_3.0)

```
S-1         module example
S-2         include 'omp_lib.h'
S-3         integer (kind=omp_lock_kind) lock
S-4         integer i
S-5
S-6         contains
S-7
S-8         subroutine work
S-9         call omp_init_lock(lock)
S-10    !$omp parallel
S-11        !$omp do
S-12        do i=1,100
S-13           !$omp task
S-14                ! Outer task
S-15                call omp_set_lock(lock)     ! lock is shared by
S-16                                            ! default in the task
S-17                      ! Capture data for the following task
S-18                      !$omp task     ! Task Scheduling Point 1
S-19                                ! do work here
S-20                      !$omp end task
S-21                 call omp_unset_lock(lock)
S-22           !$omp end task
S-23        end do
S-24    !$omp end parallel
S-25        call omp_destroy_lock(lock)
S-26        end subroutine
S-27
S-28        end module
```

Fortran

The following examples illustrate the use of the **mergeable** clause in the **task** construct. In this first example, the **task** construct has been annotated with the **mergeable** clause. The addition of this clause allows the implementation to reuse the data environment (including the ICVs) of the parent task for the task inside **foo** if the task is included or undeferred. Thus, the result of the execution may differ depending on whether the task is merged or not. Therefore the mergeable clause needs to be used with caution. In this example, the use of the mergeable clause is safe. As **x** is a shared variable the outcome does not depend on whether or not the task is merged (that is, the task will always increment the same variable and will always compute the same value for **x**).

C / C++

Example tasking.11.c (**omp_3.1**)

```
S-1   #include <stdio.h>
S-2   void foo ( )
S-3   {
S-4      int x = 2;
S-5      #pragma omp task shared(x) mergeable
S-6      {
S-7         x++;
S-8      }
S-9      #pragma omp taskwait
S-10     printf("%d\n",x);  // prints 3
S-11  }
```

C / C++

Fortran

Example tasking.11.f90 (**omp_3.1**)

```
S-1   subroutine foo()
S-2     integer :: x
S-3     x = 2
S-4   !$omp task shared(x) mergeable
S-5     x = x + 1
S-6   !$omp end task
S-7   !$omp taskwait
S-8     print *, x     ! prints 3
S-9   end subroutine
```

Fortran

This second example shows an incorrect use of the **mergeable** clause. In this example, the created task will access different instances of the variable **x** if the task is not merged, as **x** is **firstprivate**, but it will access the same variable **x** if the task is merged. As a result, the behavior of the program is unspecified and it can print two different values for **x** depending on the decisions taken by the implementation.

C / C++

Example tasking.12.c (**omp_3.1**)

```
S-1   #include <stdio.h>
S-2   void foo ( )
S-3   {
S-4      int x = 2;
S-5      #pragma omp task mergeable
S-6      {
S-7         x++;
S-8      }
S-9      #pragma omp taskwait
S-10     printf("%d\n",x);  // prints 2 or 3
S-11  }
```

C / C++

Fortran

Example tasking.12.f90 (**omp_3.1**)

```
S-1   subroutine foo()
S-2     integer :: x
S-3     x = 2
S-4   !$omp task mergeable
S-5     x = x + 1
S-6   !$omp end task
S-7   !$omp taskwait
S-8     print *, x   ! prints 2 or 3
S-9   end subroutine
```

Fortran

The following example shows the use of the **final** clause and the **omp_in_final** API call in a recursive binary search program. To reduce overhead, once a certain depth of recursion is reached the program uses the **final** clause to create only included tasks, which allow additional optimizations.

The use of the **omp_in_final** API call allows programmers to optimize their code by specifying which parts of the program are not necessary when a task can create only included tasks (that is, the code is inside a **final** task). In this example, the use of a different state variable is not necessary so once the program reaches the part of the computation that is finalized and copying from the parent state to the new state is eliminated. The allocation of **new_state** in the stack could also be avoided but it would make this example less clear. The **final** clause is most effective when used in conjunction with the **mergeable** clause since all tasks created in a **final** task region are included tasks that can be merged if the **mergeable** clause is present.

C / C++

Example tasking.13.c (**omp_3.1**)

```
S-1   #include <string.h>
S-2   #include <omp.h>
S-3   #define LIMIT  3 /* arbitrary limit on recursion depth */
S-4   void check_solution(char *);
S-5   void bin_search (int pos, int n, char *state)
S-6   {
S-7      if ( pos == n ) {
S-8         check_solution(state);
S-9         return;
S-10     }
S-11     #pragma omp task final( pos > LIMIT ) mergeable
S-12     {
S-13        char new_state[n];
S-14        if (!omp_in_final() ) {
S-15          memcpy(new_state, state, pos );
S-16          state = new_state;
S-17        }
S-18        state[pos] = 0;
S-19        bin_search(pos+1, n, state );
S-20     }
S-21     #pragma omp task final( pos > LIMIT ) mergeable
S-22     {
S-23        char new_state[n];
S-24        if (! omp_in_final() ) {
S-25          memcpy(new_state, state, pos );
S-26          state = new_state;
S-27        }
S-28        state[pos] = 1;
S-29        bin_search(pos+1, n, state );
S-30     }
S-31     #pragma omp taskwait
S-32  }
```

C / C++

Fortran

Example tasking.13.f90 (**omp_3.1**)

```
S-1   recursive subroutine bin_search(pos, n, state)
S-2     use omp_lib
S-3     integer :: pos, n
S-4     character, pointer :: state(:)
S-5     character, target, dimension(n) :: new_state1, new_state2
S-6     integer, parameter :: LIMIT = 3
S-7     if (pos .eq. n) then
S-8       call check_solution(state)
S-9       return
S-10    endif
S-11  !$omp task final(pos > LIMIT) mergeable
S-12    if (.not. omp_in_final()) then
S-13      new_state1(1:pos) = state(1:pos)
S-14      state => new_state1
S-15    endif
S-16    state(pos+1) = 'z'
S-17    call bin_search(pos+1, n, state)
S-18  !$omp end task
S-19  !$omp task final(pos > LIMIT) mergeable
S-20    if (.not. omp_in_final()) then
S-21      new_state2(1:pos) = state(1:pos)
S-22      state => new_state2
S-23    endif
S-24    state(pos+1) = 'y'
S-25    call bin_search(pos+1, n, state)
S-26  !$omp end task
S-27  !$omp taskwait
S-28  end subroutine
```

Fortran

The following example illustrates the difference between the **if** and the **final** clauses. The **if** clause has a local effect. In the first nest of tasks, the one that has the **if** clause will be undeferred but the task nested inside that task will not be affected by the **if** clause and will be created as usual. Alternatively, the **final** clause affects all **task** constructs in the **final** task region but not the **final** task itself. In the second nest of tasks, the nested tasks will be created as included tasks. Note also that the conditions for the **if** and **final** clauses are usually the opposite.

C / C++

Example tasking.14.c (`omp_3.1`)

```
S-1   void bar(void);
S-2
S-3   void foo ( )
S-4   {
S-5      int i;
S-6      #pragma omp task if(0)  // This task is undeferred
S-7      {
S-8         #pragma omp task     // This task is a regular task
S-9         for (i = 0; i < 3; i++) {
S-10            #pragma omp task     // This task is a regular task
S-11            bar();
S-12        }
S-13     }
S-14     #pragma omp task final(1) // This task is a regular task
S-15     {
S-16        #pragma omp task  // This task is included
S-17        for (i = 0; i < 3; i++) {
S-18            #pragma omp task     // This task is also included
S-19            bar();
S-20        }
S-21     }
S-22  }
```

C / C++

Fortran

Example tasking.14.f90 (`omp_3.1`)

```
S-1   subroutine foo()
S-2   integer i
S-3   !$omp task if(.FALSE.) ! This task is undeferred
S-4   !$omp task             ! This task is a regular task
S-5     do i = 1, 3
S-6       !$omp task             ! This task is a regular task
S-7         call bar()
S-8       !$omp end task
S-9     enddo
S-10  !$omp end task
S-11  !$omp end task
S-12  !$omp task final(.TRUE.) ! This task is a regular task
S-13  !$omp task               ! This task is included
S-14    do i = 1, 3
S-15      !$omp task               ! This task is also included
S-16       call bar()
S-17      !$omp end task
```

```
S-18      enddo
S-19    !$omp end task
S-20    !$omp end task
S-21    end subroutine
```

Fortran

3.2 Task Priority

In this example we compute arrays in a matrix through a *compute_array* routine. Each task has a priority value equal to the value of the loop variable *i* at the moment of its creation. A higher priority on a task means that a task is a candidate to run sooner.

The creation of tasks occurs in ascending order (according to the iteration space of the loop) but a hint, by means of the **priority** clause, is provided to reverse the execution order.

C / C++

Example task_priority.1.c (**omp_4.5**)

```
S-1   void compute_array (float *node, int M);
S-2
S-3   void compute_matrix (float *array, int N, int M)
S-4   {
S-5      int i;
S-6      #pragma omp parallel private(i)
S-7      #pragma omp single
S-8      {
S-9         for (i=0;i<N; i++) {
S-10           #pragma omp task priority(i)
S-11           compute_array(&array[i*M], M);
S-12        }
S-13     }
S-14  }
```

C / C++

Fortran

Example task_priority.1.f90 (**omp_4.5**)

```
S-1   subroutine compute_matrix(matrix, M, N)
S-2      implicit none
S-3      integer :: M, N
S-4      real :: matrix(M, N)
S-5      integer :: i
S-6      interface
S-7         subroutine compute_array(node, M)
S-8         implicit none
S-9         integer :: M
S-10        real :: node(M)
S-11        end subroutine
S-12     end interface
S-13     !$omp parallel private(i)
S-14     !$omp single
S-15     do i=1,N
S-16        !$omp task priority(i)
```

```
S-17          call compute_array(matrix(:, i), M)
S-18          !$omp end task
S-19       enddo
S-20       !$omp end single
S-21       !$omp end parallel
S-22    end subroutine compute_matrix
```

Fortran

3.3 Task Dependences

3.3.1 Flow Dependence

This example shows a simple flow dependence using a **depend** clause on the **task** construct.

C / C++

Example task_dep.1.c (**omp_4.0**)

```
S-1   #include <stdio.h>
S-2   int main() {
S-3      int x = 1;
S-4      #pragma omp parallel
S-5      #pragma omp single
S-6      {
S-7         #pragma omp task shared(x) depend(out: x)
S-8            x = 2;
S-9         #pragma omp task shared(x) depend(in: x)
S-10           printf("x = %d\n", x);
S-11     }
S-12     return 0;
S-13  }
```

C / C++

Fortran

Example task_dep.1.f90 (**omp_4.0**)

```
S-1   program example
S-2      integer :: x
S-3      x = 1
S-4      !$omp parallel
S-5      !$omp single
S-6         !$omp task shared(x) depend(out: x)
S-7            x = 2
S-8         !$omp end task
S-9         !$omp task shared(x) depend(in: x)
S-10           print*, "x = ", x
S-11        !$omp end task
S-12     !$omp end single
S-13     !$omp end parallel
S-14  end program
```

Fortran

The program will always print "x = 2", because the **depend** clauses enforce the ordering of the tasks. If the **depend** clauses had been omitted, then the tasks could execute in any order and the program and the program would have a race condition.

3.3.2 Anti-dependence

This example shows an anti-dependence using the **depend** clause on the **task** construct.

C / C++

Example task_dep.2.c (**omp_4.0**)

```
S-1   #include <stdio.h>
S-2   int main()
S-3   {
S-4      int x = 1;
S-5      #pragma omp parallel
S-6      #pragma omp single
S-7      {
S-8         #pragma omp task shared(x) depend(in: x)
S-9            printf("x = %d\n", x);
S-10        #pragma omp task shared(x) depend(out: x)
S-11           x = 2;
S-12     }
S-13     return 0;
S-14  }
```

C / C++

Fortran

Example task_dep.2.f90 (**omp_4.0**)

```
S-1   program example
S-2      integer :: x
S-3      x = 1
S-4      !$omp parallel
S-5      !$omp single
S-6         !$omp task shared(x) depend(in: x)
S-7            print*, "x = ", x
S-8         !$omp end task
S-9         !$omp task shared(x) depend(out: x)
S-10           x = 2
S-11        !$omp end task
S-12     !$omp end single
S-13     !$omp end parallel
S-14  end program
```

Fortran

The program will always print "x = 1", because the **depend** clauses enforce the ordering of the tasks. If the **depend** clauses had been omitted, then the tasks could execute in any order and the program would have a race condition.

3.3.3 Output Dependence

This example shows an output dependence using the **depend** clause on the **task** construct.

C / C++

Example task_dep.3.c (**omp_4.0**)

```
#include <stdio.h>
int main() {
   int x;
   #pragma omp parallel
   #pragma omp single
   {
      #pragma omp task shared(x) depend(out: x)
         x = 1;
      #pragma omp task shared(x) depend(out: x)
         x = 2;
      #pragma omp taskwait
      printf("x = %d\n", x);
   }
   return 0;
}
```

C / C++

Fortran

Example task_dep.3.f90 (**omp_4.0**)

```
program example
   integer :: x
   !$omp parallel
   !$omp single
      !$omp task shared(x) depend(out: x)
         x = 1
      !$omp end task
      !$omp task shared(x) depend(out: x)
         x = 2
      !$omp end task
      !$omp taskwait
      print*, "x = ", x
   !$omp end single
   !$omp end parallel
end program
```

Fortran

The program will always print "x = 2", because the **depend** clauses enforce the ordering of the tasks. If the **depend** clauses had been omitted, then the tasks could execute in any order and the program would have a race condition.

3.3.4 Concurrent Execution with Dependences

In this example we show potentially concurrent execution of tasks using multiple flow dependences expressed using the **depend** clause on the **task** construct.

C / C++

Example task_dep.4.c (**omp_4.0**)

```
S-1   #include <stdio.h>
S-2   int main() {
S-3      int x = 1;
S-4      #pragma omp parallel
S-5      #pragma omp single
S-6      {
S-7         #pragma omp task shared(x) depend(out: x)
S-8            x = 2;
S-9         #pragma omp task shared(x) depend(in: x)
S-10           printf("x + 1 = %d. ", x+1);
S-11        #pragma omp task shared(x) depend(in: x)
S-12           printf("x + 2 = %d\n", x+2);
S-13     }
S-14     return 0;
S-15  }
```

C / C++

Fortran

Example task_dep.4.f90 (**omp_4.0**)

```
S-1
S-2   program example
S-3      integer :: x
S-4
S-5      x = 1
S-6
S-7      !$omp parallel
S-8      !$omp single
S-9
S-10        !$omp task shared(x) depend(out: x)
S-11           x = 2
S-12        !$omp end task
S-13
S-14        !$omp task shared(x) depend(in: x)
S-15           print*, "x + 1 = ", x+1, "."
S-16        !$omp end task
S-17
S-18        !$omp task shared(x) depend(in: x)
S-19           print*, "x + 2 = ", x+2, "."
```

```
S-20        !$omp end task
S-21
S-22     !$omp end single
S-23     !$omp end parallel
S-24  end program
```

Fortran

The last two tasks are dependent on the first task. However there is no dependence between the last two tasks, which may execute in any order (or concurrently if more than one thread is available). Thus, the possible outputs are "x + 1 = 3. x + 2 = 4. " and "x + 2 = 4. x + 1 = 3. ". If the **depend** clauses had been omitted, then all of the tasks could execute in any order and the program would have a race condition.

3.3.5 Matrix multiplication

This example shows a task-based blocked matrix multiplication. Matrices are of NxN elements, and the multiplication is implemented using blocks of BSxBS elements.

C / C++

Example task_dep.5.c (**omp_4.0**)

```
S-1   // Assume BS divides N perfectly
S-2   void matmul_depend(int N, int BS, float A[N][N], float B[N][N], float
S-3   C[N][N] )
S-4   {
S-5      int i, j, k, ii, jj, kk;
S-6      for (i = 0; i < N; i+=BS) {
S-7         for (j = 0; j < N; j+=BS) {
S-8            for (k = 0; k < N; k+=BS) {
S-9   // Note 1: i, j, k, A, B, C are firstprivate by default
S-10  // Note 2: A, B and C are just pointers
S-11  #pragma omp task private(ii, jj, kk) \
S-12              depend ( in: A[i:BS][k:BS], B[k:BS][j:BS] ) \
S-13              depend ( inout: C[i:BS][j:BS] )
S-14              for (ii = i; ii < i+BS; ii++ )
S-15                 for (jj = j; jj < j+BS; jj++ )
S-16                    for (kk = k; kk < k+BS; kk++ )
S-17                       C[ii][jj] = C[ii][jj] + A[ii][kk] * B[kk][jj];
S-18           }
S-19        }
S-20     }
S-21  }
```

C / C++

Fortran

Example task_dep.5.f90 (`omp_4.0`)

```
S-1   ! Assume BS divides N perfectly
S-2   subroutine matmul_depend (N, BS, A, B, C)
S-3      implicit none
S-4      integer :: N, BS, BM
S-5      real, dimension(N, N) :: A, B, C
S-6      integer :: i, j, k, ii, jj, kk
S-7      BM = BS - 1
S-8      do i = 1, N, BS
S-9         do j = 1, N, BS
S-10           do k = 1, N, BS
S-11  !$omp task shared(A,B,C) private(ii,jj,kk) & ! I,J,K are firstprivate by default
S-12  !$omp depend ( in: A(i:i+BM, k:k+BM), B(k:k+BM, j:j+BM) ) &
S-13  !$omp depend ( inout: C(i:i+BM, j:j+BM) )
S-14                 do ii = i, i+BM
S-15                    do jj = j, j+BM
S-16                       do kk = k, k+BM
S-17                          C(jj,ii) = C(jj,ii) + A(kk,ii) * B(jj,kk)
S-18                       end do
S-19                    end do
S-20                 end do
S-21  !$omp end task
S-22           end do
S-23        end do
S-24     end do
S-25  end subroutine
```

Fortran

3.3.6 `taskwait` with Dependences

In this subsection three examples illustrate how the **`depend`** clause can be applied to a **`taskwait`** construct to make the generating task wait for specific child tasks to complete. This is an OpenMP 5.0 feature. In the same manner that dependences can order executions among child tasks with **`depend`** clauses on **`task`** constructs, the generating task can be scheduled to wait on child tasks at a **`taskwait`** before it can proceed.

Note: Since the **`depend`** clause on a **`taskwait`** construct relaxes the default synchronization behavior (waiting for all children to finish), it is important to realize that child tasks that are not predecessor tasks, as determined by the **`depend`** clause of the **`taskwait`** construct, may be running concurrently while the generating task is executing after the taskwait.

In the first example the generating task waits at the **taskwait** construct for the completion of the first child task because a dependence on the first task is produced by *x* with an **in** dependence type within the **depend** clause of the **taskwait** construct. Immediately after the first **taskwait** construct it is safe to access the *x* variable by the generating task, as shown in the print statement. There is no completion restraint on the second child task. Hence, immediately after the first **taskwait** it is unsafe to access the *y* variable since the second child task may still be executing. The second **taskwait** ensures that the second child task has completed; hence it is safe to access the *y* variable in the following print statement.

C / C++

Example task_dep.6.c (**omp_5.0**)

```
S-1
S-2     #include<stdio.h>
S-3
S-4     void foo()
S-5     {
S-6         int x = 0, y = 2;
S-7
S-8         #pragma omp task depend(inout: x) shared(x)
S-9         x++;                                          // 1st child task
S-10
S-11        #pragma omp task shared(y)
S-12        y--;                                          // 2nd child task
S-13
S-14        #pragma omp taskwait depend(in: x)            // 1st taskwait
S-15
S-16        printf("x=%d\n",x);
S-17
S-18        // Second task may not be finished.
S-19        // Accessing y here will create a race condition.
S-20
S-21        #pragma omp taskwait                          // 2nd taskwait
S-22
S-23        printf("y=%d\n",y);
S-24    }
S-25
S-26    int main()
S-27    {
S-28        #pragma omp parallel
S-29        #pragma omp single
S-30        foo();
S-31
S-32        return 0;
S-33    }
```

C / C++

Fortran

Example task_dep.6.f90 (`omp_5.0`)

```
S-1
S-2
S-3     subroutine foo()
S-4         implicit none
S-5         integer :: x, y
S-6
S-7         x = 0
S-8         y = 2
S-9
S-10        !$omp task depend(inout: x) shared(x)
S-11            x = x + 1                          !! 1st child task
S-12        !$omp end task
S-13
S-14        !$omp task shared(y)
S-15            y = y - 1                          !! 2nd child task
S-16        !$omp end task
S-17
S-18        !$omp taskwait depend(in: x)           !! 1st taskwait
S-19
S-20        print*, "x=", x
S-21
S-22        !! Second task may not be finished.
S-23        !! Accessing y here will create a race condition.
S-24
S-25        !$omp taskwait                         !! 2nd taskwait
S-26
S-27        print*, "y=", y
S-28
S-29    end subroutine foo
S-30
S-31    program p
S-32        implicit none
S-33        !$omp parallel
S-34        !$omp single
S-35           call foo()
S-36        !$omp end single
S-37        !$omp end parallel
S-38    end program p
```

Fortran

In this example the first two tasks are serialized, because a dependence on the first child is produced by *x* with the **in** dependence type in the **depend** clause of the second task. However, the generating task at the first **taskwait** waits only on the first child task to complete, because a

dependence on only the first child task is produced by *x* with an **in** dependence type within the **depend** clause of the **taskwait** construct. The second **taskwait** (without a **depend** clause) is included to guarantee completion of the second task before *y* is accessed. (While unnecessary, the **depend(inout: y)** clause on the 2nd child task is included to illustrate how the child task dependences can be completely annotated in a data-flow model.)

C / C++

Example task_dep.7.c (**omp_5.0**)

```
S-1
S-2     #include<stdio.h>
S-3
S-4     void foo()
S-5     {
S-6         int x = 0, y = 2;
S-7
S-8         #pragma omp task depend(inout: x) shared(x)
S-9         x++;                                              // 1st child task
S-10
S-11        #pragma omp task depend(in: x) depend(inout: y) shared(x, y)
S-12        y -= x;                                           // 2nd child task
S-13
S-14        #pragma omp taskwait depend(in: x)                // 1st taskwait
S-15
S-16        printf("x=%d\n",x);
S-17
S-18        // Second task may not be finished.
S-19        // Accessing y here would create a race condition.
S-20
S-21        #pragma omp taskwait                              // 2nd taskwait
S-22
S-23        printf("y=%d\n",y);
S-24
S-25    }
S-26
S-27    int main()
S-28    {
S-29        #pragma omp parallel
S-30        #pragma omp single
S-31        foo();
S-32
S-33        return 0;
S-34    }
```

C / C++

Fortran

Example task_dep.7.f90 (`omp_5.0`)

```
S-1
S-2
S-3     subroutine foo()
S-4     implicit none
S-5     integer :: x, y
S-6
S-7         x = 0
S-8         y = 2
S-9
S-10        !$omp task depend(inout: x) shared(x)
S-11            x = x + 1                          !! 1st child task
S-12        !$omp end task
S-13
S-14        !$omp task depend(in: x) depend(inout: y) shared(x, y)
S-15            y = y - x                          !! 2nd child task
S-16        !$omp end task
S-17
S-18        !$omp taskwait depend(in: x)           !! 1st taskwait
S-19
S-20        print*, "x=", x
S-21
S-22        !! Second task may not be finished.
S-23        !! Accessing y here would create a race condition.
S-24
S-25        !$omp taskwait                         !! 2nd taskwait
S-26
S-27        print*, "y=", y
S-28
S-29    end subroutine foo
S-30
S-31    program p
S-32    implicit none
S-33        !$omp parallel
S-34        !$omp single
S-35           call foo()
S-36        !$omp end single
S-37        !$omp end parallel
S-38    end program p
```

Fortran

This example is similar to the previous one, except the generating task is directed to also wait for completion of the second task.

The **depend** clause of the **taskwait** construct now includes an **in** dependence type for *y*. Hence the generating task must now wait on completion of any child task having *y* with an **out** (here **inout**) dependence type in its **depend** clause. So, the **depend** clause of the **taskwait** construct now constrains the second task to complete at the **taskwait**, too. (This change makes the second **taskwait** of the previous example unnecessary– it has been removed in this example.)

Note: While a taskwait construct ensures that all child tasks have completed; a depend clause on a taskwait construct only waits for specific child tasks (prescribed by the dependence type and list items in the **taskwait**'s **depend** clause). This and the previous example illustrate the need to carefully determine the dependence type of variables in the **taskwait depend** clause when selecting child tasks that the generating task must wait on, so that its execution after the taskwait does not produce race conditions on variables accessed by non-completed child tasks.

C / C++

Example task_dep.8.c (**omp_5.0**)

```
S-1
S-2     #include<stdio.h>
S-3
S-4     void foo()
S-5     {
S-6         int x = 0, y = 2;
S-7
S-8         #pragma omp task depend(inout: x) shared(x)
S-9         x++;                                              // 1st child task
S-10
S-11        #pragma omp task depend(in: x) depend(inout: y) shared(x, y)
S-12        y -= x;                                           // 2st child task
S-13
S-14        #pragma omp taskwait depend(in: x,y)
S-15
S-16        printf("x=%d\n",x);
S-17        printf("y=%d\n",y);
S-18
S-19    }
S-20
S-21    int main()
S-22    {
S-23        #pragma omp parallel
S-24        #pragma omp single
S-25        foo();
S-26
S-27        return 0;
S-28    }
```

C / C++

Fortran

Example task_dep.8.f90 (`omp_5.0`)

```
S-1
S-2
S-3     subroutine foo()
S-4     implicit nonE
S-5     integer :: x, y
S-6
S-7         x = 0
S-8         y = 2
S-9
S-10        !$omp task depend(inout: x) shared(x)
S-11            x = x + 1                          !! 1st child task
S-12        !$omp end task
S-13
S-14        !$omp task depend(in: x) depend(inout: y) shared(x, y)
S-15            y = y - x                          !! 2nd child task
S-16        !$omp end task
S-17
S-18        !$omp taskwait depend(in: x,y)
S-19
S-20        print*, "x=", x
S-21        print*, "y=", y
S-22
S-23    end subroutine foo
S-24
S-25    program p
S-26    implicit none
S-27        !$omp parallel
S-28        !$omp single
S-29           call foo()
S-30        !$omp end single
S-31        !$omp end parallel
S-32    end program p
```

Fortran

3.3.7 Mutually Exclusive Execution with Dependences

In this example we show a series of tasks, including mutually exclusive tasks, expressing dependences using the **depend** clause on the **task** construct.

The program will always print 6. Tasks T1, T2 and T3 will be scheduled first, in any order. Task T4 will be scheduled after tasks T1 and T2 are completed. T5 will be scheduled after tasks T1 and T3 are completed. Due to the **mutexinoutset** dependence type on **c**, T4 and T5 may be scheduled in any order with respect to each other, but not at the same time. Tasks T6 will be scheduled after both T4 and T5 are completed.

C / C++

Example task_dep.9.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   int main()
S-3   {
S-4      int a, b, c, d;
S-5      #pragma omp parallel
S-6      #pragma omp single
S-7      {
S-8         #pragma omp task depend(out: c)
S-9            c = 1;   /* Task T1 */
S-10        #pragma omp task depend(out: a)
S-11           a = 2;   /* Task T2 */
S-12        #pragma omp task depend(out: b)
S-13           b = 3;   /* Task T3 */
S-14        #pragma omp task depend(in: a) depend(mutexinoutset: c)
S-15           c += a;  /* Task T4 */
S-16        #pragma omp task depend(in: b) depend(mutexinoutset: c)
S-17           c += b;  /* Task T5 */
S-18        #pragma omp task depend(in: c)
S-19           d = c;   /* Task T6 */
S-20     }
S-21     printf("%d\n", d);
S-22     return 0;
S-23  }
```

C / C++

Fortran

Example task_dep.9.f90 (`omp_5.0`)

```
S-1   program example
S-2      integer :: a, b, c, d
S-3      !$omp parallel
S-4      !$omp single
S-5         !$omp task depend(out: c)
S-6         c = 1       ! Task T1
S-7         !$omp end task
S-8         !$omp task depend(out: a)
S-9         a = 2       ! Task T2
S-10        !$omp end task
S-11        !$omp task depend(out: b)
S-12        b = 3       ! Task T3
S-13        !$omp end task
S-14        !$omp task depend(in: a) depend(mutexinoutset: c)
S-15        c = c + a   ! Task T4
S-16        !$omp end task
S-17        !$omp task depend(in: b) depend(mutexinoutset: c)
S-18        c = c + b   ! Task T5
S-19        !$omp end task
S-20        !$omp task depend(in: c)
S-21        d = c       ! Task T6
S-22        !$omp end task
S-23     !$omp end single
S-24     !$omp end parallel
S-25     print *, d
S-26  end program
```

Fortran

The following example demonstrates a situation where the **`mutexinoutset`** dependence type is advantageous. If **`shortTaskB`** completes before **`longTaskA`**, the runtime can take advantage of this by scheduling **`longTaskBC`** before **`shortTaskAC`**.

C / C++

Example task_dep.10.c (`omp_5.0`)

```
S-1   extern int longTaskA(), shortTaskB();
S-2   extern int shortTaskAC(int,int), longTaskBC(int,int);
S-3   void foo (void)
S-4   {
S-5     int a, b, c;
S-6     c = 0;
S-7     #pragma omp parallel
S-8     #pragma omp single
S-9     {
```

```
S-10        #pragma omp task depend(out: a)
S-11           a = longTaskA();
S-12        #pragma omp task depend(out: b)
S-13           b = shortTaskB();
S-14        #pragma omp task depend(in: a) depend(mutexinoutset: c)
S-15           c = shortTaskAC(a,c);
S-16        #pragma omp task depend(in: b) depend(mutexinoutset: c)
S-17           c = longTaskBC(b,c);
S-18     }
S-19  }
```

C / C++

Fortran

1 *Example task_dep.10.f90* (`omp_5.0`)

```
S-1   subroutine foo
S-2      integer :: a,b,c
S-3      c = 0
S-4      !$omp parallel
S-5      !$omp single
S-6         !$omp task depend(out: a)
S-7            a = longTaskA()
S-8         !$omp end task
S-9         !$omp task depend(out: b)
S-10           b = shortTaskB()
S-11        !$omp end task
S-12        !$omp task depend(in: a) depend(mutexinoutset: c)
S-13           c = shortTaskAC(a,c)
S-14        !$omp end task
S-15        !$omp task depend(in: b) depend(mutexinoutset: c)
S-16           c = longTaskBC(b,c)
S-17        !$omp end task
S-18     !$omp end single
S-19     !$omp end parallel
S-20  end subroutine foo
```

Fortran

3.3.8 Multidependences Using Iterators

The following example uses an iterator to define a dynamic number of dependences.

In the **single** construct of a parallel region a loop generates n tasks and each task has an **out** dependence specified through an element of the *v* array. This is followed by a single task that defines an **in** dependence on each element of the array. This is accomplished by using the **iterator** modifier in the **depend** clause, supporting a dynamic number of dependences (*n* here).

The task for the *print_all_elements* function is not executed until all dependences prescribed (or registered) by the iterator are fulfilled; that is, after all the tasks generated by the loop have completed.

Note, one cannot simply use an array section in the **depend** clause of the second task construct because this would violate the **depend** clause restriction:

"List items used in **depend** clauses of the same task or sibling tasks must indicate identical storage locations or disjoint storage locations".

In this case each of the loop tasks use a single disjoint (different storage) element in their **depend** clause; however, the array-section storage area prescribed in the commented directive is neither identical nor disjoint to the storage prescibed by the elements of the loop tasks. The iterator overcomes this restriction by effectively creating n disjoint storage areas.

C / C++

Example task_dep.11.c (**omp_5.0**)

```
S-1   #include<stdio.h>
S-2
S-3   void set_an_element(int *p, int val) {
S-4       *p = val;
S-5   }
S-6
S-7   void print_all_elements(int *v, int n) {
S-8       int i;
S-9       for (i = 0; i < n; ++i) {
S-10          printf("%d, ", v[i]);
S-11      }
S-12      printf("\n");
S-13  }
S-14
S-15  void parallel_computation(int n) {
S-16      int v[n];
S-17      #pragma omp parallel
S-18      #pragma omp single
S-19      {
S-20          int i;
S-21          for (i = 0; i < n; ++i)
S-22              #pragma omp task depend(out: v[i])
```

```
S-23                set_an_element(&v[i], i);
S-24
S-25            #pragma omp task depend(iterator(it = 0:n), in: v[it])
S-26         // #pragma omp task depend(in: v[0:n]) Violates Array section restriction.
S-27            print_all_elements(v, n);
S-28
S-29        }
S-30    }
S-31
```

C / C++

Fortran

1 *Example task_dep.11.f90* (`omp_5.0`)

```
 S-1    subroutine set_an_element(e, val)
 S-2        implicit none
 S-3        integer :: e, val
 S-4
 S-5        e = val
 S-6
 S-7    end subroutine
 S-8
 S-9    subroutine print_all_elements(v, n)
S-10        implicit none
S-11        integer :: n, v(n)
S-12
S-13        print *, v
S-14
S-15    end subroutine
S-16
S-17    subroutine parallel_computation(n)
S-18        implicit none
S-19        integer :: n
S-20        integer :: i, v(n)
S-21
S-22        !$omp parallel
S-23        !$omp single
S-24            do i=1, n
S-25                !$omp task depend(out: v(i))
S-26                     call set_an_element(v(i), i)
S-27                !$omp end task
S-28            enddo
S-29
S-30            !$omp task depend(iterator(it = 1:n), in: v(it))
S-31           !!$omp task depend(in: v(1:n)) Violates Array section restriction.
S-32                call print_all_elements(v, n)
S-33            !$omp end task
```

```
!$omp end single
!$omp end parallel
end subroutine
```

Fortran

3.3.9 Dependence for Undeferred Tasks

In the following example, we show that even if a task is undeferred as specified by an **if** clause that evaluates to *false*, task dependences are still honored.

The **depend** clauses of the first and second explicit tasks specify that the first task is completed before the second task.

The second explicit task has an **if** clause that evaluates to *false*. This means that the execution of the generating task (the implicit task of the **single** region) must be suspended until the second explict task is completed. But, because of the dependence, the first explicit task must complete first, then the second explicit task can execute and complete, and only then the generating task can resume to the print statement. Thus, the program will always print "`x = 2`".

C / C++

Example task_dep.12.c (**omp_4.0**)

```
#include <stdio.h>
int main (int argc, char *argv[])
{
  int x = 0;
  #pragma omp parallel
  #pragma omp single
  {
    /* first explicit task */
    #pragma omp task shared(x) depend(out: x)
      x = 1;

    /* second explicit task */
    #pragma omp task shared(x) depend(inout: x) if(0)
      x = 2;

    /* statement executed by parent implicit task
       prints: x = 2 */
    printf("x = %d\n", x);
  }
  return 0;
}
```

C / C++

Fortran

Example task_dep.12.f90 (`omp_4.0`)

```
S-1   program example
S-2      integer :: x
S-3      x = 0
S-4      !$omp parallel
S-5      !$omp single
S-6        !... first explicit task
S-7        !$omp task shared(x) depend(out: x)
S-8           x = 1
S-9        !$omp end task
S-10
S-11       !... second explicit task
S-12       !$omp task shared(x) depend(inout: x) if(.false.)
S-13          x = 2
S-14       !$omp end task
S-15
S-16       !... statement executed by parent implicit task
S-17       ! prints: x = 2
S-18        print*, "x = ", x
S-19     !$omp end single
S-20     !$omp end parallel
S-21  end program
```

Fortran

3.4 The `taskgroup` Construct

In this example, tasks are grouped and synchronized using the **`taskgroup`** construct.

Initially, one task (the task executing the **`start_background_work()`** call) is created in the **`parallel`** region, and later a parallel tree traversal is started (the task executing the root of the recursive **`compute_tree()`** calls). While synchronizing tasks at the end of each tree traversal, using the **`taskgroup`** construct ensures that the formerly started background task does not participate in the synchronization, and is left free to execute in parallel. This is opposed to the behavior of the **`taskwait`** construct, which would include the background tasks in the synchronization.

C / C++

Example taskgroup.1.c (**`omp_4.0`**)

```
S-1   extern void start_background_work(void);
S-2   extern void check_step(void);
S-3   extern void print_results(void);
S-4   struct tree_node
S-5   {
S-6      struct tree_node *left;
S-7      struct tree_node *right;
S-8   };
S-9   typedef struct tree_node* tree_type;
S-10  extern void init_tree(tree_type);
S-11  #define max_steps 100
S-12  void compute_something(tree_type tree)
S-13  {
S-14     // some computation
S-15  }
S-16  void compute_tree(tree_type tree)
S-17  {
S-18     if (tree->left)
S-19     {
S-20       #pragma omp task
S-21         compute_tree(tree->left);
S-22     }
S-23     if (tree->right)
S-24     {
S-25       #pragma omp task
S-26         compute_tree(tree->right);
S-27     }
S-28     #pragma omp task
S-29     compute_something(tree);
S-30  }
S-31  int main()
S-32  {
```

```
  int i;
  tree_type tree;
  init_tree(tree);
  #pragma omp parallel
  #pragma omp single
  {
    #pragma omp task
      start_background_work();
    for (i = 0; i < max_steps; i++)
    {
        #pragma omp taskgroup
        {
           #pragma omp task
             compute_tree(tree);
        } // wait on tree traversal in this step
        check_step();
    }
  } // only now is background work required to be complete
  print_results();
  return 0;
}
```

C / C++

Fortran

Example taskgroup.1.f90 (`omp_4.0`)

```
module tree_type_mod
  integer, parameter :: max_steps=100
  type tree_type
    type(tree_type), pointer :: left, right
  end type
  contains
    subroutine compute_something(tree)
      type(tree_type), pointer :: tree
! some computation
    end subroutine
    recursive subroutine compute_tree(tree)
      type(tree_type), pointer :: tree
      if (associated(tree%left)) then
!$omp task
        call compute_tree(tree%left)
!$omp end task
      endif
      if (associated(tree%right)) then
!$omp task
        call compute_tree(tree%right)
!$omp end task
```

```
S-22          endif
S-23    !$omp task
S-24          call compute_something(tree)
S-25    !$omp end task
S-26        end subroutine
S-27    end module
S-28    program main
S-29      use tree_type_mod
S-30      type(tree_type), pointer :: tree
S-31      call init_tree(tree);
S-32    !$omp parallel
S-33    !$omp single
S-34    !$omp task
S-35      call start_background_work()
S-36    !$omp end task
S-37      do i=1, max_steps
S-38    !$omp taskgroup
S-39    !$omp task
S-40        call compute_tree(tree)
S-41    !$omp end task
S-42    !$omp end taskgroup ! wait on tree traversal in this step
S-43        call check_step()
S-44      enddo
S-45    !$omp end single
S-46    !$omp end parallel    ! only now is background work required to be complete
S-47      call print_results()
S-48    end program
```

Fortran

3.5 The `taskyield` Construct

The following example illustrates the use of the **taskyield** directive. The tasks in the example compute something useful and then do some computation that must be done in a critical region. By using **taskyield** when a task cannot get access to the **critical** region the implementation can suspend the current task and schedule some other task that can do something useful.

C / C++

Example taskyield.1.c (**omp_3.1**)

```
S-1   #include <omp.h>
S-2
S-3   void something_useful ( void );
S-4   void something_critical ( void );
S-5   void foo ( omp_lock_t * lock, int n )
S-6   {
S-7      int i;
S-8
S-9      for ( i = 0; i < n; i++ )
S-10        #pragma omp task
S-11        {
S-12            something_useful();
S-13            while ( !omp_test_lock(lock) ) {
S-14               #pragma omp taskyield
S-15            }
S-16            something_critical();
S-17            omp_unset_lock(lock);
S-18        }
S-19  }
```

C / C++

Fortran

Example taskyield.1.f90 (**omp_3.1**)

```
S-1   subroutine foo ( lock, n )
S-2      use omp_lib
S-3      integer (kind=omp_lock_kind) :: lock
S-4      integer n
S-5      integer i
S-6
S-7      do i = 1, n
S-8        !$omp task
S-9          call something_useful()
S-10         do while ( .not. omp_test_lock(lock) )
S-11           !$omp taskyield
S-12         end do
S-13         call something_critical()
```

```
S-14           call omp_unset_lock(lock)
S-15         !$omp end task
S-16       end do
S-17
S-18   end subroutine
```

Fortran

3.6 The `taskloop` Construct

The following example illustrates how to execute a long running task concurrently with tasks created with a **taskloop** directive for a loop having unbalanced amounts of work for its iterations.

The **grainsize** clause specifies that each task is to execute at least 500 iterations of the loop.

The **nogroup** clause removes the implicit taskgroup of the **taskloop** construct; the explicit **taskgroup** construct in the example ensures that the function is not exited before the long-running task and the loops have finished execution.

C / C++

Example taskloop.1.c (**omp_4.5**)

```
S-1   void long_running_task(void);
S-2   void loop_body(int i, int j);
S-3
S-4   void parallel_work(void) {
S-5      int i, j;
S-6   #pragma omp taskgroup
S-7      {
S-8   #pragma omp task
S-9         long_running_task(); // can execute concurrently
S-10
S-11  #pragma omp taskloop private(j) grainsize(500) nogroup
S-12        for (i = 0; i < 10000; i++) { // can execute concurrently
S-13           for (j = 0; j < i; j++) {
S-14              loop_body(i, j);
S-15           }
S-16        }
S-17     }
S-18  }
```

C / C++

Fortran

Example taskloop.1.f90 (**omp_4.5**)

```
S-1   subroutine parallel_work
S-2      integer i
S-3      integer j
S-4   !$omp taskgroup
S-5
S-6   !$omp task
S-7      call long_running_task()
S-8   !$omp end task
S-9
S-10  !$omp taskloop private(j) grainsize(500) nogroup
S-11     do i=1,10000
S-12        do j=1,i
S-13           call loop_body(i, j)
S-14        end do
S-15     end do
S-16  !$omp end taskloop
S-17
S-18  !$omp end taskgroup
S-19  end subroutine
```

Fortran

Because a **taskloop** construct encloses a loop, it is often incorrectly perceived as a worksharing construct (when it is directly nested in a **parallel** region).

While a worksharing construct distributes the loop iterations across all threads in a team, the entire loop of a **taskloop** construct is executed by every thread of the team.

In the example below the first taskloop occurs closely nested within a **parallel** region and the entire loop is executed by each of the *T* threads; hence the reduction sum is executed *T***N* times.

The loop of the second taskloop is within a **single** region and is executed by a single thread so that only *N* reduction sums occur. (The other *N*-1 threads of the **parallel** region will participate in executing the tasks. This is the common use case for the **taskloop** construct.)

In the example, the code thus prints **x1 = 16384** (*T***N*) and **x2 = 1024** (*N*).

C / C++

Example taskloop.2.c (**omp_4.5**)

```
S-1   #include <stdio.h>
S-2
S-3   #define T 16
S-4   #define N 1024
S-5
S-6   void parallel_work() {
S-7       int x1 = 0, x2 = 0;
S-8
S-9       #pragma omp parallel shared(x1,x2) num_threads(T)
S-10      {
S-11          #pragma omp taskloop
S-12          for (int i = 0; i < N; ++i) {
S-13              #pragma omp atomic
S-14              x1++;           // executed T*N times
S-15          }
S-16
S-17          #pragma omp single
S-18          #pragma omp taskloop
S-19          for (int i = 0; i < N; ++i) {
S-20              #pragma omp atomic
S-21              x2++;           // executed N times
S-22          }
S-23      }
S-24
S-25      printf("x1 = %d, x2 = %d\n", x1, x2);
S-26  }
```

C / C++

Fortran

Example taskloop.2.f90 (**omp_4.5**)

```
S-1   subroutine parallel_work
S-2       implicit none
S-3       integer :: x1, x2
S-4       integer :: i
S-5       integer, parameter :: T = 16
S-6       integer, parameter :: N = 1024
S-7
S-8       x1 = 0
S-9       x2 = 0
S-10      !$omp parallel shared(x1,x2) num_threads(T)
S-11      !$omp taskloop
S-12      do i = 1,N
S-13          !$omp atomic
```

```
S-14            x1 = x1 + 1     ! executed T*N times
S-15            !$omp end atomic
S-16        end do
S-17        !$omp end taskloop
S-18
S-19        !$omp single
S-20        !$omp taskloop
S-21        do i = 1,N
S-22            !$omp atomic
S-23            x2 = x2 + 1     ! executed N times
S-24            !$omp end atomic
S-25        end do
S-26        !$omp end taskloop
S-27        !$omp end single
S-28        !$omp end parallel
S-29
S-30        write (*,'(A,I0,A,I0)') 'x1 = ', x1, ', x2 = ',x2
S-31    end subroutine
```

Fortran

3.7 The `parallel master taskloop` Construct

In the OpenMP 5.0 Specification several combined constructs containing the **taskloop** construct were added.

Just as the **for** and **do** constructs have been combined with the **parallel** construct for convenience, so too, the combined **parallel master taskloop** and **parallel master taskloop simd** constructs have been created for convenience.

In the following example the first **taskloop** construct is enclosed by the usual **parallel** and **master** constructs to form a team of threads, and a single task generator (master thread) for the **taskloop** construct.

The same OpenMP operations for the first taskloop are accomplished by the second taskloop with the **parallel master taskloop** combined construct. The third taskloop uses the combined **parallel master taskloop simd** construct to accomplish the same behavior as closely nested **parallel master**, and **taskloop simd** constructs.

As with any combined construct the clauses of the components may be used with appropriate restrictions. The combination of the **parallel master** construct with the **taskloop** or **taskloop simd** construct produces no additional restrictions.

C / C++

Example parallel_master_taskloop.1.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #define N 100
S-4
S-5   int main()
S-6   {
S-7      int i, a[N],b[N],c[N];
S-8
S-9      for(int i=0; i<N; i++){ b[i]=i; c[i]=i; }  //init
S-10
S-11     #pragma omp parallel
S-12     #pragma omp master
S-13     #pragma omp taskloop                      // taskloop 1
S-14     for(i=0;i<N;i++){ a[i] = b[i] + c[i]; }
S-15
S-16     #pragma omp parallel master taskloop      // taskloop 2
S-17     for(i=0;i<N;i++){ b[i] = a[i] + c[i]; }
S-18
S-19     #pragma omp parallel master taskloop simd // taskloop 3
S-20     for(i=0;i<N;i++){ c[i] = a[i] + b[i]; }
S-21
S-22
```

```
S-23      printf(" %d %d\n",c[0],c[N-1]);  // 0 and 495
S-24   }
```

C / C++

Fortran

Example parallel_master_taskloop.1.f90 (`omp_5.0`)

```
S-1
S-2    program main
S-3
S-4       integer, parameter :: N=100
S-5       integer :: i, a(N),b(N),c(N)
S-6
S-7       do i=1,N                              !! initialize
S-8          b(i) = i
S-9          c(i) = i
S-10      enddo
S-11
S-12      !$omp parallel
S-13      !$omp master
S-14      !$omp taskloop                        !! taskloop 1
S-15      do i=1,N
S-16         a(i) = b(i) + c(i)
S-17      enddo
S-18      !$omp end taskloop
S-19      !$omp end master
S-20      !$omp end parallel
S-21
S-22      !$omp parallel master taskloop        !! taskloop 2
S-23      do i=1,N
S-24         b(i) = a(i) + c(i)
S-25      enddo
S-26      !$omp end parallel master taskloop
S-27
S-28      !$omp parallel master taskloop simd !! taskloop 3
S-29      do i=1,N
S-30         c(i) = a(i) + b(i)
S-31      enddo
S-32      !$omp end parallel master taskloop simd
S-33
S-34      print*,c(1),c(N)  !! 5 and 500
S-35
S-36   end program
```

Fortran

CHAPTER 4

Devices

The **target** construct consists of a **target** directive and an execution region. The **target** region is executed on the default device or the device specified in the **device** clause.

In OpenMP version 4.0, by default, all variables within the lexical scope of the construct are copied *to* and *from* the device, unless the device is the host, or the data exists on the device from a previously executed data-type construct that has created space on the device and possibly copied host data to the device storage.

The constructs that explicitly create storage, transfer data, and free storage on the device are catagorized as structured and unstructured. The **target data** construct is structured. It creates a data region around **target** constructs, and is convenient for providing persistent data throughout multiple **target** regions. The **target enter data** and **target exit data** constructs are unstructured, because they can occur anywhere and do not support a "structure" (a region) for enclosing **target** constructs, as does the **target data** construct.

The **map** clause is used on **target** constructs and the data-type constructs to map host data. It specifies the device storage and data movement **to** and **from** the device, and controls on the storage duration.

There is an important change in the OpenMP 4.5 specification that alters the data model for scalar variables and C/C++ pointer variables. The default behavior for scalar variables and C/C++ pointer variables in an 4.5 compliant code is **firstprivate**. Example codes that have been updated to reflect this new behavior are annotated with a description that describes changes required for correct execution. Often it is a simple matter of mapping the variable as **tofrom** to obtain the intended 4.0 behavior.

In OpenMP version 4.5 the mechanism for target execution is specified as occuring through a *target task*. When the **target** construct is encountered a new *target task* is generated. The *target task* completes after the **target** region has executed and all data transfers have finished.

This new specification does not affect the execution of pre-4.5 code; it is a necessary element for asynchronous execution of the **target** region when using the new **nowait** clause introduced in OpenMP 4.5.

4.1 `target` Construct

4.1.1 `target` Construct on `parallel` Construct

This following example shows how the **`target`** construct offloads a code region to a target device. The variables *p*, *v1*, *v2*, and *N* are implicitly mapped to the target device.

C / C++

Example target.1.c (**`omp_4.0`**)

```
S-1   extern void init(float*, float*, int);
S-2   extern void output(float*, int);
S-3   void vec_mult(int N)
S-4   {
S-5      int i;
S-6      float p[N], v1[N], v2[N];
S-7      init(v1, v2, N);
S-8      #pragma omp target
S-9      #pragma omp parallel for private(i)
S-10     for (i=0; i<N; i++)
S-11       p[i] = v1[i] * v2[i];
S-12     output(p, N);
S-13  }
```

C / C++

Fortran

Example target.1.f90 (**`omp_4.0`**)

```
S-1   subroutine vec_mult(N)
S-2      integer ::  i,N
S-3      real    ::  p(N), v1(N), v2(N)
S-4      call init(v1, v2, N)
S-5      !$omp target
S-6      !$omp parallel do
S-7      do i=1,N
S-8         p(i) = v1(i) * v2(i)
S-9      end do
S-10     !$omp end target
S-11     call output(p, N)
S-12  end subroutine
```

Fortran

4.1.2 `target` Construct with `map` Clause

This following example shows how the **`target`** construct offloads a code region to a target device. The variables *p*, *v1* and *v2* are explicitly mapped to the target device using the **`map`** clause. The variable *N* is implicitly mapped to the target device.

C / C++

Example target.2.c (**`omp_4.0`**)

```
S-1   extern void init(float*, float*, int);
S-2   extern void output(float*, int);
S-3   void vec_mult(int N)
S-4   {
S-5      int i;
S-6      float p[N], v1[N], v2[N];
S-7      init(v1, v2, N);
S-8      #pragma omp target map(v1, v2, p)
S-9      #pragma omp parallel for
S-10     for (i=0; i<N; i++)
S-11       p[i] = v1[i] * v2[i];
S-12     output(p, N);
S-13  }
```

C / C++

Fortran

Example target.2.f90 (**`omp_4.0`**)

```
S-1   subroutine vec_mult(N)
S-2      integer ::  i,N
S-3      real    ::  p(N), v1(N), v2(N)
S-4      call init(v1, v2, N)
S-5      !$omp target map(v1,v2,p)
S-6      !$omp parallel do
S-7      do i=1,N
S-8         p(i) = v1(i) * v2(i)
S-9      end do
S-10     !$omp end target
S-11     call output(p, N)
S-12  end subroutine
```

Fortran

4.1.3 `map` Clause with `to`/`from` map-types

The following example shows how the **`target`** construct offloads a code region to a target device. In the **`map`** clause, the **`to`** and **`from`** map-types define the mapping between the original (host) data and the target (device) data. The **`to`** map-type specifies that the data will only be read on the device, and the **`from`** map-type specifies that the data will only be written to on the device. By specifying a guaranteed access on the device, data transfers can be reduced for the **`target`** region.

The **`to`** map-type indicates that at the start of the **`target`** region the variables *v1* and *v2* are initialized with the values of the corresponding variables on the host device, and at the end of the **`target`** region the variables *v1* and *v2* are not assigned to their corresponding variables on the host device.

The **`from`** map-type indicates that at the start of the **`target`** region the variable *p* is not initialized with the value of the corresponding variable on the host device, and at the end of the **`target`** region the variable *p* is assigned to the corresponding variable on the host device.

C / C++

Example target.3.c (**`omp_4.0`**)

```
S-1   extern void init(float*, float*, int);
S-2   extern void output(float*, int);
S-3   void vec_mult(int N)
S-4   {
S-5      int i;
S-6      float p[N], v1[N], v2[N];
S-7      init(v1, v2, N);
S-8      #pragma omp target map(to: v1, v2) map(from: p)
S-9      #pragma omp parallel for
S-10     for (i=0; i<N; i++)
S-11       p[i] = v1[i] * v2[i];
S-12     output(p, N);
S-13  }
```

C / C++

The **`to`** and **`from`** map-types allow programmers to optimize data motion. Since data for the *v* arrays are not returned, and data for the *p* array are not transferred to the device, only one-half of the data is moved, compared to the default behavior of an implicit mapping.

Fortran

Example target.3.f90 (`omp_4.0`)

```
S-1   subroutine vec_mult(N)
S-2      integer ::  i,N
S-3      real    ::  p(N), v1(N), v2(N)
S-4      call init(v1, v2, N)
S-5      !$omp target map(to: v1,v2) map(from: p)
S-6      !$omp parallel do
S-7      do i=1,N
S-8         p(i) = v1(i) * v2(i)
S-9      end do
S-10     !$omp end target
S-11     call output(p, N)
S-12  end subroutine
```

Fortran

4.1.4 `map` Clause with Array Sections

The following example shows how the **`target`** construct offloads a code region to a target device. In the **`map`** clause, map-types are used to optimize the mapping of variables to the target device. Because variables *p*, *v1* and *v2* are pointers, array section notation must be used to map the arrays. The notation **`:N`** is equivalent to **`0:N`**.

C / C++

Example target.4.c (`omp_4.0`)

```
S-1   extern void init(float*, float*, int);
S-2   extern void output(float*, int);
S-3   void vec_mult(float *p, float *v1, float *v2, int N)
S-4   {
S-5      int i;
S-6      init(v1, v2, N);
S-7      #pragma omp target map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-8      #pragma omp parallel for
S-9      for (i=0; i<N; i++)
S-10       p[i] = v1[i] * v2[i];
S-11     output(p, N);
S-12  }
```

C / C++

In C, the length of the pointed-to array must be specified. In Fortran the extent of the array is known and the length need not be specified. A section of the array can be specified with the usual Fortran syntax, as shown in the following example. The value 1 is assumed for the lower bound for array section *v2(:N)*.

Fortran

Example target.4.f90 (**omp_4.0**)

```
S-1   module mults
S-2   contains
S-3   subroutine vec_mult(p,v1,v2,N)
S-4      real,pointer,dimension(:) :: p, v1, v2
S-5      integer                   :: N,i
S-6      call init(v1, v2, N)
S-7      !$omp target map(to: v1(1:N), v2(:N)) map(from: p(1:N))
S-8      !$omp parallel do
S-9      do i=1,N
S-10        p(i) = v1(i) * v2(i)
S-11     end do
S-12     !$omp end target
S-13     call output(p, N)
S-14  end subroutine
S-15  end module
```

Fortran

A more realistic situation in which an assumed-size array is passed to **vec_mult** requires that the length of the arrays be specified, because the compiler does not know the size of the storage. A section of the array must be specified with the usual Fortran syntax, as shown in the following example. The value 1 is assumed for the lower bound for array section *v2(:N)*.

Fortran

Example target.4b.f90 (**omp_4.0**)

```
S-1   module mults
S-2   contains
S-3   subroutine vec_mult(p,v1,v2,N)
S-4      real,dimension(*) :: p, v1, v2
S-5      integer           :: N,i
S-6      call init(v1, v2, N)
S-7      !$omp target map(to: v1(1:N), v2(:N)) map(from: p(1:N))
S-8      !$omp parallel do
S-9      do i=1,N
S-10        p(i) = v1(i) * v2(i)
S-11     end do
S-12     !$omp end target
S-13     call output(p, N)
```

```
end subroutine
end module
```

Fortran

4.1.5 `target` Construct with `if` Clause

The following example shows how the **target** construct offloads a code region to a target device.

The **if** clause on the **target** construct indicates that if the variable *N* is smaller than a given threshold, then the **target** region will be executed by the host device.

The **if** clause on the **parallel** construct indicates that if the variable *N* is smaller than a second threshold then the **parallel** region is inactive.

C / C++

Example target.5.c (**omp_4.0**)

```
#define THRESHOLD1 1000000
#define THRESHOLD2 1000

extern void init(float*, float*, int);
extern void output(float*, int);

void vec_mult(float *p, float *v1, float *v2, int N)
{
   int i;

   init(v1, v2, N);

   #pragma omp target if(N>THRESHOLD1) map(to: v1[0:N], v2[:N])\
        map(from: p[0:N])
   #pragma omp parallel for if(N>THRESHOLD2)
   for (i=0; i<N; i++)
     p[i] = v1[i] * v2[i];

   output(p, N);
}
```

C / C++

Fortran

Example target.5.f90 (**omp_4.0**)

```
S-1   module params
S-2   integer,parameter :: THRESHOLD1=1000000, THRESHHOLD2=1000
S-3   end module
S-4
S-5   subroutine vec_mult(p, v1, v2, N)
S-6      use params
S-7      real    ::  p(N), v1(N), v2(N)
S-8      integer ::  i
S-9
S-10     call init(v1, v2, N)
S-11
S-12     !$omp target if(N>THRESHHOLD1) map(to: v1, v2 ) map(from: p)
S-13        !$omp parallel do if(N>THRESHOLD2)
S-14        do i=1,N
S-15           p(i) = v1(i) * v2(i)
S-16        end do
S-17     !$omp end target
S-18
S-19     call output(p, N)
S-20  end subroutine
```

Fortran

The following example is a modification of the above *target.5* code to show the combined **target** and parallel loop directives. It uses the *directive-name* modifier in multiple **if** clauses to specify the component directive to which it applies.

The **if** clause with the **target** modifier applies to the **target** component of the combined directive, and the **if** clause with the **parallel** modifier applies to the **parallel** component of the combined directive.

C / C++

Example target.6.c (**omp_4.5**)

```
S-1   #define THRESHOLD1 1000000
S-2   #define THRESHOLD2 1000
S-3
S-4   extern void init(float*, float*, int);
S-5   extern void output(float*, int);
S-6
S-7   void vec_mult(float *p, float *v1, float *v2, int N)
S-8   {
S-9      int i;
S-10
S-11     init(v1, v2, N);
```

```
S-12
S-13        #pragma omp target parallel for \
S-14             if(target: N>THRESHOLD1) if(parallel: N>THRESHOLD2) \
S-15             map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-16        for (i=0; i<N; i++)
S-17          p[i] = v1[i] * v2[i];
S-18
S-19        output(p, N);
S-20    }
```

C / C++

Fortran

Example target.6.f90 (**omp_4.5**)

```
S-1     module params
S-2     integer,parameter :: THRESHOLD1=1000000, THRESHHOLD2=1000
S-3     end module
S-4
S-5     subroutine vec_mult(p, v1, v2, N)
S-6        use params
S-7        real    ::  p(N), v1(N), v2(N)
S-8        integer ::  i
S-9
S-10       call init(v1, v2, N)
S-11
S-12       !$omp target parallel do  &
S-13       !$omp&   if(target: N>THRESHHOLD1) if(parallel: N>THRESHOLD2) &
S-14       !$omp&   map(to: v1, v2 ) map(from: p)
S-15          do i=1,N
S-16             p(i) = v1(i) * v2(i)
S-17          end do
S-18       !$omp end target parallel do
S-19
S-20       call output(p, N)
S-21    end subroutine
```

Fortran

4.1.6 target Reverse Offload

Beginning with OpenMP 5.0, implementations are allowed to offload back to the host (reverse offload).

In the example below the *error_handler* function is executed back on the host, if an erroneous value is detected in the *A* array on the device.

This is accomplished by specifying the *device-modifier* **ancestor** modifier, along with a device number of **1**, to indicate that the execution is to be performed on the immediate parent (*1st ancestor*)– the host.

The **requires** directive (another 5.0 feature) uses the **reverse_offload** clause to guarantee that the reverse offload is implemented.

Note that the **declare target** directive uses the **device_type** clause (another 5.0 feature) to specify that the *error_handler* function is compiled to execute on the *host* only. This ensures that no attempt will be made to create a device version of the function. This feature may be necessary if the function exists in another compile unit.

C / C++

Example target_reverse_offload.7.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #include <stdlib.h>
S-4
S-5   #define N 100
S-6
S-7   #pragma omp requires reverse_offload
S-8
S-9   void error_handler(int wrong_value, int index)
S-10  {
S-11     printf(" Error in offload: A[%d]=%d\n", index,wrong_value);
S-12     printf("        Expecting: A[i ]=i\n");
S-13     exit(1);
S-14  // output:  Error in offload: A[99]=-1
S-15  //                Expecting: A[i ]=i
S-16
S-17  }
S-18  #pragma omp declare target device_type(host) to(error_handler)
S-19
S-20  int main()
S-21  {
S-22     int A[N];
S-23
S-24     for (int i=0; i<N; i++) A[i] = i;
S-25
```

```
   A[N-1]=-1;

   #pragma omp target map(A)
   {
      for (int i=0; i<N; i++)
      {
         if (A[i] != i)
         {
            #pragma omp target device(ancestor: 1) map(always,to: A[i:1])
               error_handler(A[i], i);
         }
      }
   }
   return 0;
}
```

C / C++

Fortran

Example target_reverse_offload.7.f90 (`omp_5.0`)

```
!$omp requires reverse_offload

subroutine error_handler(wrong_value, index)
  integer :: wrong_value,index
  !$omp declare target device_type(host)

   write( *,'("Error in offload: A(",i3,")=",i3)' ) index,wrong_value
   write( *,'("       Expecting: A(  i)=  i")'    )
   stop
     !!output: Error in offload: A( 99)= -1
     !!               Expecting: A(  i)=  i
end subroutine

program rev_off
  use omp_lib
  integer, parameter :: N=100
  integer            :: A(N) = (/ (i, i=1,100) /)

   A(N-1)=-1

   !$omp target map(A)
      do i=1,N
         if (A(i) /= i)  then
           !$omp target device(ancestor: 1) map(always,to :A(i))
               call error_handler(A(i), i)
           !$omp end target
```

```
S-28              endif
S-29            end do
S-30         !$omp end target
S-31
S-32      end program
```

Fortran

4.2 `defaultmap` Clause

The implicitly-determined, data-mapping and data-sharing attribute rules of variables referenced in a **`target`** construct can be changed by the **`defaultmap`** clause introduced in OpenMP 5.0. The implicit behavior is specified as **`alloc`**, **`to`**, **`from`**, **`tofrom`**, **`firstprivate`**, **`none`**, **`default`** or **`present`**, and is applied to a variable-category, where **`scalar`**, **`aggregate`**, **`allocatable`**, and **`pointer`** are the variable categories.

In OpenMP, a "category" has a common data-mapping and data-sharing behavior for variable types within the category. In C/C++, **`scalar`** refers to base-language scalar variables, except pointers. In Fortran it refers to a scalar variable, as defined by the base language, with intrinsic type, and excludes character type.

Also, **`aggregate`** refers to arrays and structures (C/C++) and derived types (Fortran). Fortran has the additional category of **`allocatable`**.

In the example below, the first **`target`** construct uses **`defaultmap`** clauses to set data-mapping and possibly data-sharing attributes that reproduce the default implicit mapping (data-mapping and data-sharing attributes). That is, if the **`defaultmap`** clauses were removed, the results would be identical.

In the second **`target`** construct all implicit behavior is removed by specifying the **`none`** implicit behavior in the **`defaultmap`** clause. Hence, all variables must be explicitly mapped. In the C/C++ code a scalar (`s`), an array (`A`) and a structure (`S`) are explicitly mapped **`tofrom`**. The Fortran code uses a derived type (`D`) in lieu of structure.

The third **`target`** construct shows another usual case for using the **`defaultmap`** clause. The default mapping for (non-pointer) scalar variables is specified as **`tofrom`**. Here, the default implicit mapping for `s3` is **`tofrom`** as specified in the **`defaultmap`** clause, and `s1` and `s2` are explicitly mapped with the **`firstprivate`** data-sharing attribute.

In the fourth **`target`** construct all arrays, structures (C/C++) and derived types (Fortran) are mapped with **`firstprivate`** data-sharing behavior by a **`defaultmap`** clause with an **`aggregate`** variable category. For the `H` allocated array in the Fortran code, the **`allocable`** category must be used in a separate **`defaultmap`** clause to acquire **`firsprivate`** data-sharing behavior (`H` has the Fortran allocatable attribute).

C / C++

Example target_defaultmap.1.c (**`omp_5.0`**)

```
S-1   #include <stdlib.h>
S-2   #include <stdio.h>
S-3   #define N 2
S-4
S-5   int main(){
S-6     typedef struct S_struct { int s; int A[N]; } S_struct_t;
S-7
S-8
```

```
S-9       int         s;       //scalar int variable (scalar)
S-10      int         A[N];    //aggregate variable  (array)
S-11      S_struct_t  S;       //aggregate variable  (structure)
S-12      int         *ptr;    //scalar, pointer variable (pointer)
S-13
S-14      int         s1, s2, s3;
S-15
S-16    // Initialize everything to zero;
S-17       s=2; s1=s2=s3=0;
S-18       A[0]=0; A[1]=0;
S-19       S.s=0; S.A[0]=0; S.A[1]=0;
S-20
S-21    // Target Region 1
S-22                       // Uses defaultmap to set scalars, aggregates & pointers
S-23                       // to normal defaults.
S-24        #pragma omp target \
S-25                defaultmap(firstprivate: scalar)     //could also be default \
S-26                defaultmap(tofrom:       aggregate)  //could also be default \
S-27                defaultmap(default:      pointer)    //must be default       \
S-28                map(ptr2m[:N])
S-29        {
S-30            s       = 3;               // SCALAR firstprivate, value not returned
S-31
S-32            A[0]    = 3;  A[1] = 3;    // AGGREGATE array, default map tofrom
S-33
S-34                                        // AGGREGATE structure, default tofrom
S-35            S.s     = 2;
S-36            S.A[0]  = 2;  S.A[1] = 2;
S-37
S-38            ptr = &A[0];                // POINTER is private
S-39            ptr[0] = 2;   ptr[1] = 2;
S-40
S-41        }
S-42       if(s==2 && A[0]==2 && S.s==2 && S.A[0]==2)
S-43          printf(" PASSED 1 of 4\n");
S-44
S-45
S-46    // Target Region 2
S-47                       // no implicit mapping allowed.
S-48        #pragma omp target defaultmap(none) map(tofrom: s, A, S)
S-49        {
S-50            s     +=5;           // All variables must be explicitly mapped
S-51            A[0]  +=5; A[1]+=5;
S-52            S.s   +=5;
S-53            S.A[0]+=5; S.A[1]+=5;
S-54        }
S-55        if(s==7 && A[0]==7 && S.s==7 && S.A[0]==7) printf(" PASSED 2 of 4\n");
```

```
S-56
S-57
S-58    // Target Region 3
S-59                     // defaultmap & explicit map with variables in same category
S-60        s1=s2=s3=1;
S-61        #pragma  omp defaultmap(tofrom: scalar) map(firstprivate: s1,s2)
S-62        {
S-63            s1 += 5;           // firstprivate (s1 value not returned to host)
S-64            s2 += 5;           // firstprivate (s2 value not returned to host)
S-65            s3 += s1 + s2;     // mapped as tofrom
S-66        }
S-67        if(s1==1 && s2==1 && s3==13 ) printf(" PASSED 3 of 4\n");
S-68
S-69
S-70    // Target Region 4
S-71        A[0]=0; A[1]=0;
S-72        S.A[0]=0; S.A[1]=0;
S-73
S-74        // arrays and structure are firstprivate, and scalars are from
S-75        #pragma omp target defaultmap(firstprivate: aggregate) map(from: s1, s2)
S-76        {
S-77
S-78            A[0]+=1; S.A[0]+=1; //Aggregate changes not returned to host
S-79            A[1]+=1; S.A[1]+=1; //Aggregate changes not returned to host
S-80            s1 = A[0]+S.A[0]; //s1 value returned to host
S-81            s2 = A[1]+S.A[1]; //s1 value returned to host
S-82        }
S-83        if( A[0]==0 && S.A[0]==0 && s1==2 ) printf(" PASSED 4 of 4\n");
S-84
S-85    }
```

C / C++

Fortran

1 *Example target_defaultmap.1.f90* (`omp_5.0`)

```
S-1     program defaultmap
S-2       integer, parameter :: N=2
S-3
S-4       type DDT_sA
S-5         integer  :: s
S-6         integer  :: A(N)
S-7       end type
S-8
S-9       integer             :: s,s1,s2,s3 !! SCALAR: variable (integer)
S-10
S-11      integer,target      :: A(N)       !! AGGREGATE: Array
S-12      type(DDT_sA)        :: D          !! AGGREGATE: Derived Data Type (D)
```

```
S-13
S-14      integer,allocatable :: H(:)       !! ALLOCATABLE: Heap allocated array
S-15
S-16      integer,pointer     :: ptrA(:)    !! POINTER: points to Array
S-17
S-18      ! Assign vaues to scalar, Array, Allocatable, and Pointers
S-19
S-20        s=2;
S-21        s1=0;   s2=0;     s3=0
S-22        D%s=0;  D%A(1)=0; D%A(2)=0
S-23        A(1)=0; A(2)=0
S-24
S-25        allocate( H(2) )
S-26        H(1)=0; H(2)=0
S-27
S-28     !! Target Region 1
S-29                        !! Using defaultmap to set scalars, aggregates & pointers
S-30                        !! and allocatables to normal defaults.
S-31        !$omp target                                          &
S-32        !$omp&       defaultmap( firstprivate: scalar)      &
S-33        !$omp&       defaultmap( tofrom:       aggregate)   &
S-34        !$omp&       defaultmap( tofrom:       allocatable) &
S-35        !$omp&       defaultmap( default:      pointer)
S-36
S-37            s = 3                     !! SCALAR firstprivate, value not returned
S-38
S-39            A(1) = 3                  !! AGGREGATE array, default map tofrom
S-40            A(2) = 3
S-41
S-42            D%s = 2                   !! AGGR. Derived Types, default map tofrom
S-43            D%A(1) = 2;  D%A(2) = 2
S-44
S-45            H(1) = 2;    H(2) = 2     !! ALLOCATABLE, default map tofrom
S-46
S-47            ptrA=>A                   !! POINTER is private
S-48            ptrA(1) = 2; ptrA(2) = 2
S-49
S-50        !$omp end target
S-51
S-52        if(s==2 .and. A(1)==2 .and. D%s==2 .and. D%A(1)==2 .and. H(1) == 2) &
S-53           print*," PASSED 1 of 4"
S-54
S-55     !! Target Region 2
S-56                        !! no implicit mapping allowed
S-57        !$omp target defaultmap(none) map(tofrom: s, A, D)
S-58
S-59            s=s+5                     !! All variables must be explicitly mapped
```

```
S-60            A(1)=A(1)+5;       A(2)=A(2)+5
S-61            D%s=D%s+5
S-62            D%A(1)=D%A(1)+5; D%A(2)=D%A(2)+5
S-63
S-64        !$omp end target
S-65        if(s==7 .and. A(1)==7 .and. D%s==7 .and. D%A(1)==7) print*," PASSED 2 of 4"
S-66
S-67    !! Target Region 3
S-68                        !!defaultmap & explicit map with variables in same category
S-69        s1=1; s2=1; s3=1
S-70        !$omp defaultmap(tofrom: scalar) map(firstprivate: s1,s2)
S-71
S-72            s1 = s1+5;           !! firstprivate (s1 value not returned to host)
S-73            s2 = s2+5;           !! firstprivate (s2 value not returned to host)
S-74            s3 = s3 +s1 + s2;    !! mapped as tofrom
S-75
S-76        !$omp end target
S-77        if(s1==1 .and. s2==1 .and. s3==13) print*," PASSED 3 of 4"
S-78
S-79    !! Target Region 4
S-80        A(1)=0;   A(2)=0
S-81        D%A(1)=0; D%A(2)=0
S-82        H(1)=0;   H(2)=0
S-83                    !! non-allocated arrays & derived types are in AGGREGATE cat.
S-84                    !! Allocatable Arrays are in ALLOCATABLE category
S-85                    !! Scalars are explicitly mapped from
S-86        !$omp target defaultmap(firstprivate: aggregate )  &
S-87        !$omp&       defaultmap(firstprivate: allocatable) &
S-88        !$omp&       map(from: s1, s2)
S-89
S-90            A(1)=A(1)+1; D%A(1)=D%A(1)+1; H(1)=H(1)+1 !!changes not returned to host
S-91            A(2)=A(2)+1; D%A(2)=D%A(2)+1; H(2)=H(2)+1 !!changes not returned to host
S-92            s1 = A(1)+D%A(1)+H(1)                     !!s1 returned to host
S-93            s2 = A(2)+D%A(2)+H(1)                     !!s2 returned to host
S-94
S-95        !$omp end target
S-96        if(A(1)==0 .and. D%A(1)==0 .and. H(1)==0 .and. s1==3) &
S-97           print*," PASSED 4 of 4"
S-98
S-99    end program
```

Fortran

4.3 Pointer mapping

Pointers that contain host addresses require that those addresses are translated to device addresses for them to be useful in the context of a device data environment. Broadly speaking, there are two scenarios where this is important.

The first scenario is where the pointer is mapped to the device data environment, such that references to the pointer inside a **target** region are to the corresponding pointer. Pointer attachment ensures that the corresponding pointer will contain a device address when all of the following conditions are true:

- the pointer is mapped by directive A to a device;
- a list item that uses the pointer as its base pointer (call it the *pointee*) is mapped, to the same device, by directive B, which may be the same as A;
- the effect of directive B is to create either the corresponding pointer or pointee in the device data environment of the device.

Given the above conditions, pointer attachment is initiated as a result of directive B and subsequent references to the pointee list item in a target region that use the pointer will access the corresponding pointee. The corresponding pointer remains in this *attached* state until it is removed from the device data environment.

The second scenario, which is only applicable for C/C++, is where the pointer is implicitly privatized inside a **target** construct when it appears as the base pointer to a list item on the construct and does not appear explicitly as a list item in a **map** clause, **is_device_ptr** clause, or data-sharing attribute clause. This scenario can be further split into two cases: the list item is a zero-length array section (e.g., *p[:0]*) or it is not.

If it is a zero-length array section, this will trigger a runtime check on entry to the **target** region for a previously mapped list item where the value of the pointer falls within the range of its base address and ending address. If such a match is found the private pointer is initialized to the device address corresponding to the value of the original pointer, and otherwise it is initialized to NULL (or retains its original value if the **unified_address** requirement is specified for that compilation unit).

If the list item (again, call it the *pointee*) is not a zero-length array section, the private pointer will be initialized such that references in the **target** region to the pointee list item that use the pointer will access the corresponding pointee.

The following example shows the basics of mapping pointers with and without associated storage on the host.

Storage for pointers *ptr1* and *ptr2* is created on the host. To map storage that is associated with a pointer on the host, the data can be explicitly mapped as an array section so that the compiler knows the amount of data to be assigned in the device (to the "corresponding" data storage area). On the **target** construct array sections are mapped; however, the pointer *ptr1* is mapped, while *ptr2* is

not. Since *ptr2* is not explicitly mapped, it is firstprivate. This creates a subtle difference in the way these pointers can be used.

As a firstprivate pointer, *ptr2* can be manipulated on the device; however, as an explicitly mapped pointer, *ptr1* becomes an *attached* pointer and cannot be manipulated. In both cases the host pointer is not updated with the device pointer address—as one would expect for distributed memory. The storage data on the host is updated from the corresponding device data at the end of the **target** region.

As a comparison, note that the *aray* array is automatically mapped, since the compiler knows the extent of the array.

The pointer *ptr3* is used inside the **target** construct, but it does not appear in a data-mapping or data-sharing clause. Nor is there a **defaultmap** clause on the construct to indicate what its implicit data-mapping or data-sharing attribute should be. For such a case, *ptr3* will be implicitly privatized within the construct and there will be a runtime check to see if the host memory to which it is pointing has corresponding memory in the device data environment. If this runtime check passes, the private *ptr3* would be initialized to point to the corresponding memory. But in this case the check does not pass and so it is initialized to null. Since *ptr3* is private, the value to which it is assigned in the **target** region is not returned into the original *ptr3* on the host.

C / C++

Example target_ptr_map.1.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   #include <stdlib.h>
S-3   #define N 100
S-4
S-5   int main()
S-6   {
S-7     int *ptr1;
S-8     int *ptr2;
S-9     int *ptr3;
S-10    int aray[N];
S-11
S-12    ptr1 = (int *)malloc(sizeof(int)*N);
S-13    ptr2 = (int *)malloc(sizeof(int)*N);
S-14
S-15    #pragma omp target map(ptr1, ptr1[:N]) map(ptr2[:N] )
S-16    {
S-17       for (int i=0; i<N; i++)
S-18       {
S-19           ptr1[i] = i;
S-20           ptr2[i] = i;
S-21           aray[i] = i;
S-22       }
S-23
```

```
S-24        //*(++ptr1) = 9;  //NOT ALLOWED since ptr1 is an attached pointer
S-25          *(++ptr2) = 9;  //    allowed since ptr2 is firstprivate
S-26
S-27          ptr3=(int *)malloc(sizeof(int)*N); // ptr3 is firstprivate
S-28                                             // ptr3 value not returned
S-29          for (int i=0; i<N; i++) ptr3[i] = 5;
S-30
S-31          for (int i=0; i<N; i++) ptr1[i] += ptr3[i];
S-32
S-33          free(ptr3);     // explicitly free allocated storage on device
S-34        }
S-35
S-36        printf(" %d %d\n",ptr1[1],ptr2[1]);
S-37        //        6  9
S-38
S-39        free(ptr1);
S-40        free(ptr2);
S-41        return 0;
S-42      }
```

C / C++

In the following example the global pointer *p* appears in a **declare target** directive. Hence, the pointer *p* will persist on the device throughout executions in all **target** regions.

The pointer is also used in an array section of a **map** clause on a **target** construct. When storage associated with a **declare target** pointer is mapped, as for the array section *p[:N]* in the **target** construct, the array section on the device is *attached* to the device pointer *p* on entry to the construct, and the value of the device pointer *p* becomes undefined on exit. (Of course, storage allocation for the array section on the device will occur before the pointer on the device is *attached*.)

C / C++

Example target_ptr_map.2.c (**omp_5.0**)

```
S-1     #include <stdio.h>
S-2     #include <stdlib.h>
S-3     #define N 100
S-4
S-5     #pragma omp declare target
S-6     int *p;
S-7     extern void use_arg_p(int *p, int n);
S-8     extern void use_global_p(     int n);
S-9     #pragma omp end declare target
S-10
S-11    int main()
S-12    {
S-13      int i;
S-14      p = (int *)malloc(sizeof(int)*N);
```

```
S-15
S-16      #pragma omp target map(p[:N])  // device p attached to array section
S-17      {
S-18        for (i=0; i<N; i++) p[i] = i;
S-19        use_arg_p(p, N);
S-20        use_global_p(N);
S-21      }                               // value of host p is preserved
S-22
S-23      printf(" %3.3d %3.3d\n", p[1], p[N-1]);
S-24             // 003   297   <- output
S-25
S-26      free(p);
S-27      return 0;
S-28    }
S-29
S-30    //#pragma omp declare target (optional here because of prototype spec)
S-31    void use_arg_p(int *q, int n)
S-32    {
S-33      int i;
S-34      for (i=0; i<n; i++)
S-35        q[i] *= 2;
S-36    }
S-37
S-38    void use_global_p(int n)
S-39    {
S-40      int i;
S-41      for (i=0; i<n; i++)
S-42        p[i] += i;   // valid since p is in declare target and called from
S-43                     // inside target region where p was attached to valid memory
S-44    }
S-45    //#pragma omp end declare target (optional here because of prototype spec)
```

C / C++

The following two examples illustrate subtle differences in pointer attachment to device address because of the order of data mapping.

In example *target_ptr_map.3a* the global pointer *p1* points to array *x* and *p2* points to array *y* on the host. The array section *x[:N]* is mapped by the **target enter data** directive while array *y* is mapped on the **target** construct. Since the **declare target** directive is applied to the declaration of *p1*, *p1* is a treated like a mapped variable on the **target** construct and references to *p1* inside the construct will be to the corresponding *p1* that exists on the device. However, the corresponding *p1* will be undefined since there is no pointer attachment for it. Pointer attachment for *p1* would require that (1) *p1* (or an lvalue expression that refers to the same storage as *p1*) appears as a base pointer to a list item in a **map** clause, and (2) the construct that has the **map** clause causes the list item to transition from *not mapped* to *mapped*. The conditions are clearly not satisifed for this example.

The problem for *p2* in this example is also subtle. It will be privatized inside the **target** construct, with a runtime check for whether the memory to which it is pointing has corresponding memory that is accessible on the device. If this check is successful then the *p2* inside the construct would be appropriately initialized to point to that corresponding memory. Unfortunately, despite there being an implicit map of the array *y* (to which *p2* is pointing) on the construct, the order of this map relative to the initialization of *p2* is unspecified. Therefore, the initial value of *p2* will also be undefined.

Thus, referencing values via either *p1* or *p2* inside the **target** region would be invalid.

C / C++

Example target_ptr_map.3a.c (**omp_5.0**)

```
#define N 100

int x[N], y[N];
#pragma omp declare target
int *p1;
#pragma omp end declare target
int *p2;

int foo()
{
  p1 = &x[0];
  p2 = &y[0];

  // Explicitly map array section x[:N]
  #pragma omp target enter data map(x[:N])

  #pragma omp target  // as if .. map(p1) map(p1[:0]) map(p2[:0]) map(y)
  {
    // Accessing the mapped arrays x,y is OK here.
    x[0] = 1;
    y[1] = 2;

    // Pointer attachment for p1 does not occur here
    //   because p1[:0] does not allocate a new array section and
    //   array x is present on the target construct as it was mapped
    //   before by the target enter data directive.
    p1[0] = 3;      // accessing p1 is undefined

    // The initial value of p2 in the target region is undefined
    //   because map(y) may occur after map(p2[:0]).
    p2[1] = 4;      // accessing p2 is undefined
  }

```

```
  return 0;
}
```

C / C++

In example *target_ptr_map.3b* the mapping orders for arrays *x* and *y* were rearranged to allow proper pointer attachments. On the **target** construct, the **map(x)** clause triggers pointer attachment for *p1* to the device address of *x*. Pointer *p2* is assigned the device address of the previously mapped array *y*. Referencing values via either *p1* or *p2* inside the **target** region is now valid.

C / C++

Example target_ptr_map.3b.c (**omp_5.0**)

```
#define N 100

int x[N], y[N];
#pragma omp declare target
int *p1;
#pragma omp end declare target
int *p2;

int foo()
{
  p1 = &x[0];
  p2 = &y[0];

  // Explicitly map array section y[:N]
  #pragma omp target enter data map(y[:N])

  #pragma omp target map(x[:N]) map(p1[:N]) map(p2[:0])
  {
    // Accessing the mapped arrays x,y is OK here.
    x[0] = 1;
    y[1] = 2;

    // Pointer attachment for p1 occurs here when array x is mapped
    //   on the target construct (as p1 = &x[0] on the device)
    p1[0] = 3;      // accessing p1 is OK

    // p2 in the target region is initialized to &y[0]
    p2[1] = 4;      // accessing p2 is OK
  }

  return 0;
}
```

C / C++

4.4 Structure mapping

In the example below, only structure elements *S.a*, *S.b* and *S.p* of the *S* structure appear in **map** clauses of a **target** construct. Only these components have corresponding variables and storage on the device. Hence, the large arrays, *S.buffera* and *S.bufferb*, and the *S.x* component have no storage on the device and cannot be accessed.

Also, since the pointer member *S.p* is used in an array section of a **map** clause, the array storage of the array section on the device, *S.p[:N]*, is *attached* to the pointer member *S.p* on the device. Explicitly mapping the pointer member *S.p* is optional in this case.

Note: The buffer arrays and the *x* variable have been grouped together, so that the components that will reside on the device are all together (without gaps). This allows the runtime to optimize the transfer and the storage footprint on the device.

C / C++

Example target_struct_map.1.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   #include <stdlib.h>
S-3   #define N 100
S-4   #define BAZILLION 2000000
S-5
S-6   struct foo {
S-7     char buffera[BAZILLION];
S-8     char bufferb[BAZILLION];
S-9     float x;
S-10    float a, b;
S-11    float *p;
S-12  };
S-13
S-14  #pragma omp declare target
S-15  void saxpyfun(struct foo *S)
S-16  {
S-17    int i;
S-18    for(i=0; i<N; i++)
S-19      S->p[i] = S->p[i]*S->a + S->b;
S-20  }
S-21  #pragma omp end declare target
S-22
S-23  int main()
S-24  {
S-25    struct foo S;
S-26    int i;
S-27
S-28    S.a = 2.0;
S-29    S.b = 4.0;
```

```
S-30      S.p = (float *)malloc(sizeof(float)*N);
S-31      for(i=0; i<N; i++) S.p[i] = i;
S-32
S-33      #pragma omp target map(alloc:S.p) map(S.p[:N]) map(to:S.a, S.b)
S-34      saxpyfun(&S);
S-35
S-36      printf(" %4.0f %4.0f\n", S.p[0], S.p[N-1]);
S-37            //     4  202  <- output
S-38      return 0;
S-39    }
```

C / C++

The following example is a slight modification of the above example for a C++ class. In the member function *SAXPY::driver* the array section *p[:N]* is *attached* to the pointer member *p* on the device.

C++

Example target_struct_map.2.cpp (`omp_5.0`)

```
S-1     #include <cstdio>
S-2     #include <cstdlib>
S-3     #define N 100
S-4
S-5     class SAXPY {
S-6       private:
S-7        float a, b, *p;
S-8       public:
S-9        float buffer[N];
S-10
S-11       SAXPY(float arg_a, float arg_b){ a=arg_a; b=arg_b; }
S-12       void driver();
S-13       void saxpyfun(float *p);
S-14    };
S-15
S-16    #pragma omp declare target
S-17    void SAXPY::saxpyfun(float *q)
S-18    {
S-19      for(int i=0; i<N; i++)
S-20        buffer[i] = q[i]*a + b;
S-21    }
S-22    #pragma omp end declare target
S-23
S-24    void SAXPY::driver()
S-25    {
S-26      p = (float *) malloc(N*sizeof(float));
S-27      for(int i=0; i<N; i++) p[i]=i;
S-28
S-29      #pragma omp target map(alloc:p) map(to:p[:N]) map(to:a,b) \
```

```
S-30                  map(from:buffer[:N])   // attach(p) to device_malloc()
S-31      {
S-32        saxpyfun(p);
S-33      }
S-34
S-35      free(p);
S-36    }
S-37
S-38    int main()
S-39    {
S-40      SAXPY my_saxpy(2.0,4.0);
S-41
S-42      my_saxpy.driver();
S-43
S-44      printf(" %4.0f %4.0f\n", my_saxpy.buffer[0], my_saxpy.buffer[N-1]);
S-45             //   4   202     <- output
S-46
S-47      return 0;
S-48    }
```

C++

4.5 Array Sections in Device Constructs

The following examples show the usage of array sections in **map** clauses on **target** and **target data** constructs.

This example shows the invalid usage of two separate sections of the same array inside of a **target** construct.

C / C++

Example array_sections.1.c (**omp_4.0**)

```
S-1   void foo ()
S-2   {
S-3      int A[30];
S-4   #pragma omp target data map( A[0:4] )
S-5   {
S-6      /* Cannot map distinct parts of the same array */
S-7      #pragma omp target map( A[7:20] )
S-8      {
S-9         A[2] = 0;
S-10     }
S-11  }
S-12  }
```

C / C++

Fortran

Example array_sections.1.f90 (**omp_4.0**)

```
S-1   subroutine foo()
S-2   integer :: A(30)
S-3      A = 1
S-4      !$omp target data map( A(1:4) )
S-5        ! Cannot map distinct parts of the same array
S-6        !$omp target map( A(8:27) )
S-7           A(3) = 0
S-8        !$omp end target
S-9      !$omp end target data
S-10  end subroutine
```

Fortran

This example shows the invalid usage of two separate sections of the same array inside of a **target** construct.

C / C++

Example array_sections.2.c (**omp_4.0**)

```
S-1   void foo ()
S-2   {
S-3      int A[30], *p;
S-4   #pragma omp target data map( A[0:4] )
S-5   {
S-6      p = &A[0];
S-7      /* invalid because p[3] and A[3] are the same
S-8       * location on the host but the array section
S-9       * specified via p[...] is not a subset of A[0:4] */
S-10     #pragma omp target map( p[3:20] )
S-11     {
S-12        A[2] = 0;
S-13        p[8] = 0;
S-14     }
S-15  }
S-16  }
```

C / C++

Fortran

Example array_sections.2.f90 (**omp_4.0**)

```
S-1   subroutine foo()
S-2   integer,target  :: A(30)
S-3   integer,pointer :: p(:)
S-4      A=1
S-5      !$omp target data map( A(1:4) )
S-6        p=>A
S-7        ! invalid because p(4) and A(4) are the same
S-8        ! location on the host but the array section
S-9        ! specified via p(...) is not a subset of A(1:4)
S-10       !$omp target map( p(4:23) )
S-11          A(3) = 0
S-12          p(9) = 0
S-13       !$omp end target
S-14     !$omp end target data
S-15  end subroutine
```

Fortran

This example shows the valid usage of two separate sections of the same array inside of a **target** construct.

C / C++

Example array_sections.3.c (**omp_4.0**)

```
S-1   void foo ()
S-2   {
S-3      int A[30], *p;
S-4   #pragma omp target data map( A[0:4] )
S-5   {
S-6      p = &A[0];
S-7      #pragma omp target map( p[7:20] )
S-8      {
S-9         A[2] = 0;
S-10        p[8] = 0;
S-11     }
S-12  }
S-13  }
```

C / C++

Fortran

Example array_sections.3.f90 (**omp_4.0**)

```
S-1   subroutine foo()
S-2   integer,target  :: A(30)
S-3   integer,pointer :: p(:)
S-4      !$omp target data map( A(1:4) )
S-5        p=>A
S-6        !$omp target map( p(8:27) )
S-7           A(3) = 0
S-8           p(9) = 0
S-9        !$omp end target
S-10     !$omp end target data
S-11  end subroutine
```

Fortran

This example shows the valid usage of a wholly contained array section of an already mapped array section inside of a **target** construct.

C / C++

Example array_sections.4.c (**omp_4.0**)

```
S-1   void foo ()
S-2   {
S-3      int A[30], *p;
S-4   #pragma omp target data map( A[0:10] )
S-5   {
S-6      p = &A[0];
S-7      #pragma omp target map( p[3:7] )
S-8      {
S-9         A[2] = 0;
S-10        p[8] = 0;
S-11        A[8] = 1;
S-12     }
S-13  }
S-14  }
```

C / C++

Fortran

Example array_sections.4.f90 (**omp_4.0**)

```
S-1   subroutine foo()
S-2   integer,target  :: A(30)
S-3   integer,pointer :: p(:)
S-4      !$omp target data map( A(1:10) )
S-5        p=>A
S-6        !$omp target map( p(4:10) )
S-7           A(3) = 0
S-8           p(9) = 0
S-9           A(9) = 1
S-10       !$omp end target
S-11     !$omp end target data
S-12  end subroutine
```

Fortran

4.6 Array Shaping

C / C++

A pointer variable can be shaped to a multi-dimensional array to facilitate data access. This is achieved by a *shape-operator* casted in front of a pointer (lvalue expression):

([s_1]**[**s_2**]...[**s_n**])***pointer*

where each s_i is an integral-type expression of positive value. The shape-operator can appear in either the *motion-clause* of the **target update** directive or the **depend** clause.

The following example shows the use of the shape-operator in the **target update** directive. The shape-operator **([nx][ny+2])** casts pointer variable a to a 2-dimentional array of size *nx×(ny+2)*. The resulting array is then accessed as array sections (such as **[0:nx][1]** and **[0:nx][ny]**) in the **from** or **to** clause for transferring two columns of noncontiguous boundary data from or to the device. Note the use of additional parentheses around the shape-operator and a to ensure the correct precedence over array-section operations.

Example array_shaping.1.c (**omp_5.0**)

```
S-1   #pragma omp declare target
S-2   int do_work(double *a, int nx, int ny);
S-3   int other_work(double *a, int nx, int ny);
S-4   #pragma omp end declare target
S-5
S-6   void exch_data(double *a, int nx, int ny);
S-7
S-8   void array_shaping(double *a, int nx, int ny)
S-9   {
S-10     // map data to device and do work
S-11     #pragma omp target data map(a[0:nx*(ny+2)])
S-12     {
S-13        // do work on the device
S-14        #pragma omp target  // map(a[0:nx*(ny+2)]) is optional here
S-15        do_work(a, nx, ny);
S-16
S-17        // update boundary points (two columns of 2D array) on the host
S-18        // pointer is shaped to 2D array using the shape-operator
S-19        #pragma omp target update from( (([nx][ny+2])a)[0:nx][1], \
S-20                                        (([nx][ny+2])a)[0:nx][ny] )
S-21
S-22        // exchange ghost points with neighbors
S-23        exch_data(a, nx, ny);
S-24
S-25        // update ghost points (two columns of 2D array) on the device
S-26        // pointer is shaped to 2D array using the shape-operator
S-27        #pragma omp target update to( (([nx][ny+2])a)[0:nx][0], \
```

```
S-28                                                (([nx][ny+2])a)[0:nx][ny+1] )
S-29
S-30          // perform other work on the device
S-31          #pragma omp target  // map(a[0:nx*(ny+2)]) is optional here
S-32          other_work(a, nx, ny);
S-33       }
S-34    }
```

C / C++

The shape operator is not defined for Fortran. Explicit array shaping of procedure arguments can be used instead to achieve a similar goal. Below is the Fortran-equivalent of the above example that illustrates the support of transferring two rows of noncontiguous boundary data in the **target update** directive.

Fortran

Example array_shaping.1.f90 (**omp_5.0**)

```
S-1     module m
S-2        interface
S-3           subroutine do_work(a, nx, ny)
S-4              !$omp declare target to(do_work)
S-5              integer, intent(in) :: nx, ny
S-6              double precision a(0:nx+1,ny)
S-7           end subroutine do_work
S-8
S-9           subroutine other_work(a, nx, ny)
S-10             !$omp declare target to(other_work)
S-11             integer, intent(in) :: nx, ny
S-12             double precision a(0:nx+1,ny)
S-13          end subroutine other_work
S-14
S-15          subroutine exch_data(a, nx, ny)
S-16             integer, intent(in) :: nx, ny
S-17             double precision a(0:nx+1,ny)
S-18          end subroutine exch_data
S-19       end interface
S-20    end module m
S-21
S-22    subroutine array_shaping(a, nx, ny)
S-23       use m
S-24       implicit none
S-25       integer, intent(in) :: nx, ny
S-26       double precision a(0:nx+1,ny)
S-27
S-28       ! map data to device and do work
S-29       !$omp target data map(a)
S-30
```

```
S-31          ! do work on the device
S-32          !$omp target      ! map(a) is optional here
S-33          call do_work(a, nx, ny)
S-34          !$omp end target
S-35
S-36          ! update boundary points (two rows of 2D array) on the host.
S-37          ! data transferred are noncontiguous
S-38          !$omp target update from( a(1,1:ny), a(nx,1:ny) )
S-39
S-40          ! exchange ghost points with neighbors
S-41          call exch_data(a, nx, ny)
S-42
S-43          ! update ghost points (two rows of 2D array) on the device.
S-44          ! data transferred are noncontiguous
S-45          !$omp target update to( a(0,1:ny), a(nx+1,1:ny) )
S-46
S-47          ! perform other work on the device
S-48          !$omp target      ! map(a) is optional here
S-49          call other_work(a, nx, ny)
S-50          !$omp end target
S-51
S-52       !$omp end target data
S-53
S-54    end subroutine
```

Fortran

4.7 `declare mapper` Construct

The following examples show how to use the **`declare mapper`** directive to prescribe a map for later use. It is also quite useful for pre-defining partitioned and nested structure elements.

In the first example the **`declare mapper`** directive specifies that any structure of type *myvec_t* for which implicit data-mapping rules apply will be mapped according to its **`map`** clause. The variable *v* is used for referencing the structure and its elements within the **`map`** clause. Within the **`map`** clause the *v* variable specifies that all elements of the structure are to be mapped. Additionally, the array section *v.data[0:v.len]* specifies that the dynamic storage for data is to be mapped.

Within the main program the *s* variable is typed as *myvec_t*. Since the variable is found within the target region and the type has a mapping prescribed by a **`declare mapper`** directive, it will be automatically mapped according to its prescription: full structure, plus the dynamic storage of the *data* element.

C / C++

Example target_mapper.1.c (**`omp_5.0`**)

```
S-1   #include  <stdlib.h>
S-2   #include   <stdio.h>
S-3   #define N 100
S-4
S-5   typedef struct myvec{
S-6       size_t len;
S-7       double *data;
S-8   } myvec_t;
S-9
S-10  #pragma omp declare mapper(myvec_t v) \
S-11                      map(v, v.data[0:v.len])
S-12  void init(myvec_t *s);
S-13
S-14  int main(){
S-15     myvec_t s;
S-16
S-17     s.data = (double *)calloc(N,sizeof(double));
S-18     s.len  = N;
S-19
S-20     #pragma omp target
S-21     init(&s);
S-22
S-23     printf("s.data[%d]=%lf\n",N-1,s.data[N-1]);  //s.data[99]=99.000000
S-24  }
S-25
S-26  void init(myvec_t *s)
S-27  { for(int i=0; i<s->len; i++) s->data[i]=i; }
```

C / C++

Fortran

Example target_mapper.1.f90 (**omp_5.0**)

```
S-1
S-2   module my_structures
S-3     type myvec_t
S-4       integer                      :: len
S-5       double precision, pointer    :: data(:)
S-6     end type
S-7   end module
S-8
S-9   program main
S-10    use my_structures
S-11    integer, parameter :: N=100
S-12
S-13    !$omp  declare mapper(myvec_t :: v) &
S-14    !$omp&         map(v, v%data(1:v%len))
S-15
S-16    type(myvec_t) :: s
S-17
S-18    allocate(s%data(N))
S-19    s%data(1:N) = 0.0d0
S-20    s%len = N
S-21
S-22    !$omp target
S-23    call init(s)
S-24    !$omp end target
S-25
S-26    print*,"s%data(",N,")=",s%data(N)  !! s%data( 100 )=100.000000000000
S-27  end program
S-28
S-29  subroutine init(s)
S-30    use my_structures
S-31    type(myvec_t) :: s
S-32
S-33    s%data = [ (i, i=1,s%len) ]
S-34  end subroutine
```

Fortran

The next example illustrates the use of the *mapper-identifier* and deep copy within a structure. The structure, *dzmat_t*, represents a complex matrix, with separate real (*r_m*) and imaginary (*i_m*) elements. Two map identifiers are created for partitioning the *dzmat_t* structure.

For the C/C++ code the first identifier is named *top_id* and maps the top half of two matrices of type *dzmat_t*; while the second identifier, *bottom_id*, maps the lower half of two matrices. Each identifier is applied to a different **target** construct, as **map(mapper(top_id), tofrom: a,b)** and **map(mapper(bottom_id), tofrom: a,b)**. Each target offload is allowed to

execute concurrently on two different devices (*0* and *1*) through the **nowait** clause. The OpenMP 5.0 **parallel master** construct creates a region of two threads for these **target** constructs, with a single thread (*master*) generator.

The Fortran code uses the *left_id* and *right_id* map identifiers in the **map(mapper(left_id),tofrom: a,b)** and **map(mapper(right_id),tofrom: a,b)** map clauses. The array sections for these left and right contiguous portions of the matrices were defined previously in the **declare mapper** directive.

Note, the *is* and *ie* scalars are firstprivate by default for a target region, but are declared firstprivate anyway to remind the user of important firstprivate data-sharing properties required here.

C / C++

Example target_mapper.2.c (**omp_5.0**)

```
#include <stdio.h>
//                    N MUST BE EVEN
#define N  100

  typedef struct dzmat
  {
     double r_m[N][N];
     double i_m[N][N];
  } dzmat_t;

  #pragma omp declare mapper( top_id: dzmat_t v) \
                      map(v.r_m[0:N/2][0:N],     \
                          v.i_m[0:N/2][0:N]      )

  #pragma omp declare mapper(bottom_id: dzmat_t v) \
                      map(v.r_m[N/2:N/2][0:N],     \
                          v.i_m[N/2:N/2][0:N]      )

void dzmat_init(dzmat_t *z, int is, int ie, int n);          //initialization
void host_add(  dzmat_t *a, dzmat_t *b, dzmat_t *c, int n); //matrix add: c=a+b


int main()
{
  dzmat_t a,b,c;
  int     is,ie;

  is=0; ie=N/2-1;       //top N/2 rows on device 0
  #pragma omp target map(mapper(top_id), tofrom: a,b) device(0) \
                      firstprivate(is,ie) nowait
  {
    dzmat_init(&a,is,ie,N);
    dzmat_init(&b,is,ie,N);
```

```
S-34      }
S-35
S-36      is=N/2; ie=N-1;       //bottom N/2 rows on device 1
S-37      #pragma omp target map(mapper(bottom_id), tofrom: a,b) device(1) \
S-38                          firstprivate(is,ie) nowait
S-39      {
S-40        dzmat_init(&a,is,ie,N);
S-41        dzmat_init(&b,is,ie,N);
S-42      }
S-43
S-44      #pragma omp taskwait
S-45
S-46      host_add(&a,&b,&c,N);
S-47    }
```

C / C++

Fortran

1 *Example target_mapper.2.f90* (`omp_5.0`)

```
S-1     module complex_mats
S-2
S-3       integer, parameter :: N=100    !N must be even
S-4       type dzmat_t
S-5         double precision ::  r_m(N,N), i_m(N,N)
S-6       end type
S-7
S-8       !$omp  declare mapper( left_id: dzmat_t :: v) map( v%r_m(N,  1:N/2), &
S-9       !$omp&                                             v%i_m(N,  1:N/2))
S-10
S-11      !$omp  declare mapper(right_id: dzmat_t :: v) map( v%r_m(N,N/2+1:N), &
S-12      !$omp&                                             v%i_m(N,N/2+1:N))
S-13
S-14    end module
S-15
S-16
S-17    program main
S-18      use  complex_mats
S-19      type(dzmat_t) :: a,b,c
S-20      external dzmat_init, host_add  !initialization and matrix add: a=b+c
S-21
S-22      integer :: is,ie
S-23
S-24
S-25      is=1; ie=N/2            !left N/2 columns on device 0
S-26      !$omp target map(mapper( left_id), tofrom: a,b) device(0) &
S-27      !$omp&       firstprivate(is,ie) nowait
S-28        call dzmat_init(a,is,ie)
```

```
S-29        call dzmat_init(b,is,ie)
S-30      !$omp end target
S-31
S-32      is=N/2+1; ie=N          !right N/2 columns on device 1
S-33      !$omp target map(mapper(right_id), tofrom: a,b) device(1) &
S-34      !$omp&       firstprivate(is,ie) nowait
S-35        call dzmat_init(a,is,ie)
S-36        call dzmat_init(b,is,ie)
S-37      !$omp end target
S-38
S-39      !omp taskwait
S-40
S-41      call host_add(a,b,c)
S-42
S-43    end program main
```

Fortran

In the third example *myvec* structures are nested within a *mypoints* structure. The *myvec_t* type is mapped as in the first example. Following the *mypoints* structure declaration, the *mypoints_t* type is mapped by a **declare mapper** directive. For this structure the *hostonly_data* element will not be mapped; also the array section of *x* (*v.x[:1]*) and *x* will be mapped; and *scratch* will be allocated and used as scratch storage on the device. The default map-type mapping, **tofrom**, applies to the *x* array section, but not to *scratch* which is explicitly mapped with the **alloc** map-type. Note: the variable *v* is not included in the map list (otherwise the *hostonly_data* would be mapped)– just the elements to be mapped are listed.

The two mappers are combined when a *mypoints_t* structure type is mapped, because the mapper *myvec_t* structure type is used within a *mypoints_t* type structure.

C / C++

Example target_mapper.3.c (**omp_5.0**)

```
S-1
S-2     #include <stdlib.h>
S-3     #include  <stdio.h>
S-4
S-5     #define N 100
S-6
S-7     typedef struct myvec {
S-8         size_t len;
S-9         double *data;
S-10    } myvec_t;
S-11
S-12    #pragma omp declare mapper(myvec_t v) \
S-13                        map(v, v.data[0:v.len])
S-14
S-15    typedef struct mypoints {
```

```
    struct myvec scratch;
    struct myvec *x;
    double hostonly_data[500000];
} mypoints_t;

#pragma omp declare mapper(mypoints_t v)  \
                   map(v.x, v.x[0] ) map(alloc:v.scratch)

void init_mypts_array(mypoints_t *P, int n);
void eval_mypts_array(mypoints_t *P, int n);

int main(){

   mypoints_t P;

   init_mypts_array(&P, N);

   #pragma omp target map(P)
   eval_mypts_array(&P, N);

}
```

C / C++

Fortran

Example target_mapper.3.f90 (`omp_5.0`)

```

module my_structures
  type myvec_t
    integer                    :: len
    double precision, pointer :: data(:)
  end type
  !$omp  declare mapper(myvec_t :: v) &
  !$omp&         map(v)

  type mypoints_t
     type(myvec_t)             :: scratch
     type(myvec_t), pointer    :: x(:)
     double precision          :: hostonly_data(500000)
  end  type
 !$omp  declare mapper(mypoints_t :: v)  &
 !$omp&         map(v%x, v%x(1)) map(alloc:v%scratch)

end module


program main
```

```
S-22      use my_structures
S-23      external  init_mypts_array, eval_mypts_array
S-24
S-25      type(mypoints_t) :: P
S-26
S-27       call init_mypts_array(P)
S-28
S-29       !$omp target map(P)
S-30       call eval_mypts_array(P)
S-31
S-32    end program
S-33
```

Fortran

4.8 `target data` Construct

4.8.1 Simple `target data` Construct

This example shows how the **`target data`** construct maps variables to a device data environment. The **`target data`** construct creates a new device data environment and maps the variables *v1*, *v2*, and *p* to the new device data environment. The **`target`** construct enclosed in the **`target data`** region creates a new device data environment, which inherits the variables *v1*, *v2*, and *p* from the enclosing device data environment. The variable *N* is mapped into the new device data environment from the encountering task's data environment.

C / C++

Example target_data.1.c (**`omp_4.0`**)

```
S-1   extern void init(float*, float*, int);
S-2   extern void output(float*, int);
S-3   void vec_mult(float *p, float *v1, float *v2, int N)
S-4   {
S-5      int i;
S-6      init(v1, v2, N);
S-7      #pragma omp target data map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-8      {
S-9         #pragma omp target
S-10        #pragma omp parallel for
S-11        for (i=0; i<N; i++)
S-12          p[i] = v1[i] * v2[i];
S-13     }
S-14     output(p, N);
S-15  }
```

C / C++

The Fortran code passes a reference and specifies the extent of the arrays in the declaration. No length information is necessary in the map clause, as is required with C/C++ pointers.

Fortran

Example target_data.1.f90 (**omp_4.0**)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2      real    ::  p(N), v1(N), v2(N)
S-3      integer ::  i
S-4      call init(v1, v2, N)
S-5      !$omp target data map(to: v1, v2) map(from: p)
S-6      !$omp target
S-7      !$omp parallel do
S-8         do i=1,N
S-9            p(i) = v1(i) * v2(i)
S-10        end do
S-11     !$omp end target
S-12     !$omp end target data
S-13     call output(p, N)
S-14  end subroutine
```

Fortran

4.8.2 `target data` Region Enclosing Multiple `target` Regions

The following examples show how the **target data** construct maps variables to a device data environment of a **target** region. The **target data** construct creates a device data environment and encloses **target** regions, which have their own device data environments. The device data environment of the **target data** region is inherited by the device data environment of an enclosed **target** region. The **target data** construct is used to create variables that will persist throughout the **target data** region.

In the following example the variables *v1* and *v2* are mapped at each **target** construct. Instead of mapping the variable *p* twice, once at each **target** construct, *p* is mapped once by the **target data** construct.

C / C++

Example target_data.2.c (**omp_4.0**)

```
S-1   extern void init(float*, float*, int);
S-2   extern void init_again(float*, float*, int);
S-3   extern void output(float*, int);
S-4   void vec_mult(float *p, float *v1, float *v2, int N)
S-5   {
S-6      int i;
S-7      init(v1, v2, N);
S-8      #pragma omp target data map(from: p[0:N])
S-9      {
S-10        #pragma omp target map(to: v1[:N], v2[:N])
S-11        #pragma omp parallel for
S-12        for (i=0; i<N; i++)
S-13          p[i] = v1[i] * v2[i];
S-14        init_again(v1, v2, N);
S-15        #pragma omp target map(to: v1[:N], v2[:N])
S-16        #pragma omp parallel for
S-17        for (i=0; i<N; i++)
S-18          p[i] = p[i] + (v1[i] * v2[i]);
S-19     }
S-20     output(p, N);
S-21  }
```

C / C++

The Fortran code uses reference and specifies the extent of the *p*, *v1* and *v2* arrays. No length information is necessary in the **map** clause, as is required with C/C++ pointers. The arrays *v1* and *v2* are mapped at each **target** construct. Instead of mapping the array *p* twice, once at each target construct, *p* is mapped once by the **target data** construct.

Fortran

Example target_data.2.f90 (**omp_4.0**)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2      real    ::  p(N), v1(N), v2(N)
S-3      integer ::  i
S-4      call init(v1, v2, N)
S-5      !$omp target data map(from: p)
S-6         !$omp target map(to: v1, v2 )
S-7            !$omp parallel do
S-8            do i=1,N
S-9               p(i) = v1(i) * v2(i)
S-10           end do
S-11        !$omp end target
S-12        call init_again(v1, v2, N)
```

```
S-13          !$omp target map(to: v1, v2 )
S-14             !$omp parallel do
S-15             do i=1,N
S-16                p(i) = p(i) + v1(i) * v2(i)
S-17             end do
S-18          !$omp end target
S-19       !$omp end target data
S-20       call output(p, N)
S-21    end subroutine
```

Fortran

In the following example, the array *Q* is mapped once at the enclosing **target data** region instead of at each **target** construct. In OpenMP 4.0, a scalar variable is implicitly mapped with the **tofrom** map-type. But since OpenMP 4.5, a scalar variable, such as the *tmp* variable, has to be explicitly mapped with the **tofrom** map-type at the first **target** construct in order to return its reduced value from the parallel loop construct to the host. The variable defaults to firstprivate at the second **target** construct.

C / C++

Example target_data.3.c (**omp_4.0**)

```
S-1
S-2     #include <math.h>
S-3     #define COLS 100
S-4
S-5     void gramSchmidt(float Q[][COLS], const int rows)
S-6     {
S-7         int cols = COLS;
S-8         #pragma omp target data map(Q[0:rows][0:cols])
S-9         for(int k=0; k < cols; k++)
S-10        {
S-11            double tmp = 0.0;
S-12
S-13            #pragma omp target map(tofrom: tmp)
S-14            #pragma omp parallel for reduction(+:tmp)
S-15            for(int i=0; i < rows; i++)
S-16                tmp += (Q[i][k] * Q[i][k]);
S-17
S-18            tmp = 1/sqrt(tmp);
S-19
S-20            #pragma omp target
S-21            #pragma omp parallel for
S-22            for(int i=0; i < rows; i++)
S-23                Q[i][k] *= tmp;
S-24        }
S-25    }
S-26
```

```
/* Note:  The variable tmp is now mapped with tofrom, for correct
   execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
 */
```

C / C++

Fortran

Example target_data.3.f90 (**omp_4.0**)

```
subroutine gramSchmidt(Q,rows,cols)
integer             ::   rows,cols,  i,k
double precision    :: Q(rows,cols), tmp
      !$omp target data map(Q)
      do k=1,cols
         tmp = 0.0d0
        !$omp target map(tofrom: tmp)
           !$omp parallel do reduction(+:tmp)
           do i=1,rows
              tmp = tmp + (Q(i,k) * Q(i,k))
           end do
        !$omp end target

          tmp = 1.0d0/sqrt(tmp)

        !$omp target
           !$omp parallel do
           do i=1,rows
               Q(i,k) = Q(i,k)*tmp
           enddo
        !$omp end target
      end do
      !$omp end target data
end subroutine

! Note:  The variable tmp is now mapped with tofrom, for correct
! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.8.3 `target data` Construct with Orphaned Call

The following two examples show how the **`target data`** construct maps variables to a device data environment. The **`target data`** construct's device data environment encloses the **`target`** construct's device data environment in the function **`vec_mult()`**.

When the type of the variable appearing in an array section is pointer, the pointer variable and the storage location of the corresponding array section are mapped to the device data environment. The pointer variable is treated as if it had appeared in a **`map`** clause with a map-type of **`alloc`**. The array section's storage location is mapped according to the map-type in the **`map`** clause (the default map-type is **`tofrom`**).

The **`target`** construct's device data environment inherits the storage locations of the array sections *v1[0:N]*, *v2[:n]*, and *p0[0:N]* from the enclosing target data construct's device data environment. Neither initialization nor assignment is performed for the array sections in the new device data environment.

The pointer variables *p1*, *v3*, and *v4* are mapped into the target construct's device data environment with an implicit map-type of alloc and they are assigned the address of the storage location associated with their corresponding array sections. Note that the following pairs of array section storage locations are equivalent (*p0[:N]*, *p1[:N]*), (*v1[:N]*,*v3[:N]*), and (*v2[:N]*,*v4[:N]*).

C / C++

Example target_data.4.c (**`omp_4.0`**)

```
S-1
S-2
S-3     void vec_mult(float*, float*, float*, int);
S-4
S-5     extern void init(float*, float*, int);
S-6     extern void output(float*, int);
S-7
S-8
S-9     void foo(float *p0, float *v1, float *v2, int N)
S-10    {
S-11       init(v1, v2, N);
S-12
S-13       #pragma omp target data map(to: v1[0:N], v2[:N]) map(from: p0[0:N])
S-14       {
S-15          vec_mult(p0, v1, v2, N);
S-16       }
S-17
S-18       output(p0, N);
S-19    }
S-20
S-21
S-22
S-23    void vec_mult(float *p1, float *v3, float *v4, int N)
```

```
{
   int i;
   #pragma omp target map(to: v3[0:N], v4[:N]) map(from: p1[0:N])
   #pragma omp parallel for
   for (i=0; i<N; i++)
   {
     p1[i] = v3[i] * v4[i];
   }
}
```

C / C++

The Fortran code maps the pointers and storage in an identical manner (same extent, but uses indices from 1 to *N*).

The **target** construct's device data environment inherits the storage locations of the arrays *v1*, *v2* and *p0* from the enclosing **target data** constructs's device data environment. However, in Fortran the associated data of the pointer is known, and the shape is not required.

The pointer variables *p1*, *v3*, and *v4* are mapped into the **target** construct's device data environment with an implicit map-type of **alloc** and they are assigned the address of the storage location associated with their corresponding array sections. Note that the following pair of array storage locations are equivalent (*p0*,*p1*), (*v1*,*v3*), and (*v2*,*v4*).

Fortran

Example target_data.4.f90 (**omp_4.0**)

```
module mults
contains
subroutine foo(p0,v1,v2,N)
real,pointer,dimension(:) :: p0, v1, v2
integer                   :: N,i

   call init(v1, v2, N)

   !$omp target data map(to: v1, v2) map(from: p0)
    call vec_mult(p0,v1,v2,N)
   !$omp end target data

   call output(p0, N)

end subroutine

subroutine vec_mult(p1,v3,v4,N)
real,pointer,dimension(:) :: p1, v3, v4
integer                   :: N,i
```

```
   !$omp target map(to: v3, v4) map(from: p1)
   !$omp parallel do
   do i=1,N
      p1(i) = v3(i) * v4(i)
   end do
   !$omp end target

end subroutine
end module
```

Fortran

In the following example, the variables *p1*, *v3*, and *v4* are references to the pointer variables *p0*, *v1* and *v2* respectively. The **target** construct's device data environment inherits the pointer variables *p0*, *v1*, and *v2* from the enclosing **target data** construct's device data environment. Thus, *p1*, *v3*, and *v4* are already present in the device data environment.

C++

Example target_data.5.cpp (**omp_4.0**)

```
void vec_mult(float* &, float* &, float* &, int &);
extern void init(float*, float*, int);
extern void output(float*, int);
void foo(float *p0, float *v1, float *v2, int N)
{
   init(v1, v2, N);
   #pragma omp target data map(to: v1[0:N], v2[:N]) map(from: p0[0:N])
   {
      vec_mult(p0, v1, v2, N);
   }
   output(p0, N);
}
void vec_mult(float* &p1, float* &v3, float* &v4, int &N)
{
   int i;
   #pragma omp target map(to: v3[0:N], v4[:N]) map(from: p1[0:N])
   #pragma omp parallel for
   for (i=0; i<N; i++)
     p1[i] = v3[i] * v4[i];
}
```

C++

In the following example, the usual Fortran approach is used for dynamic memory. The *p0*, *v1*, and *v2* arrays are allocated in the main program and passed as references from one routine to another. In **vec_mult**, *p1*, *v3* and *v4* are references to the *p0*, *v1*, and *v2* arrays, respectively. The **target** construct's device data environment inherits the arrays *p0*, *v1*, and *v2* from the enclosing target data

construct's device data environment. Thus, *p1*, *v3*, and *v4* are already present in the device data environment.

Fortran

Example target_data.5.f90 (`omp_4.0`)

```
S-1   module my_mult
S-2   contains
S-3   subroutine foo(p0,v1,v2,N)
S-4   real,dimension(:) :: p0, v1, v2
S-5   integer           :: N,i
S-6      call init(v1, v2, N)
S-7      !$omp target data map(to: v1, v2) map(from: p0)
S-8       call vec_mult(p0,v1,v2,N)
S-9      !$omp end target data
S-10     call output(p0, N)
S-11  end subroutine
S-12  subroutine vec_mult(p1,v3,v4,N)
S-13  real,dimension(:) :: p1, v3, v4
S-14  integer           :: N,i
S-15     !$omp target map(to: v3, v4) map(from: p1)
S-16     !$omp parallel do
S-17     do i=1,N
S-18        p1(i) = v3(i) * v4(i)
S-19     end do
S-20     !$omp end target
S-21  end subroutine
S-22  end module
S-23  program main
S-24  use my_mult
S-25  integer, parameter :: N=1024
S-26  real,allocatable, dimension(:) :: p, v1, v2
S-27     allocate( p(N), v1(N), v2(N) )
S-28     call foo(p,v1,v2,N)
S-29  end program
```

Fortran

4.8.4 `target data` Construct with `if` Clause

The following two examples show how the **`target data`** construct maps variables to a device data environment.

In the following example, the if clause on the **`target data`** construct indicates that if the variable *N* is smaller than a given threshold, then the **`target data`** construct will not create a device data environment.

The **`target`** constructs enclosed in the **`target data`** region must also use an **`if`** clause on the same condition, otherwise the pointer variable *p* is implicitly mapped with a map-type of **`tofrom`**, but the storage location for the array section *p[0:N]* will not be mapped in the device data environments of the **`target`** constructs.

C / C++

Example target_data.6.c (**`omp_4.0`**)

```
S-1   #define THRESHOLD 1000000
S-2   extern void init(float*, float*, int);
S-3   extern void init_again(float*, float*, int);
S-4   extern void output(float*, int);
S-5   void vec_mult(float *p, float *v1, float *v2, int N)
S-6   {
S-7      int i;
S-8      init(v1, v2, N);
S-9      #pragma omp target data if(N>THRESHOLD) map(from: p[0:N])
S-10     {
S-11        #pragma omp target if (N>THRESHOLD) map(to: v1[:N], v2[:N])
S-12        #pragma omp parallel for
S-13        for (i=0; i<N; i++)
S-14          p[i] = v1[i] * v2[i];
S-15        init_again(v1, v2, N);
S-16        #pragma omp target if (N>THRESHOLD) map(to: v1[:N], v2[:N])
S-17        #pragma omp parallel for
S-18        for (i=0; i<N; i++)
S-19          p[i] = p[i] + (v1[i] * v2[i]);
S-20     }
S-21     output(p, N);
S-22  }
```

C / C++

The **if** clauses work the same way for the following Fortran code. The **target** constructs enclosed in the **target data** region should also use an **if** clause with the same condition, so that the **target data** region and the **target** region are either both created for the device, or are both ignored.

Fortran

Example target_data.6.f90 (**omp_4.0**)

```
S-1   module params
S-2   integer,parameter :: THRESHOLD=1000000
S-3   end module
S-4   subroutine vec_mult(p, v1, v2, N)
S-5      use params
S-6      real    ::  p(N), v1(N), v2(N)
S-7      integer ::  i
S-8      call init(v1, v2, N)
S-9      !$omp target data if(N>THRESHOLD) map(from: p)
S-10        !$omp target if(N>THRESHOLD) map(to: v1, v2)
S-11           !$omp parallel do
S-12           do i=1,N
S-13              p(i) = v1(i) * v2(i)
S-14           end do
S-15        !$omp end target
S-16        call init_again(v1, v2, N)
S-17        !$omp target if(N>THRESHOLD) map(to: v1, v2)
S-18           !$omp parallel do
S-19           do i=1,N
S-20              p(i) = p(i) + v1(i) * v2(i)
S-21           end do
S-22        !$omp end target
S-23     !$omp end target data
S-24     call output(p, N)
S-25  end subroutine
```

Fortran

In the following example, when the **if** clause conditional expression on the **target** construct evaluates to *false*, the target region will execute on the host device. However, the **target data** construct created an enclosing device data environment that mapped *p[0:N]* to a device data environment on the default device. At the end of the **target data** region the array section *p[0:N]* will be assigned from the device data environment to the corresponding variable in the data environment of the task that encountered the **target data** construct, resulting in undefined values in *p[0:N]*.

C / C++

Example target_data.7.c (**omp_4.0**)

```
#define THRESHOLD 1000000
extern void init(float*, float*, int);
extern void output(float*, int);
void vec_mult(float *p, float *v1, float *v2, int N)
{
   int i;
   init(v1, v2, N);
   #pragma omp target data map(from: p[0:N])
   {
      #pragma omp target if (N>THRESHOLD) map(to: v1[:N], v2[:N])
      #pragma omp parallel for
      for (i=0; i<N; i++)
        p[i] = v1[i] * v2[i];
   } /* UNDEFINED behavior if N<=THRESHOLD */
   output(p, N);
}
```

C / C++

The **if** clauses work the same way for the following Fortran code. When the **if** clause conditional expression on the **target** construct evaluates to *false*, the **target** region will execute on the host device. However, the **target data** construct created an enclosing device data environment that mapped the *p* array (and *v1* and *v2*) to a device data environment on the default target device. At the end of the **target data** region the *p* array will be assigned from the device data environment to the corresponding variable in the data environment of the task that encountered the **target data** construct, resulting in undefined values in *p*.

Fortran

Example target_data.7.f90 (**omp_4.0**)

```
S-1   module params
S-2   integer, parameter :: THRESHOLD=1000000
S-3   end module
S-4   subroutine vec_mult(p, v1, v2, N)
S-5      use params
S-6      real    ::  p(N), v1(N), v2(N)
S-7      integer ::  i
S-8      call init(v1, v2, N)
S-9      !$omp target data map(from: p)
S-10        !$omp target if(N>THRESHOLD) map(to: v1, v2)
S-11           !$omp parallel do
S-12           do i=1,N
S-13              p(i) = v1(i) * v2(i)
S-14           end do
S-15        !$omp end target
S-16     !$omp end target data
S-17     call output(p, N)  !*** UNDEFINED behavior if N<=THRESHOLD
S-18  end subroutine
```

Fortran

4.9 `target enter data` and `target exit data` Constructs

The structured data construct (**`target data`**) provides persistent data on a device for subsequent **`target`** constructs as shown in the **`target data`** examples above. This is accomplished by creating a single **`target data`** region containing **`target`** constructs.

The unstructured data constructs allow the creation and deletion of data on the device at any appropriate point within the host code, as shown below with the **`target enter data`** and **`target exit data`** constructs.

The following C++ code creates/deletes a vector in a constructor/destructor of a class. The constructor creates a vector with **`target enter data`** and uses an **`alloc`** modifier in the **`map`** clause to avoid copying values to the device. The destructor deletes the data (**`target exit data`**) and uses the **`delete`** modifier in the **`map`** clause to avoid copying data back to the host. Note, the stand-alone **`target enter data`** occurs after the host vector is created, and the **`target exit data`** construct occurs before the host data is deleted.

C++

Example target_unstructured_data.1.cpp (**`omp_4.5`**)

```
S-1   class Matrix
S-2   {
S-3
S-4     Matrix(int n) {
S-5       len = n;
S-6       v = new double[len];
S-7       #pragma omp target enter data map(alloc:v[0:len])
S-8     }
S-9
S-10    ~Matrix() {
S-11      // NOTE: delete map type should be used, since the corresponding
S-12      // host data will cease to exist after the deconstructor is called.
S-13
S-14      #pragma omp target exit data map(delete:v[0:len])
S-15      delete[] v;
S-16    }
S-17
S-18    private:
S-19    double* v;
S-20    int len;
S-21
S-22  };
```

C++

The following C code allocates and frees the data member of a Matrix structure. The **init_matrix** function allocates the memory used in the structure and uses the **target enter data** directive to map it to the target device. The **free_matrix** function removes the mapped array from the target device and then frees the memory on the host. Note, the stand-alone **target enter data** occurs after the host memory is allocated, and the **target exit data** construct occurs before the host data is freed.

C / C++

Example target_unstructured_data.1.c (**omp_4.5**)

```
#include <stdlib.h>
typedef struct {
  double *A;
  int N;
} Matrix;

void init_matrix(Matrix *mat, int n)
{
  mat->A = (double *)malloc(n*sizeof(double));
  mat->N = n;
  #pragma omp target enter data map(alloc:mat->A[:n])
}

void free_matrix(Matrix *mat)
{
  #pragma omp target exit data map(delete:mat->A[:mat->N])
  mat->N = 0;
  free(mat->A);
  mat->A = NULL;
}
```

C / C++

The following Fortran code allocates and deallocates a module array. The **initialize** subroutine allocates the module array and uses the **target enter data** directive to map it to the target device. The **finalize** subroutine removes the mapped array from the target device and then deallocates the array on the host. Note, the stand-alone **target enter data** occurs after the host memory is allocated, and the **target exit data** construct occurs before the host data is deallocated.

Fortran

Example target_unstructured_data.1.f90 (**omp_4.5**)

```
S-1   module example
S-2     real(8), allocatable :: A(:)
S-3
S-4     contains
S-5       subroutine initialize(N)
S-6         integer :: N
S-7
S-8         allocate(A(N))
S-9         !$omp target enter data map(alloc:A)
S-10
S-11      end subroutine initialize
S-12
S-13      subroutine finalize()
S-14
S-15        !$omp target exit data map(delete:A)
S-16        deallocate(A)
S-17
S-18      end subroutine finalize
S-19  end module example
```

Fortran

4.10 `target update` Construct

4.10.1 Simple `target data` and `target update` Constructs

The following example shows how the **`target update`** construct updates variables in a device data environment.

The **`target data`** construct maps array sections *v1[:N]* and *v2[:N]* (arrays *v1* and *v2* in the Fortran code) into a device data environment.

The task executing on the host device encounters the first **`target`** region and waits for the completion of the region.

After the execution of the first **`target`** region, the task executing on the host device then assigns new values to *v1[:N]* and *v2[:N]* (*v1* and *v2* arrays in Fortran code) in the task's data environment by calling the function **`init_again()`**.

The **`target update`** construct assigns the new values of *v1* and *v2* from the task's data environment to the corresponding mapped array sections in the device data environment of the **`target data`** construct.

The task executing on the host device then encounters the second **`target`** region and waits for the completion of the region.

The second **`target`** region uses the updated values of *v1[:N]* and *v2[:N]*.

C / C++

Example target_update.1.c (**`omp_4.0`**)

```
*/
extern void init(float *, float *, int);
extern void init_again(float *, float *, int);
extern void output(float *, int);
void vec_mult(float *p, float *v1, float *v2, int N)
{
   int i;
   init(v1, v2, N);
   #pragma omp target data map(to: v1[:N], v2[:N]) map(from: p[0:N])
   {
      #pragma omp target
      #pragma omp parallel for
      for (i=0; i<N; i++)
        p[i] = v1[i] * v2[i];
      init_again(v1, v2, N);
      #pragma omp target update to(v1[:N], v2[:N])
      #pragma omp target
      #pragma omp parallel for
      for (i=0; i<N; i++)
        p[i] = p[i] + (v1[i] * v2[i]);
```

```
S-21        }
S-22        output(p, N);
S-23     }
```

C / C++

Fortran

Example target_update.1.f90 (`omp_4.0`)

```
S-1      subroutine vec_mult(p, v1, v2, N)
S-2         real    ::  p(N), v1(N), v2(N)
S-3         integer ::  i
S-4         call init(v1, v2, N)
S-5         !$omp target data map(to: v1, v2) map(from: p)
S-6            !$omp target
S-7            !$omp parallel do
S-8               do i=1,N
S-9                  p(i) = v1(i) * v2(i)
S-10              end do
S-11           !$omp end target
S-12           call init_again(v1, v2, N)
S-13           !$omp target update to(v1, v2)
S-14           !$omp target
S-15           !$omp parallel do
S-16              do i=1,N
S-17                 p(i) = p(i) + v1(i) * v2(i)
S-18              end do
S-19           !$omp end target
S-20        !$omp end target data
S-21        call output(p, N)
S-22     end subroutine
```

Fortran

4.10.2 `target update` Construct with `if` Clause

The following example shows how the **`target update`** construct updates variables in a device data environment.

The **`target data`** construct maps array sections *v1[:N]* and *v2[:N]* (arrays *v1* and *v2* in the Fortran code) into a device data environment. In between the two **`target`** regions, the task executing on the host device conditionally assigns new values to *v1* and *v2* in the task's data environment. The function **`maybe_init_again()`** returns *true* if new data is written.

When the conditional expression (the return value of **`maybe_init_again()`**) in the **`if`** clause is *true*, the **`target update`** construct assigns the new values of *v1* and *v2* from the task's data environment to the corresponding mapped array sections in the **`target data`** construct's device data environment.

C / C++

Example target_update.2.c (**`omp_4.0`**)

```
S-1   extern void init(float *, float *, int);
S-2   extern int maybe_init_again(float *, int);
S-3   extern void output(float *, int);
S-4   void vec_mult(float *p, float *v1, float *v2, int N)
S-5   {
S-6      int i;
S-7      init(v1, v2, N);
S-8      #pragma omp target data map(to: v1[:N], v2[:N]) map(from: p[0:N])
S-9      {
S-10        int changed;
S-11        #pragma omp target
S-12        #pragma omp parallel for
S-13        for (i=0; i<N; i++)
S-14          p[i] = v1[i] * v2[i];
S-15        changed = maybe_init_again(v1,  N);
S-16        #pragma omp target update if (changed) to(v1[:N])
S-17        changed = maybe_init_again(v2,  N);
S-18        #pragma omp target update if (changed) to(v2[:N])
S-19        #pragma omp target
S-20        #pragma omp parallel for
S-21        for (i=0; i<N; i++)
S-22          p[i] = p[i] + (v1[i] * v2[i]);
S-23     }
S-24     output(p, N);
S-25  }
```

C / C++

Fortran

Example target_update.2.f90 (`omp_4.0`)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2      interface
S-3         logical function maybe_init_again (v1, N)
S-4         real :: v1(N)
S-5         integer :: N
S-6         end function
S-7      end interface
S-8      real    ::  p(N), v1(N), v2(N)
S-9      integer ::  i
S-10     logical :: changed
S-11     call init(v1, v2, N)
S-12     !$omp target data map(to: v1, v2) map(from: p)
S-13        !$omp target
S-14           !$omp parallel do
S-15           do i=1, N
S-16              p(i) = v1(i) * v2(i)
S-17           end do
S-18        !$omp end target
S-19        changed = maybe_init_again(v1, N)
S-20        !$omp target update if(changed) to(v1(:N))
S-21        changed = maybe_init_again(v2, N)
S-22        !$omp target update if(changed) to(v2(:N))
S-23        !$omp target
S-24           !$omp parallel do
S-25           do i=1, N
S-26              p(i) = p(i) + v1(i) * v2(i)
S-27           end do
S-28        !$omp end target
S-29     !$omp end target data
S-30     call output(p, N)
S-31  end subroutine
```

Fortran

4.11 `declare target` Construct

4.11.1 `declare target` and `end declare target` for a Function

The following example shows how the **`declare target`** directive is used to indicate that the corresponding call inside a **`target`** region is to a **`fib`** function that can execute on the default target device.

A version of the function is also available on the host device. When the **`if`** clause conditional expression on the **`target`** construct evaluates to *false*, the **`target`** region (thus **`fib`**) will execute on the host device.

For C/C++ codes the declaration of the function **`fib`** appears between the **`declare target`** and **`end declare target`** directives.

C / C++

Example declare_target.1.c (**`omp_4.0`**)

```
S-1   #pragma omp declare target
S-2   extern void fib(int N);
S-3   #pragma omp end declare target
S-4   #define THRESHOLD 1000000
S-5   void fib_wrapper(int n)
S-6   {
S-7      #pragma omp target if(n > THRESHOLD)
S-8      {
S-9         fib(n);
S-10     }
S-11  }
```

C / C++

The Fortran **`fib`** subroutine contains a **`declare target`** declaration to indicate to the compiler to create an device executable version of the procedure. The subroutine name has not been included on the **`declare target`** directive and is, therefore, implicitly assumed.

The program uses the **`module_fib`** module, which presents an explicit interface to the compiler with the **`declare target`** declarations for processing the **`fib`** call.

Fortran

Example declare_target.1.f90 (**omp_4.0**)

```
S-1   module module_fib
S-2   contains
S-3      subroutine fib(N)
S-4         integer :: N
S-5         !$omp declare target
S-6         !...
S-7      end subroutine
S-8   end module
S-9   module params
S-10  integer :: THRESHOLD=1000000
S-11  end module
S-12  program my_fib
S-13  use params
S-14  use module_fib
S-15     !$omp target if( N > THRESHOLD )
S-16        call fib(N)
S-17     !$omp end target
S-18  end program
```

Fortran

The next Fortran example shows the use of an external subroutine. Without an explicit interface (through module use or an interface block) the **declare target** declarations within a external subroutine are unknown to the main program unit; therefore, a **declare target** must be provided within the program scope for the compiler to determine that a target binary should be available.

Fortran

Example declare_target.2.f90 (**omp_4.0**)

```
S-1   program my_fib
S-2   integer :: N = 8
S-3   !$omp declare target(fib)
S-4      !$omp target
S-5         call fib(N)
S-6      !$omp end target
S-7   end program
S-8   subroutine fib(N)
S-9   integer :: N
S-10  !$omp declare target
S-11       print*,"hello from fib"
S-12       !...
S-13  end subroutine
```

Fortran

4.11.2 `declare target` Construct for Class Type

C++

The following example shows how the **`declare target`** and **`end declare target`** directives are used to enclose the declaration of a variable *varY* with a class type **`typeY`**. The member function **`typeY::foo()`** cannot be accessed on a target device because its declaration did not appear between **`declare target`** and **`end declare target`** directives.

Example declare_target.2.cpp (**`omp_4.0`**)

```
S-1   struct typeX
S-2   {
S-3      int a;
S-4   };
S-5   class typeY
S-6   {
S-7      int a;
S-8     public:
S-9      int foo() { return a^0x01;}
S-10  };
S-11  #pragma omp declare target
S-12  struct typeX varX;  // ok
S-13  class typeY varY; // ok if varY.foo() not called on target device
S-14  #pragma omp end declare target
S-15  void foo()
S-16  {
S-17     #pragma omp target
S-18     {
S-19        varX.a = 100; // ok
S-20        varY.foo(); // error foo() is not available on a target device
S-21     }
S-22  }
```

C++

4.11.3 `declare target` and `end declare target` for Variables

The following examples show how the **`declare target`** and **`end declare target`** directives are used to indicate that global variables are mapped to the implicit device data environment of each target device.

In the following example, the declarations of the variables *p*, *v1*, and *v2* appear between **`declare target`** and **`end declare target`** directives indicating that the variables are mapped to the implicit device data environment of each target device. The **`target update`** directive is then

used to manage the consistency of the variables *p*, *v1*, and *v2* between the data environment of the encountering host device task and the implicit device data environment of the default target device.

C / C++

Example declare_target.3.c (**omp_4.0**)

```
S-1   #define N 1000
S-2   #pragma omp declare target
S-3   float p[N], v1[N], v2[N];
S-4   #pragma omp end declare target
S-5   extern void init(float *, float *, int);
S-6   extern void output(float *, int);
S-7   void vec_mult()
S-8   {
S-9      int i;
S-10     init(v1, v2, N);
S-11     #pragma omp target update to(v1, v2)
S-12     #pragma omp target
S-13     #pragma omp parallel for
S-14     for (i=0; i<N; i++)
S-15       p[i] = v1[i] * v2[i];
S-16     #pragma omp target update from(p)
S-17     output(p, N);
S-18  }
```

C / C++

The Fortran version of the above C code uses a different syntax. Fortran modules use a list syntax on the **declare target** directive to declare mapped variables.

Fortran

Example declare_target.3.f90 (**omp_4.0**)

```
S-1   module my_arrays
S-2   !$omp declare target (N, p, v1, v2)
S-3   integer, parameter :: N=1000
S-4   real               :: p(N), v1(N), v2(N)
S-5   end module
S-6   subroutine vec_mult()
S-7   use my_arrays
S-8      integer :: i
S-9      call init(v1, v2, N);
S-10     !$omp target update to(v1, v2)
S-11     !$omp target
S-12     !$omp parallel do
S-13     do i = 1,N
S-14       p(i) = v1(i) * v2(i)
S-15     end do
```

```
!$omp end target
!$omp target update from (p)
call output(p, N)
end subroutine
```

Fortran

The following example also indicates that the function **Pfun()** is available on the target device, as well as the variable *Q*, which is mapped to the implicit device data environment of each target device. The **target update** directive is then used to manage the consistency of the variable *Q* between the data environment of the encountering host device task and the implicit device data environment of the default target device.

In the following example, the function and variable declarations appear between the **declare target** and **end declare target** directives.

C / C++

Example declare_target.4.c (**omp_4.0**)

```
#define N 10000
#pragma omp declare target
float Q[N][N];
float Pfun(const int i, const int k)
{ return Q[i][k] * Q[k][i]; }
#pragma omp end declare target
float accum(int k)
{
    float tmp = 0.0;
    #pragma omp target update to(Q)
    #pragma omp target map(tofrom: tmp)
    #pragma omp parallel for reduction(+:tmp)
    for(int i=0; i < N; i++)
        tmp += Pfun(i,k);
    return tmp;
}

/* Note:  The variable tmp is now mapped with tofrom, for correct
   execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
 */
```

C / C++

The Fortran version of the above C code uses a different syntax. In Fortran modules a list syntax on the **declare target** directive is used to declare mapped variables and procedures. The *N* and *Q* variables are declared as a comma separated list. When the **declare target** directive is used to declare just the procedure, the procedure name need not be listed – it is implicitly assumed, as illustrated in the **Pfun()** function.

Fortran

Example declare_target.4.f90 (`omp_4.0`)

```
S-1   module my_global_array
S-2   !$omp declare target (N,Q)
S-3   integer, parameter :: N=10
S-4   real               :: Q(N,N)
S-5   contains
S-6   function Pfun(i,k)
S-7   !$omp declare target
S-8   real               :: Pfun
S-9   integer,intent(in) :: i,k
S-10     Pfun=(Q(i,k) * Q(k,i))
S-11  end function
S-12  end module
S-13
S-14  function accum(k) result(tmp)
S-15  use my_global_array
S-16  real    :: tmp
S-17  integer :: i, k
S-18     tmp = 0.0e0
S-19     !$omp target map(tofrom: tmp)
S-20     !$omp parallel do reduction(+:tmp)
S-21     do i=1,N
S-22        tmp = tmp + Pfun(k,i)
S-23     end do
S-24     !$omp end target
S-25  end function
S-26
S-27  ! Note:  The variable tmp is now mapped with tofrom, for correct
S-28  ! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.11.4 `declare target` and `end declare target` with `declare simd`

The following example shows how the **`declare target`** and **`end declare target`** directives are used to indicate that a function is available on a target device. The **`declare simd`** directive indicates that there is a SIMD version of the function **`P()`** that is available on the target device as well as one that is available on the host device.

C / C++

Example declare_target.5.c (**omp_4.0**)

```
S-1   #define N 10000
S-2   #define M 1024
S-3   #pragma omp declare target
S-4   float Q[N][N];
S-5   #pragma omp declare simd uniform(i) linear(k) notinbranch
S-6   float P(const int i, const int k)
S-7   {
S-8     return Q[i][k] * Q[k][i];
S-9   }
S-10  #pragma omp end declare target
S-11
S-12  float accum(void)
S-13  {
S-14    float tmp = 0.0;
S-15    int i, k;
S-16  #pragma omp target map(tofrom: tmp)
S-17  #pragma omp parallel for reduction(+:tmp)
S-18    for (i=0; i < N; i++) {
S-19      float tmp1 = 0.0;
S-20  #pragma omp simd reduction(+:tmp1)
S-21      for (k=0; k < M; k++) {
S-22        tmp1 += P(i,k);
S-23      }
S-24      tmp += tmp1;
S-25    }
S-26    return tmp;
S-27  }
S-28
S-29  /* Note:  The variable tmp is now mapped with tofrom, for correct
S-30     execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
S-31   */
```

C / C++

The Fortran version of the above C code uses a different syntax. Fortran modules use a list syntax of the **declare target** declaration for the mapping. Here the *N* and *Q* variables are declared in the list form as a comma separated list. The function declaration does not use a list and implicitly assumes the function name. In this Fortran example row and column indices are reversed relative to the C/C++ example, as is usual for codes optimized for memory access.

Fortran

Example declare_target.5.f90 (`omp_4.0`)

```
S-1   module my_global_array
S-2   !$omp declare target (N,Q)
S-3   integer, parameter :: N=10000, M=1024
S-4   real               :: Q(N,N)
S-5   contains
S-6   function P(k,i)
S-7   !$omp declare simd uniform(i) linear(k) notinbranch
S-8   !$omp declare target
S-9   real               :: P
S-10  integer,intent(in) :: k,i
S-11     P=(Q(k,i) * Q(i,k))
S-12  end function
S-13  end module
S-14
S-15  function accum() result(tmp)
S-16  use my_global_array
S-17  real    :: tmp, tmp1
S-18  integer :: i
S-19     tmp = 0.0e0
S-20     !$omp target map(tofrom: tmp)
S-21     !$omp parallel do private(tmp1) reduction(+:tmp)
S-22     do i=1,N
S-23        tmp1 = 0.0e0
S-24        !$omp simd reduction(+:tmp1)
S-25        do k = 1,M
S-26           tmp1 = tmp1 + P(k,i)
S-27        end do
S-28        tmp = tmp + tmp1
S-29     end do
S-30     !$omp end target
S-31  end function
S-32
S-33  ! Note:  The variable tmp is now mapped with tofrom, for correct
S-34  ! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.11.5 `declare target` Directive with `link` Clause

In the OpenMP 4.5 standard the **`declare target`** directive was extended to allow static data to be mapped, *when needed*, through a **`link`** clause.

Data storage for items listed in the **`link`** clause becomes available on the device when it is mapped implicitly or explicitly in a **`map`** clause, and it persists for the scope of the mapping (as specified by a **`target`** construct, a **`target data`** construct, or **`target enter/exit data`** constructs).

Tip: When all the global data items will not fit on a device and are not needed simultaneously, use the **`link`** clause and map the data only when it is needed.

The following C and Fortran examples show two sets of data (single precision and double precision) that are global on the host for the entire execution on the host; but are only used globally on the device for part of the program execution. The single precision data are allocated and persist only for the first **`target`** region. Similarly, the double precision data are in scope on the device only for the second **`target`** region.

C / C++

Example declare_target.6.c (**`omp_4.5`**)

```
S-1   #define N 100000000
S-2
S-3   float  sp[N], sv1[N], sv2[N];
S-4   double dp[N], dv1[N], dv2[N];
S-5   #pragma omp declare target link(sp,sv1,sv2) \
S-6                              link(dp,dv1,dv2)
S-7
S-8   void s_init(float *, float *, int);
S-9   void d_init(double *, double *, int);
S-10  void s_output(float *, int);
S-11  void d_output(double *, int);
S-12
S-13  #pragma omp declare target
S-14  void s_vec_mult_accum()
S-15  {
S-16     int i;
S-17
S-18     #pragma omp parallel for
S-19     for (i=0; i<N; i++)
S-20       sp[i] = sv1[i] * sv2[i];
S-21  }
S-22
S-23  void d_vec_mult_accum()
S-24  {
S-25     int i;
S-26
S-27     #pragma omp parallel for
```

```
S-28        for (i=0; i<N; i++)
S-29          dp[i] = dv1[i] * dv2[i];
S-30    }
S-31    #pragma omp end declare target
S-32
S-33    int main()
S-34    {
S-35       s_init(sv1, sv2, N);
S-36       #pragma omp target map(to:sv1,sv2) map(from:sp)
S-37          s_vec_mult_accum();
S-38       s_output(sp, N);
S-39
S-40       d_init(dv1, dv2, N);
S-41       #pragma omp target map(to:dv1,dv2) map(from:dp)
S-42          d_vec_mult_accum();
S-43       d_output(dp, N);
S-44
S-45      return 0;
S-46    }
```

C / C++

Fortran

Example declare_target.6.f90 (`omp_4.5`)

```
S-1     module m_dat
S-2        integer, parameter :: N=100000000
S-3        !$omp declare target link(sp,sv1,sv2)
S-4        real :: sp(N), sv1(N), sv2(N)
S-5
S-6        !$omp declare target link(dp,dv1,dv2)
S-7        double precision :: dp(N), dv1(N), dv2(N)
S-8
S-9     contains
S-10       subroutine s_vec_mult_accum()
S-11       !$omp declare target
S-12          integer :: i
S-13
S-14          !$omp parallel do
S-15          do i = 1,N
S-16            sp(i) = sv1(i) * sv2(i)
S-17          end do
S-18
S-19       end subroutine s_vec_mult_accum
S-20
S-21       subroutine d_vec_mult_accum()
S-22       !$omp declare target
S-23          integer :: i
```

```
S-24
S-25          !$omp parallel do
S-26          do i = 1,N
S-27            dp(i) = dv1(i) * dv2(i)
S-28          end do
S-29
S-30       end subroutine
S-31    end module m_dat
S-32
S-33    program prec_vec_mult
S-34       use m_dat
S-35
S-36       call s_init(sv1, sv2, N)
S-37       !$omp target map(to:sv1,sv2) map(from:sp)
S-38         call s_vec_mult_accum()
S-39       !$omp end target
S-40       call s_output(sp, N)
S-41
S-42       call d_init(dv1, dv2, N)
S-43       !$omp target map(to:dv1,dv2) map(from:dp)
S-44         call d_vec_mult_accum()
S-45       !$omp end target
S-46       call d_output(dp, N)
S-47
S-48    end program
```

Fortran

4.12 `teams` Constructs

4.12.1 `target` and `teams` Constructs with `omp_get_num_teams` and `omp_get_team_num` Routines

The following example shows how the **`target`** and **`teams`** constructs are used to create a league of thread teams that execute a region. The **`teams`** construct creates a league of at most two teams where the master thread of each team executes the **`teams`** region.

The **`omp_get_num_teams`** routine returns the number of teams executing in a **`teams`** region. The **`omp_get_team_num`** routine returns the team number, which is an integer between 0 and one less than the value returned by **`omp_get_num_teams`**. The following example manually distributes a loop across two teams.

C / C++

Example teams.1.c (**`omp_4.0`**)

```
S-1   #include <stdlib.h>
S-2   #include <omp.h>
S-3   float dotprod(float B[], float C[], int N)
S-4   {
S-5      float sum0 = 0.0;
S-6      float sum1 = 0.0;
S-7      #pragma omp target map(to: B[:N], C[:N]) map(tofrom: sum0, sum1)
S-8      #pragma omp teams num_teams(2)
S-9      {
S-10        int i;
S-11        if (omp_get_num_teams() != 2)
S-12           abort();
S-13        if (omp_get_team_num() == 0)
S-14        {
S-15           #pragma omp parallel for reduction(+:sum0)
S-16           for (i=0; i<N/2; i++)
S-17              sum0 += B[i] * C[i];
S-18        }
S-19        else if (omp_get_team_num() == 1)
S-20        {
S-21           #pragma omp parallel for reduction(+:sum1)
S-22           for (i=N/2; i<N; i++)
S-23              sum1 += B[i] * C[i];
S-24        }
S-25     }
S-26     return sum0 + sum1;
S-27  }
S-28
S-29  /* Note:  The variables sum0,sum1 are now mapped with tofrom, for correct
```

```
   execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
 */
```

C / C++

Fortran

Example teams.1.f90 (`omp_4.0`)

```
function dotprod(B,C,N) result(sum)
use omp_lib, ONLY : omp_get_num_teams, omp_get_team_num
    real    :: B(N), C(N), sum,sum0, sum1
    integer :: N, i
    sum0 = 0.0e0
    sum1 = 0.0e0
    !$omp target map(to: B, C) map(tofrom: sum0, sum1)
    !$omp teams num_teams(2)
      if (omp_get_num_teams() /= 2) stop "2 teams required"
      if (omp_get_team_num() == 0) then
         !$omp parallel do reduction(+:sum0)
         do i=1,N/2
            sum0 = sum0 + B(i) * C(i)
         end do
      else if (omp_get_team_num() == 1) then
         !$omp parallel do reduction(+:sum1)
         do i=N/2+1,N
            sum1 = sum1 + B(i) * C(i)
         end do
      end if
    !$omp end teams
    !$omp end target
    sum = sum0 + sum1
end function

! Note:  The variables sum0,sum1 are now mapped with tofrom, for correct
! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.12.2 `target`, `teams`, and `distribute` Constructs

The following example shows how the **`target`**, **`teams`**, and **`distribute`** constructs are used to execute a loop nest in a **`target`** region. The **`teams`** construct creates a league and the master thread of each team executes the **`teams`** region. The **`distribute`** construct schedules the subsequent loop iterations across the master threads of each team.

The number of teams in the league is less than or equal to the variable *num_blocks*. Each team in the league has a number of threads less than or equal to the variable *block_threads*. The iterations in the outer loop are distributed among the master threads of each team.

When a team's master thread encounters the parallel loop construct before the inner loop, the other threads in its team are activated. The team executes the **`parallel`** region and then workshares the execution of the loop.

Each master thread executing the **`teams`** region has a private copy of the variable *sum* that is created by the **`reduction`** clause on the **`teams`** construct. The master thread and all threads in its team have a private copy of the variable *sum* that is created by the **`reduction`** clause on the parallel loop construct. The second private *sum* is reduced into the master thread's private copy of *sum* created by the **`teams`** construct. At the end of the **`teams`** region, each master thread's private copy of *sum* is reduced into the final *sum* that is implicitly mapped into the **`target`** region.

C / C++

Example teams.2.c (**`omp_4.0`**)

```
S-1   #define min(x, y) (((x) < (y)) ? (x) : (y))
S-2
S-3   float dotprod(float B[], float C[], int N, int block_size,
S-4     int num_teams, int block_threads)
S-5   {
S-6       float sum = 0.0;
S-7       int i, i0;
S-8       #pragma omp target map(to: B[0:N], C[0:N]) map(tofrom: sum)
S-9       #pragma omp teams num_teams(num_teams) thread_limit(block_threads) \
S-10        reduction(+:sum)
S-11      #pragma omp distribute
S-12      for (i0=0; i0<N; i0 += block_size)
S-13         #pragma omp parallel for reduction(+:sum)
S-14         for (i=i0; i< min(i0+block_size,N); i++)
S-15             sum += B[i] * C[i];
S-16      return sum;
S-17  }
S-18
S-19  /* Note:  The variable sum is now mapped with tofrom, for correct
S-20     execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
S-21   */
```

C / C++

Fortran

Example teams.2.f90 (**omp_4.0**)

```
function dotprod(B,C,N, block_size, num_teams, block_threads) result(sum)
implicit none
    real    :: B(N), C(N), sum
    integer :: N, block_size, num_teams, block_threads, i, i0
    sum = 0.0e0
    !$omp target map(to: B, C) map(tofrom: sum)
    !$omp teams num_teams(num_teams) thread_limit(block_threads) &
    !$omp&  reduction(+:sum)
    !$omp distribute
       do i0=1,N, block_size
          !$omp parallel do reduction(+:sum)
          do i = i0, min(i0+block_size,N)
             sum = sum + B(i) * C(i)
          end do
       end do
    !$omp end teams
    !$omp end target
end function

! Note:  The variable sum is now mapped with tofrom, for correct
! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.12.3 `target teams`, and Distribute Parallel Loop Constructs

The following example shows how the **target teams** and distribute parallel loop constructs are used to execute a **target** region. The **target teams** construct creates a league of teams where the master thread of each team executes the **teams** region.

The distribute parallel loop construct schedules the loop iterations across the master threads of each team and then across the threads of each team.

C / C++

1 *Example teams.3.c* (**omp_4.5**)

```
S-1   float dotprod(float B[], float C[], int N)
S-2   {
S-3      float sum = 0;
S-4      int i;
S-5      #pragma omp target teams map(to: B[0:N], C[0:N]) \
S-6                               defaultmap(tofrom:scalar) reduction(+:sum)
S-7      #pragma omp distribute parallel for reduction(+:sum)
S-8      for (i=0; i<N; i++)
S-9         sum += B[i] * C[i];
S-10     return sum;
S-11  }
S-12
S-13  /* Note:  The variable sum is now mapped with tofrom from the defaultmap
S-14     clause on the combined target teams construct, for correct
S-15     execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
S-16   */
```

C / C++

Fortran

2 *Example teams.3.f90* (**omp_4.5**)

```
S-1   function dotprod(B,C,N) result(sum)
S-2      real    :: B(N), C(N), sum
S-3      integer :: N, i
S-4      sum = 0.0e0
S-5      !$omp target teams map(to: B, C)  &
S-6      !$omp&             defaultmap(tofrom:scalar) reduction(+:sum)
S-7      !$omp distribute parallel do reduction(+:sum)
S-8         do i = 1,N
S-9            sum = sum + B(i) * C(i)
S-10        end do
S-11     !$omp end target teams
S-12  end function
S-13
S-14  ! Note:  The variable sum is now mapped with tofrom from the defaultmap
S-15  !  clause on the combined target teams construct, for correct
S-16  !  execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.12.4 `target teams` and Distribute Parallel Loop Constructs with Scheduling Clauses

The following example shows how the **`target teams`** and distribute parallel loop constructs are used to execute a **`target`** region. The **`teams`** construct creates a league of at most eight teams where the master thread of each team executes the **`teams`** region. The number of threads in each team is less than or equal to 16.

The **`distribute`** parallel loop construct schedules the subsequent loop iterations across the master threads of each team and then across the threads of each team.

The **`dist_schedule`** clause on the distribute parallel loop construct indicates that loop iterations are distributed to the master thread of each team in chunks of 1024 iterations.

The **`schedule`** clause indicates that the 1024 iterations distributed to a master thread are then assigned to the threads in its associated team in chunks of 64 iterations.

C / C++

Example teams.4.c (**`omp_4.0`**)

```
S-1   #define N 1024*1024
S-2   float dotprod(float B[], float C[])
S-3   {
S-4       float sum = 0.0;
S-5       int i;
S-6       #pragma omp target map(to: B[0:N], C[0:N]) map(tofrom: sum)
S-7       #pragma omp teams num_teams(8) thread_limit(16) reduction(+:sum)
S-8       #pragma omp distribute parallel for reduction(+:sum) \
S-9                   dist_schedule(static, 1024) schedule(static, 64)
S-10      for (i=0; i<N; i++)
S-11          sum += B[i] * C[i];
S-12      return sum;
S-13  }
S-14
S-15  /* Note:  The variable sum is now mapped with tofrom, for correct
S-16     execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
S-17   */
```

C / C++

Fortran

Example teams.4.f90 (`omp_4.0`)

```
S-1   module arrays
S-2   integer,parameter :: N=1024*1024
S-3   real :: B(N), C(N)
S-4   end module
S-5   function dotprod() result(sum)
S-6   use arrays
S-7      real    :: sum
S-8      integer :: i
S-9      sum = 0.0e0
S-10     !$omp target map(to: B, C) map(tofrom: sum)
S-11     !$omp teams num_teams(8) thread_limit(16) reduction(+:sum)
S-12     !$omp distribute parallel do reduction(+:sum) &
S-13     !$omp&  dist_schedule(static, 1024) schedule(static, 64)
S-14        do i = 1,N
S-15           sum = sum + B(i) * C(i)
S-16        end do
S-17     !$omp end teams
S-18     !$omp end target
S-19  end function
S-20
S-21  ! Note:  The variable sum is now mapped with tofrom, for correct
S-22  ! execution with 4.5 (and pre-4.5) compliant compilers. See Devices Intro.
```

Fortran

4.12.5 `target teams` and `distribute simd` Constructs

The following example shows how the **`target teams`** and **`distribute simd`** constructs are used to execute a loop in a **`target`** region. The **`target teams`** construct creates a league of teams where the master thread of each team executes the **`teams`** region.

The **`distribute simd`** construct schedules the loop iterations across the master thread of each team and then uses SIMD parallelism to execute the iterations.

C / C++

Example teams.5.c (**omp_4.0**)

```
S-1   extern void init(float *, float *, int);
S-2   extern void output(float *, int);
S-3   void vec_mult(float *p, float *v1, float *v2, int N)
S-4   {
S-5      int i;
S-6      init(v1, v2, N);
S-7      #pragma omp target teams map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-8      #pragma omp distribute simd
S-9      for (i=0; i<N; i++)
S-10       p[i] = v1[i] * v2[i];
S-11     output(p, N);
S-12  }
```

C / C++

Fortran

Example teams.5.f90 (**omp_4.0**)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2      real    ::  p(N), v1(N), v2(N)
S-3      integer ::  i
S-4      call init(v1, v2, N)
S-5      !$omp target teams map(to: v1, v2) map(from: p)
S-6         !$omp distribute simd
S-7            do i=1,N
S-8               p(i) = v1(i) * v2(i)
S-9            end do
S-10     !$omp end target teams
S-11     call output(p, N)
S-12  end subroutine
```

Fortran

4.12.6 `target teams` and Distribute Parallel Loop SIMD Constructs

The following example shows how the **`target teams`** and the distribute parallel loop SIMD constructs are used to execute a loop in a **`target teams`** region. The **`target teams`** construct creates a league of teams where the master thread of each team executes the **`teams`** region.

The distribute parallel loop SIMD construct schedules the loop iterations across the master thread of each team and then across the threads of each team where each thread uses SIMD parallelism.

C / C++

Example teams.6.c (**`omp_4.0`**)

```
S-1   extern void init(float *, float *, int);
S-2   extern void output(float *, int);
S-3   void vec_mult(float *p, float *v1, float *v2, int N)
S-4   {
S-5      int i;
S-6      init(v1, v2, N);
S-7      #pragma omp target teams map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-8      #pragma omp distribute parallel for simd
S-9      for (i=0; i<N; i++)
S-10       p[i] = v1[i] * v2[i];
S-11     output(p, N);
S-12  }
```

C / C++

Fortran

Example teams.6.f90 (**`omp_4.0`**)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2      real    ::  p(N), v1(N), v2(N)
S-3      integer ::  i
S-4      call init(v1, v2, N)
S-5      !$omp target teams map(to: v1, v2) map(from: p)
S-6         !$omp distribute parallel do simd
S-7            do i=1,N
S-8               p(i) = v1(i) * v2(i)
S-9            end do
S-10     !$omp end target teams
S-11     call output(p, N)
S-12  end subroutine
```

Fortran

4.13 Asynchronous `target` Execution and Dependences

Asynchronous execution of a **target** region can be accomplished by creating an explicit task around the **target** region. Examples with explicit tasks are shown at the beginning of this section.

As of OpenMP 4.5 and beyond the **nowait** clause can be used on the **target** directive for asynchronous execution. Examples with **nowait** clauses follow the explicit **task** examples.

This section also shows the use of **depend** clauses to order executions through dependences.

4.13.1 Asynchronous `target` with Tasks

The following example shows how the **task** and **target** constructs are used to execute multiple **target** regions asynchronously. The task that encounters the **task** construct generates an explicit task that contains a **target** region. The thread executing the explicit task encounters a task scheduling point while waiting for the execution of the **target** region to complete, allowing the thread to switch back to the execution of the encountering task or one of the previously generated explicit tasks.

C / C++

Example async_target.1.c (**omp_4.0**)

```
S-1   #pragma omp declare target
S-2   float F(float);
S-3   #pragma omp end declare target
S-4   #define N 1000000000
S-5   #define CHUNKSZ 1000000
S-6   void init(float *, int);
S-7   float Z[N];
S-8   void pipedF(){
S-9      int C, i;
S-10     init(Z, N);
S-11     for (C=0; C<N; C+=CHUNKSZ){
S-12        #pragma omp task shared(Z)
S-13        #pragma omp target map(Z[C:CHUNKSZ])
S-14        #pragma omp parallel for
S-15        for (i=0; i<CHUNKSZ; i++) Z[i] = F(Z[i]);
S-16     }
S-17     #pragma omp taskwait
S-18  }
```

C / C++

The Fortran version has an interface block that contains the **declare target**. An identical statement exists in the function declaration (not shown here).

Fortran

Example async_target.1.f90 (**omp_4.0**)

```
S-1   module parameters
S-2   integer, parameter :: N=1000000000, CHUNKSZ=1000000
S-3   end module
S-4   subroutine pipedF()
S-5   use parameters, ONLY: N, CHUNKSZ
S-6   integer             :: C, i
S-7   real                :: z(N)
S-8
S-9   interface
S-10     function F(z)
S-11     !$omp declare target
S-12       real, intent(IN) ::z
S-13       real             ::F
S-14     end function F
S-15  end interface
S-16
S-17     call init(z,N)
S-18
S-19     do C=1,N,CHUNKSZ
S-20
S-21        !$omp task shared(z)
S-22        !$omp target map(z(C:C+CHUNKSZ-1))
S-23        !$omp parallel do
S-24           do i=C,C+CHUNKSZ-1
S-25              z(i) = F(z(i))
S-26           end do
S-27        !$omp end target
S-28        !$omp end task
S-29
S-30     end do
S-31     !$omp taskwait
S-32     print*, z
S-33
S-34  end subroutine pipedF
```

Fortran

The following example shows how the **task** and **target** constructs are used to execute multiple **target** regions asynchronously. The task dependence ensures that the storage is allocated and initialized on the device before it is accessed.

C / C++

Example async_target.2.c (**omp_4.0**)

```
S-1   #include <stdlib.h>
S-2   #include <omp.h>
S-3   #pragma omp declare target
S-4   extern void init(float *, float *, int);
S-5   #pragma omp end declare target
S-6   extern void foo();
S-7   extern void output(float *, int);
S-8   void vec_mult(float *p, int N, int dev)
S-9   {
S-10     float *v1, *v2;
S-11     int i;
S-12     #pragma omp task shared(v1, v2) depend(out: v1, v2)
S-13     #pragma omp target device(dev) map(v1, v2)
S-14     {
S-15         // check whether on device dev
S-16         if (omp_is_initial_device())
S-17            abort();
S-18         v1 = (float *)malloc(N*sizeof(float));
S-19         v2 = (float *)malloc(N*sizeof(float));
S-20         init(v1, v2, N);
S-21     }
S-22     foo(); // execute other work asychronously
S-23     #pragma omp task shared(v1, v2, p) depend(in: v1, v2)
S-24     #pragma omp target device(dev) map(to: v1, v2) map(from: p[0:N])
S-25     {
S-26         // check whether on device dev
S-27         if (omp_is_initial_device())
S-28            abort();
S-29         #pragma omp parallel for
S-30         for (i=0; i<N; i++)
S-31           p[i] = v1[i] * v2[i];
S-32         free(v1);
S-33         free(v2);
S-34     }
S-35     #pragma omp taskwait
S-36     output(p, N);
S-37  }
```

C / C++

The Fortran example below is similar to the C version above. Instead of pointers, though, it uses the convenience of Fortran allocatable arrays on the device. In order to preserve the arrays allocated on the device across multiple **target** regions, a **target data** region is used in this case.

If there is no shape specified for an allocatable array in a **map** clause, only the array descriptor (also called a dope vector) is mapped. That is, device space is created for the descriptor, and it is initially populated with host values. In this case, the *v1* and *v2* arrays will be in a non-associated state on the device. When space for *v1* and *v2* is allocated on the device in the first **target** region the addresses to the space will be included in their descriptors.

At the end of the first **target** region, the arrays *v1* and *v2* are preserved on the device for access in the second **target** region. At the end of the second **target** region, the data in array *p* is copied back, the arrays *v1* and *v2* are not.

A **depend** clause is used in the **task** directive to provide a wait at the beginning of the second **target** region, to insure that there is no race condition with *v1* and *v2* in the two tasks. It would be noncompliant to use *v1* and/or *v2* in lieu of *N* in the **depend** clauses, because the use of non-allocated allocatable arrays as list items in a **depend** clause would lead to unspecified behavior.

Note – This example is not strictly compliant with the OpenMP 4.5 specification since the allocation status of allocatable arrays *v1* and *v2* is changed inside the **target** region, which is not allowed. (See the restrictions for the **map** clause in the *Data-mapping Attribute Rules and Clauses* section of the specification.) However, the intention is to relax the restrictions on mapping of allocatable variables in the next release of the specification so that the example will be compliant.

Fortran

Example async_target.2.f90 (**omp_4.0**)

```
S-1   subroutine mult(p,  N, idev)
S-2     use omp_lib, ONLY: omp_is_initial_device
S-3     real              :: p(N)
S-4     real,allocatable :: v1(:), v2(:)
S-5     integer ::  i, idev
S-6     !$omp declare target (init)
S-7
S-8     !$omp target data map(v1,v2)
S-9
S-10    !$omp task shared(v1,v2) depend(out: N)
S-11       !$omp target device(idev)
S-12          if( omp_is_initial_device() ) &
S-13             stop "not executing on target device"
S-14          allocate(v1(N), v2(N))
S-15          call init(v1,v2,N)
S-16       !$omp end target
S-17    !$omp end task
S-18
S-19    call foo()  ! execute other work asychronously
S-20
S-21    !$omp task shared(v1,v2,p) depend(in: N)
S-22       !$omp target device(idev) map(from: p)
```

```
S-23            if( omp_is_initial_device() ) &
S-24               stop "not executing on target device"
S-25            !$omp parallel do
S-26               do i = 1,N
S-27                  p(i) = v1(i) * v2(i)
S-28               end do
S-29            deallocate(v1,v2)
S-30
S-31         !$omp end target
S-32      !$omp end task
S-33
S-34      !$omp taskwait
S-35
S-36      !$omp end target data
S-37
S-38      call output(p, N)
S-39
S-40   end subroutine
```

Fortran

4.13.2 `nowait` Clause on `target` Construct

The following example shows how to execute code asynchronously on a device without an explicit task. The **`nowait`** clause on a **`target`** construct allows the thread of the *target task* to perform other work while waiting for the **`target`** region execution to complete. Hence, the the **`target`** region can execute asynchronously on the device (without requiring a host thread to idle while waiting for the *target task* execution to complete).

In this example the product of two vectors (arrays), *v1* and *v2*, is formed. One half of the operations is performed on the device, and the last half on the host, concurrently.

After a team of threads is formed the master thread generates the *target task* while the other threads can continue on, without a barrier, to the execution of the host portion of the vector product. The completion of the *target task* (asynchronous target execution) is guaranteed by the synchronization in the implicit barrier at the end of the host vector-product worksharing loop region. See the **`barrier`** glossary entry in the OpenMP specification for details.

The host loop scheduling is **`dynamic`**, to balance the host thread executions, since one thread is being used for offload generation. In the situation where little time is spent by the *target task* in setting up and tearing down the the target execution, **`static`** scheduling may be desired.

C / C++

Example async_target.3.c (**omp_4.5**)

```
#include <stdio.h>

#define N 1000000      //N must be even
void init(int n, float *v1, float *v2);

int main(){
   int   i, n=N;
   int   chunk=1000;
   float v1[N],v2[N],vxv[N];

   init(n, v1,v2);

   #pragma omp parallel
   {

      #pragma omp master
      #pragma omp target teams distribute parallel for nowait \
                                map(to: v1[0:n/2]) \
                                map(to: v2[0:n/2]) \
                                map(from: vxv[0:n/2])
      for(i=0; i<n/2; i++){ vxv[i] = v1[i]*v2[i]; }

      #pragma omp for schedule(dynamic,chunk)
      for(i=n/2; i<n; i++){ vxv[i] = v1[i]*v2[i]; }

   }
   printf(" vxv[0] vxv[n-1] %f %f\n", vxv[0], vxv[n-1]);
   return 0;
}
```

C / C++

Fortran

Example async_target.3.f90 (**omp_4.5**)

```
program concurrent_async
   use omp_lib
   integer,parameter :: n=1000000  !!n must be even
   integer           :: i, chunk=1000
   real              :: v1(n),v2(n),vxv(n)

   call init(n, v1,v2)

```

```
S-10        !$omp parallel
S-11
S-12           !$omp master
S-13           !$omp target teams distribute parallel do nowait &
S-14           !$omp&                      map(to: v1(1:n/2))   &
S-15           !$omp&                      map(to: v2(1:n/2))   &
S-16           !$omp&                      map(from: vxv(1:n/2))
S-17           do i = 1,n/2;    vxv(i) = v1(i)*v2(i); end do
S-18           !$omp end master
S-19
S-20           !$omp do schedule(dynamic,chunk)
S-21           do i = n/2+1,n;  vxv(i) = v1(i)*v2(i); end do
S-22
S-23        !$omp end parallel
S-24
S-25        print*, " vxv(1) vxv(n) :", vxv(1), vxv(n)
S-26
S-27     end program
```

Fortran

4.13.3 Asynchronous `target` with `nowait` and `depend` Clauses

More details on dependences can be found in Section 3.3 on page 90, Task Dependences. In this example, there are three flow dependences. In the first two dependences the target task does not execute until the preceding explicit tasks have finished. These dependences are produced by arrays *v1* and *v2* with the **`out`** dependence type in the first two tasks, and the **`in`** dependence type in the target task.

The last dependence is produced by array *p* with the **`out`** dependence type in the target task, and the **`in`** dependence type in the last task. The last task does not execute until the target task finishes.

The **`nowait`** clause on the **`target`** construct creates a deferrable *target task*, allowing the encountering task to continue execution without waiting for the completion of the *target task*.

C / C++

Example async_target.4.c (**omp_4.5**)

```
S-1
S-2   extern void init(  float*, int);
S-3   extern void output(float*, int);
S-4
S-5   void vec_mult(int N)
S-6   {
S-7      int i;
S-8      float p[N], v1[N], v2[N];
S-9
S-10     #pragma omp parallel num_threads(2)
S-11     {
S-12        #pragma omp single
S-13        {
S-14           #pragma omp task depend(out:v1)
S-15           init(v1, N);
S-16
S-17           #pragma omp task depend(out:v2)
S-18           init(v2, N);
S-19
S-20           #pragma omp target nowait depend(in:v1,v2) depend(out:p) \
S-21                                        map(to:v1,v2) map( from: p)
S-22           #pragma omp parallel for private(i)
S-23           for (i=0; i<N; i++)
S-24              p[i] = v1[i] * v2[i];
S-25
S-26           #pragma omp task depend(in:p)
S-27           output(p, N);
S-28        }
S-29     }
S-30  }
```

C / C++

Fortran

Example async_target.4.f90 (**omp_4.5**)

```
S-1
S-2   subroutine vec_mult(N)
S-3      implicit none
S-4      integer           :: i, N
S-5      real, allocatable :: p(:), v1(:), v2(:)
S-6      allocate( p(N), v1(N), v2(N) )
S-7
S-8      !$omp parallel num_threads(2)
S-9
```

```
S-10          !$omp single
S-11
S-12             !$omp task depend(out:v1)
S-13             call init(v1, N)
S-14             !$omp end task
S-15
S-16             !$omp task depend(out:v2)
S-17             call init(v2, N)
S-18             !$omp end task
S-19
S-20             !$omp target nowait depend(in:v1,v2) depend(out:p) &
S-21             !$omp&                  map(to:v1,v2)  map(from: p)
S-22             !$omp parallel do
S-23             do i=1,N
S-24                p(i) = v1(i) * v2(i)
S-25             end do
S-26             !$omp end target
S-27
S-28
S-29             !$omp task depend(in:p)
S-30             call output(p, N)
S-31             !$omp end task
S-32
S-33           !$omp end single
S-34        !$omp end parallel
S-35
S-36        deallocate( p, v1, v2 )
S-37
S-38     end subroutine
```

Fortran

4.14 Device Routines

4.14.1 `omp_is_initial_device` Routine

The following example shows how the **`omp_is_initial_device`** runtime library routine can be used to query if a code is executing on the initial host device or on a target device. The example then sets the number of threads in the **`parallel`** region based on where the code is executing.

C / C++

Example device.1.c (**`omp_4.0`**)

```
S-1   #include <stdio.h>
S-2   #include <omp.h>
S-3
S-4   #pragma omp declare target
S-5   void vec_mult(float *p, float *v1, float *v2, int N);
S-6   extern float *p, *v1, *v2;
S-7   extern int N;
S-8   #pragma omp end declare target
S-9
S-10  extern void init_vars(float *, float *, int);
S-11  extern void output(float *, int);
S-12
S-13  void foo()
S-14  {
S-15     init_vars(v1, v2, N);
S-16     #pragma omp target device(42) map(p[:N], v1[:N], v2[:N])
S-17     {
S-18        vec_mult(p, v1, v2, N);
S-19     }
S-20     output(p, N);
S-21  }
S-22
S-23  void vec_mult(float *p, float *v1, float *v2, int N)
S-24  {
S-25     int i;
S-26     int nthreads;
S-27     if (!omp_is_initial_device())
S-28     {
S-29        printf("1024 threads on target device\n");
S-30        nthreads = 1024;
S-31     }
S-32     else
S-33     {
S-34        printf("8 threads on initial device\n");
S-35        nthreads = 8;
S-36     }
```

```
S-37      #pragma omp parallel for private(i) num_threads(nthreads)
S-38      for (i=0; i<N; i++)
S-39        p[i] = v1[i] * v2[i];
S-40   }
```

C / C++

Fortran

1 *Example device.1.f90* (`omp_4.0`)

```
 S-1   module params
 S-2      integer,parameter :: N=1024
 S-3   end module params
 S-4   module vmult
 S-5   contains
 S-6      subroutine vec_mult(p, v1, v2, N)
 S-7      use omp_lib, ONLY : omp_is_initial_device
 S-8      !$omp declare target
 S-9      real    :: p(N), v1(N), v2(N)
S-10      integer :: i, nthreads, N
S-11         if (.not. omp_is_initial_device()) then
S-12            print*, "1024 threads on target device"
S-13            nthreads = 1024
S-14         else
S-15            print*, "8 threads on initial device"
S-16            nthreads = 8
S-17         endif
S-18         !$omp parallel do private(i) num_threads(nthreads)
S-19         do i = 1,N
S-20           p(i) = v1(i) * v2(i)
S-21         end do
S-22      end subroutine vec_mult
S-23   end module vmult
S-24   program prog_vec_mult
S-25   use params
S-26   use vmult
S-27   real :: p(N), v1(N), v2(N)
S-28      call init(v1,v2,N)
S-29      !$omp target device(42) map(p, v1, v2)
S-30         call vec_mult(p, v1, v2, N)
S-31      !$omp end target
S-32      call output(p, N)
S-33   end program
```

Fortran

4.14.2 `omp_get_num_devices` Routine

The following example shows how the **`omp_get_num_devices`** runtime library routine can be used to determine the number of devices.

C / C++

Example device.2.c (**`omp_4.0`**)

```
S-1   #include <omp.h>
S-2   extern void init(float *, float *, int);
S-3   extern void output(float *, int);
S-4   void vec_mult(float *p, float *v1, float *v2, int N)
S-5   {
S-6      int i;
S-7      init(v1, v2, N);
S-8      int ndev = omp_get_num_devices();
S-9      int do_offload = (ndev>0 && N>1000000);
S-10     #pragma omp target if(do_offload) map(to: v1[0:N], v2[:N]) map(from: p[0:N])
S-11     #pragma omp parallel for if(N>1000) private(i)
S-12     for (i=0; i<N; i++)
S-13       p[i] = v1[i] * v2[i];
S-14     output(p, N);
S-15  }
```

C / C++

Fortran

Example device.2.f90 (**`omp_4.0`**)

```
S-1   subroutine vec_mult(p, v1, v2, N)
S-2   use omp_lib, ONLY : omp_get_num_devices
S-3   real    :: p(N), v1(N), v2(N)
S-4   integer :: N, i, ndev
S-5   logical :: do_offload
S-6      call init(v1, v2, N)
S-7      ndev = omp_get_num_devices()
S-8      do_offload = (ndev>0) .and. (N>1000000)
S-9      !$omp target if(do_offload) map(to: v1, v2) map(from: p)
S-10     !$omp parallel do if(N>1000)
S-11        do i=1,N
S-12           p(i) = v1(i) * v2(i)
S-13        end do
S-14     !$omp end target
S-15     call output(p, N)
S-16  end subroutine
```

Fortran

4.14.3 `omp_set_default_device` and `omp_get_default_device` Routines

The following example shows how the **`omp_set_default_device`** and **`omp_get_default_device`** runtime library routines can be used to set the default device and determine the default device respectively.

C / C++

Example device.3.c (**`omp_4.0`**)

```
S-1   #include <omp.h>
S-2   #include <stdio.h>
S-3   void foo(void)
S-4   {
S-5      int default_device = omp_get_default_device();
S-6      printf("Default device = %d\n", default_device);
S-7      omp_set_default_device(default_device+1);
S-8      if (omp_get_default_device() != default_device+1)
S-9         printf("Default device is still = %d\n", default_device);
S-10  }
```

C / C++

Fortran

Example device.3.f90 (**`omp_4.0`**)

```
S-1   program foo
S-2   use omp_lib, ONLY : omp_get_default_device, omp_set_default_device
S-3   integer :: old_default_device, new_default_device
S-4      old_default_device = omp_get_default_device()
S-5      print*, "Default device = ", old_default_device
S-6      new_default_device = old_default_device + 1
S-7      call omp_set_default_device(new_default_device)
S-8      if (omp_get_default_device() == old_default_device) &
S-9         print*,"Default device is STILL = ", old_default_device
S-10  end program
```

Fortran

4.14.4 Target Memory and Device Pointers Routines

The following example shows how to create space on a device, transfer data to and from that space, and free the space, using API calls. The API calls directly execute allocation, copy and free operations on the device, without invoking any mapping through a **target** directive. The **omp_target_alloc** routine allocates space and returns a device pointer for referencing the space in the **omp_target_memcpy** API routine on the host. The **omp_target_free** routine frees the space on the device.

The example also illustrates how to access that space in a **target** region by exposing the device pointer in an **is_device_ptr** clause.

The example creates an array of cosine values on the default device, to be used on the host device. The function fails if a default device is not available.

C / C++

Example device.4.c (**omp_4.5**)

```
S-1   #include <stdio.h>
S-2   #include <math.h>
S-3   #include <stdlib.h>
S-4   #include <omp.h>
S-5
S-6   void get_dev_cos(double *mem, size_t s)
S-7   {
S-8      int h, t, i;
S-9      double * mem_dev_cpy;
S-10     h = omp_get_initial_device();
S-11     t = omp_get_default_device();
S-12
S-13     if (omp_get_num_devices() < 1 || t < 0){
S-14        printf(" ERROR: No device found.\n");
S-15        exit(1);
S-16     }
S-17
S-18     mem_dev_cpy = (double *)omp_target_alloc( sizeof(double) * s, t);
S-19     if(mem_dev_cpy == NULL){
S-20        printf(" ERROR: No space left on device.\n");
S-21        exit(1);
S-22     }
S-23
S-24                            /* dst  src */
S-25     omp_target_memcpy(mem_dev_cpy, mem, sizeof(double)*s,
S-26                                0,    0,
S-27                                t,   h);
S-28
S-29     #pragma omp target is_device_ptr(mem_dev_cpy) device(t)
S-30     #pragma omp teams distribute parallel for
```

```
S-31        for(i=0;i<s;i++){ mem_dev_cpy[i] = cos((double)i); } /* init data */
S-32
S-33                        /* dst  src */
S-34        omp_target_memcpy(mem, mem_dev_cpy, sizeof(double)*s,
S-35                          0,              0,
S-36                          h,              t);
S-37
S-38        omp_target_free(mem_dev_cpy, t);
S-39    }
```

C / C++

CHAPTER 5

SIMD

Single instruction, multiple data (SIMD) is a form of parallel execution in which the same operation is performed on multiple data elements independently in hardware vector processing units (VPU), also called SIMD units. The addition of two vectors to form a third vector is a SIMD operation. Many processors have SIMD (vector) units that can perform simultaneously 2, 4, 8 or more executions of the same operation (by a single SIMD unit).

Loops without loop-carried backward dependency (or with dependency preserved using ordered simd) are candidates for vectorization by the compiler for execution with SIMD units. In addition, with state-of-the-art vectorization technology and **`declare simd`** construct extensions for function vectorization in the OpenMP 4.5 specification, loops with function calls can be vectorized as well. The basic idea is that a scalar function call in a loop can be replaced by a vector version of the function, and the loop can be vectorized simultaneously by combining a loop vectorization (**`simd`** directive on the loop) and a function vectorization (**`declare simd`** directive on the function).

A **`simd`** construct states that SIMD operations be performed on the data within the loop. A number of clauses are available to provide data-sharing attributes (**`private`**, **`linear`**, **`reduction`** and **`lastprivate`**). Other clauses provide vector length preference/restrictions (**`simdlen`** / **`safelen`**), loop fusion (**`collapse`**), and data alignment (**`aligned`**).

The **`declare simd`** directive designates that a vector version of the function should also be constructed for execution within loops that contain the function and have a **`simd`** directive. Clauses provide argument specifications (**`linear`**, **`uniform`**, and **`aligned`**), a requested vector length (**`simdlen`**), and designate whether the function is always/never called conditionally in a loop (**`branch`**/**`inbranch`**). The latter is for optimizing performance.

Also, the **`simd`** construct has been combined with the worksharing loop constructs (**`for simd`** and **`do simd`**) to enable simultaneous thread execution in different SIMD units.

5.1 `simd` and `declare simd` Constructs

The following example illustrates the basic use of the **`simd`** construct to assure the compiler that the loop can be vectorized.

C / C++

Example SIMD.1.c (**`omp_4.0`**)

```
S-1   void star( double *a, double *b, double *c, int n, int *ioff )
S-2   {
S-3      int i;
S-4      #pragma omp simd
S-5      for ( i = 0; i < n; i++ )
S-6         a[i] *= b[i] * c[i+ *ioff];
S-7   }
```

C / C++

Fortran

Example SIMD.1.f90 (**`omp_4.0`**)

```
S-1   subroutine star(a,b,c,n,ioff_ptr)
S-2      implicit none
S-3      double precision :: a(*),b(*),c(*)
S-4      integer          :: n, i
S-5      integer, pointer :: ioff_ptr
S-6
S-7      !$omp simd
S-8      do i = 1,n
S-9         a(i) = a(i) * b(i) * c(i+ioff_ptr)
S-10     end do
S-11
S-12  end subroutine
```

Fortran

When a function can be inlined within a loop the compiler has an opportunity to vectorize the loop. By guaranteeing SIMD behavior of a function's operations, characterizing the arguments of the function and privatizing temporary variables of the loop, the compiler can often create faster, vector code for the loop. In the examples below the **`declare simd`** construct is used on the *add1* and *add2* functions to enable creation of their corresponding SIMD function versions for execution within the associated SIMD loop. The functions characterize two different approaches of accessing data within the function: by a single variable and as an element in a data array, respectively. The *add3* C function uses dereferencing.

The **`declare simd`** constructs also illustrate the use of **`uniform`** and **`linear`** clauses. The **`uniform(fact)`** clause indicates that the variable *fact* is invariant across the SIMD lanes. In the *add2* function *a* and *b* are included in the **`uniform`** list because the C pointer and the Fortran array references are constant. The *i* index used in the *add2* function is included in a **`linear`** clause with a constant-linear-step of 1, to guarantee a unity increment of the associated loop. In the **`declare simd`** construct for the *add3* C function the **`linear(a,b:1)`** clause instructs the compiler to generate unit-stride loads across the SIMD lanes; otherwise, costly *gather* instructions would be generated for the unknown sequence of access of the pointer dereferences.

In the **`simd`** constructs for the loops the **`private(tmp)`** clause is necessary to assure that the each vector operation has its own *tmp* variable.

C / C++

Example SIMD.2.c (**`omp_4.0`**)

```
S-1   #include <stdio.h>
S-2
S-3   #pragma omp declare simd uniform(fact)
S-4   double add1(double a, double b, double fact)
S-5   {
S-6      double c;
S-7      c = a + b + fact;
S-8      return c;
S-9   }
S-10
S-11  #pragma omp declare simd uniform(a,b,fact) linear(i:1)
S-12  double add2(double *a, double *b, int i, double fact)
S-13  {
S-14     double c;
S-15     c = a[i] + b[i] + fact;
S-16     return c;
S-17  }
S-18
S-19  #pragma omp declare simd uniform(fact) linear(a,b:1)
S-20  double add3(double *a, double *b, double fact)
S-21  {
S-22     double c;
S-23     c = *a + *b + fact;
```

```
S-24        return c;
S-25     }
S-26
S-27     void work( double *a, double *b, int n )
S-28     {
S-29        int i;
S-30        double tmp;
S-31        #pragma omp simd private(tmp)
S-32        for ( i = 0; i < n; i++ ) {
S-33           tmp  = add1( a[i],  b[i], 1.0);
S-34           a[i] = add2( a,     b, i, 1.0) + tmp;
S-35           a[i] = add3(&a[i], &b[i], 1.0);
S-36        }
S-37     }
S-38
S-39     int main(){
S-40        int i;
S-41        const int N=32;
S-42        double a[N], b[N];
S-43
S-44        for ( i=0; i<N; i++ ) {
S-45           a[i] = i; b[i] = N-i;
S-46        }
S-47
S-48        work(a, b, N );
S-49
S-50        for ( i=0; i<N; i++ ) {
S-51           printf("%d %f\n", i, a[i]);
S-52        }
S-53
S-54        return 0;
S-55     }
```

C / C++

Fortran

1 *Example SIMD.2.f90* (`omp_4.0`)

```
S-1      program main
S-2         implicit none
S-3         integer, parameter :: N=32
S-4         integer :: i
S-5         double precision   :: a(N), b(N)
S-6         do i = 1,N
S-7            a(i) = i-1
S-8            b(i) = N-(i-1)
S-9         end do
S-10        call work(a, b, N )
```

```
S-11        do i = 1,N
S-12           print*, i,a(i)
S-13        end do
S-14     end program
S-15
S-16     function add1(a,b,fact) result(c)
S-17        implicit none
S-18     !$omp declare simd(add1) uniform(fact)
S-19        double precision :: a,b,fact, c
S-20        c = a + b + fact
S-21     end function
S-22
S-23     function add2(a,b,i, fact) result(c)
S-24        implicit none
S-25     !$omp declare simd(add2) uniform(a,b,fact) linear(i:1)
S-26        integer          :: i
S-27        double precision :: a(*),b(*),fact, c
S-28        c = a(i) + b(i) + fact
S-29     end function
S-30
S-31     subroutine work(a, b, n )
S-32        implicit none
S-33        double precision           :: a(n),b(n), tmp
S-34        integer                    :: n, i
S-35        double precision, external :: add1, add2
S-36
S-37        !$omp simd private(tmp)
S-38        do i = 1,n
S-39           tmp  = add1(a(i), b(i), 1.0d0)
S-40           a(i) = add2(a,    b, i, 1.0d0) + tmp
S-41           a(i) = a(i) + b(i) + 1.0d0
S-42        end do
S-43     end subroutine
```

Fortran

A thread that encounters a SIMD construct executes a vectorized code of the iterations. Similar to the concerns of a worksharing loop a loop vectorized with a SIMD construct must assure that temporary and reduction variables are privatized and declared as reductions with clauses. The example below illustrates the use of **`private`** and **`reduction`** clauses in a SIMD construct.

C / C++

Example SIMD.3.c (**`omp_4.0`**)

```
S-1   double work( double *a, double *b, int n )
S-2   {
S-3      int i;
S-4      double tmp, sum;
S-5      sum = 0.0;
S-6      #pragma omp simd private(tmp) reduction(+:sum)
S-7      for (i = 0; i < n; i++) {
S-8         tmp = a[i] + b[i];
S-9         sum += tmp;
S-10     }
S-11     return sum;
S-12  }
```

C / C++

Fortran

Example SIMD.3.f90 (**`omp_4.0`**)

```
S-1   subroutine work( a, b, n, sum )
S-2      implicit none
S-3      integer :: i, n
S-4      double precision :: a(n), b(n), sum, tmp
S-5
S-6      sum = 0.0d0
S-7      !$omp simd private(tmp) reduction(+:sum)
S-8      do i = 1,n
S-9         tmp = a(i) + b(i)
S-10        sum = sum + tmp
S-11     end do
S-12
S-13  end subroutine work
```

Fortran

A **safelen(N)** clause in a **simd** construct assures the compiler that there are no loop-carried dependencies for vectors of size *N* or below. If the **safelen** clause is not specified, then the default safelen value is the number of loop iterations.

The **safelen(16)** clause in the example below guarantees that the vector code is safe for vectors up to and including size 16. In the loop, *m* can be 16 or greater, for correct code execution. If the value of *m* is less than 16, the behavior is undefined.

C / C++

Example SIMD.4.c (**omp_4.0**)

```
S-1   void work( float *b, int n, int m )
S-2   {
S-3      int i;
S-4      #pragma omp simd safelen(16)
S-5      for (i = m; i < n; i++)
S-6         b[i] = b[i-m] - 1.0f;
S-7   }
```

C / C++

Fortran

Example SIMD.4.f90 (**omp_4.0**)

```
S-1   subroutine work( b, n, m )
S-2      implicit none
S-3      real        :: b(n)
S-4      integer     :: i,n,m
S-5
S-6      !$omp simd safelen(16)
S-7      do i = m+1, n
S-8         b(i) = b(i-m) - 1.0
S-9      end do
S-10  end subroutine work
```

Fortran

The following SIMD construct instructs the compiler to collapse the *i* and *j* loops into a single SIMD loop in which SIMD chunks are executed by threads of the team. Within the workshared loop chunks of a thread, the SIMD chunks are executed in the lanes of the vector units.

C / C++

Example SIMD.5.c (`omp_4.0`)

```
S-1   void work( double **a, double **b, double **c, int n )
S-2   {
S-3      int i, j;
S-4      double tmp;
S-5      #pragma omp for simd collapse(2) private(tmp)
S-6      for (i = 0; i < n; i++) {
S-7         for (j = 0; j < n; j++) {
S-8            tmp = a[i][j] + b[i][j];
S-9            c[i][j] = tmp;
S-10        }
S-11     }
S-12  }
```

C / C++

Fortran

Example SIMD.5.f90 (`omp_4.0`)

```
S-1   subroutine work( a, b, c,  n )
S-2      implicit none
S-3      integer :: i,j,n
S-4      double precision :: a(n,n), b(n,n), c(n,n), tmp
S-5
S-6      !$omp do simd collapse(2) private(tmp)
S-7      do j = 1,n
S-8         do i = 1,n
S-9            tmp = a(i,j) + b(i,j)
S-10           c(i,j) = tmp
S-11        end do
S-12     end do
S-13
S-14  end subroutine work
```

Fortran

5.2 `inbranch` and `notinbranch` Clauses

The following examples illustrate the use of the **`declare simd`** construct with the **`inbranch`** and **`notinbranch`** clauses. The **`notinbranch`** clause informs the compiler that the function *foo* is never called conditionally in the SIMD loop of the function *myaddint*. On the other hand, the **`inbranch`** clause for the function goo indicates that the function is always called conditionally in the SIMD loop inside the function *myaddfloat*.

C / C++

Example SIMD.6.c (**`omp_4.0`**)

```
S-1   #pragma omp declare simd linear(p:1) notinbranch
S-2   int foo(int *p){
S-3     *p = *p + 10;
S-4     return *p;
S-5   }
S-6
S-7   int myaddint(int *a, int *b, int n)
S-8   {
S-9   #pragma omp simd
S-10    for (int i=0; i<n; i++){
S-11        a[i]  = foo(&b[i]);  /* foo is not called under a condition */
S-12    }
S-13    return a[n-1];
S-14  }
S-15
S-16  #pragma omp declare simd linear(p:1) inbranch
S-17  float goo(float *p){
S-18    *p = *p + 18.5f;
S-19    return *p;
S-20  }
S-21
S-22  int myaddfloat(float *x, float *y, int n)
S-23  {
S-24  #pragma omp simd
S-25    for (int i=0; i<n; i++){
S-26       x[i] = (x[i] > y[i]) ? goo(&y[i]) : y[i];
S-27         /* goo is called under the condition (or within a branch) */
S-28    }
S-29    return x[n-1];
S-30  }
```

C / C++

Example SIMD.6.f90 (`omp_4.0`)

```
S-1   function foo(p) result(r)
S-2     implicit none
S-3   !$omp declare simd(foo) notinbranch
S-4     integer :: p, r
S-5     p = p + 10
S-6     r = p
S-7   end function foo
S-8
S-9   function myaddint(a, b,  n) result(r)
S-10    implicit none
S-11    integer :: a(*), b(*), n, r
S-12    integer :: i
S-13    integer, external :: foo
S-14
S-15    !$omp simd
S-16    do i=1, n
S-17        a(i) = foo(b(i))  ! foo is not called under a condition
S-18    end do
S-19    r = a(n)
S-20
S-21  end function myaddint
S-22
S-23  function goo(p) result(r)
S-24    implicit none
S-25  !$omp declare simd(goo) inbranch
S-26    real :: p, r
S-27    p = p + 18.5
S-28    r = p
S-29  end function goo
S-30
S-31  function myaddfloat(x, y, n) result(r)
S-32    implicit none
S-33    real :: x(*), y(*), r
S-34    integer :: n
S-35    integer :: i
S-36    real, external :: goo
S-37
S-38    !$omp simd
S-39    do i=1, n
S-40       if (x(i) > y(i)) then
S-41          x(i) = goo(y(i))
S-42          ! goo is called under the condition (or within a branch)
S-43       else
S-44          x(i) = y(i)
```

```
S-45        endif
S-46      end do
S-47
S-48      r = x(n)
S-49    end function myaddfloat
```

Fortran

In the code below, the function *fib()* is called in the main program and also recursively called in the function *fib()* within an **if** condition. The compiler creates a masked vector version and a non-masked vector version for the function *fib()* while retaining the original scalar version of the *fib()* function.

C / C++

Example SIMD.7.c (**omp_4.0**)

```
S-1     #include <stdio.h>
S-2     #include <stdlib.h>
S-3
S-4     #define N 45
S-5     int a[N], b[N], c[N];
S-6
S-7     #pragma omp declare simd inbranch
S-8     int fib( int n )
S-9     {
S-10       if (n <= 1)
S-11          return n;
S-12       else {
S-13          return fib(n-1) + fib(n-2);
S-14       }
S-15    }
S-16
S-17    int main(void)
S-18    {
S-19       int i;
S-20
S-21       #pragma omp simd
S-22       for (i=0; i < N; i++) b[i] = i;
S-23
S-24       #pragma omp simd
S-25       for (i=0; i < N; i++) {
S-26          a[i] = fib(b[i]);
S-27       }
S-28       printf("Done a[%d] = %d\n", N-1, a[N-1]);
S-29       return 0;
S-30    }
```

C / C++

Fortran

Example SIMD.7.f90 (`omp_4.0`)

```
S-1   program fibonacci
S-2      implicit none
S-3      integer,parameter :: N=45
S-4      integer           :: a(0:N-1), b(0:N-1)
S-5      integer           :: i
S-6      integer, external :: fib
S-7
S-8      !$omp simd
S-9      do i = 0,N-1
S-10        b(i) = i
S-11     end do
S-12
S-13     !$omp simd
S-14     do i=0,N-1
S-15        a(i) = fib(b(i))
S-16     end do
S-17
S-18     write(*,*) "Done a(", N-1, ") = ", a(N-1)
S-19                          ! 44  701408733
S-20  end program
S-21
S-22  recursive function fib(n) result(r)
S-23     implicit none
S-24  !$omp declare simd(fib) inbranch
S-25     integer  :: n, r
S-26
S-27     if (n <= 1) then
S-28        r = n
S-29     else
S-30        r = fib(n-1) + fib(n-2)
S-31     endif
S-32
S-33  end function fib
```

Fortran

5.3 Loop-Carried Lexical Forward Dependence

The following example tests the restriction on an SIMD loop with the loop-carried lexical forward-dependence. This dependence must be preserved for the correct execution of SIMD loops.

A loop can be vectorized even though the iterations are not completely independent when it has loop-carried dependences that are forward lexical dependences, indicated in the code below by the read of *A[j+1]* and the write to *A[j]* in C/C++ code (or *A(j+1)* and *A(j)* in Fortran). That is, the read of *A[j+1]* (or *A(j+1)* in Fortran) before the write to *A[j]* (or *A(j)* in Fortran) ordering must be preserved for each iteration in *j* for valid SIMD code generation.

This test assures that the compiler preserves the loop carried lexical forward-dependence for generating a correct SIMD code.

C / C++

Example SIMD.8.c (**omp_4.0**)

```
S-1   #include <stdio.h>
S-2   #include <math.h>
S-3
S-4   int   P[1000];
S-5   float A[1000];
S-6
S-7   float do_work(float *arr)
S-8   {
S-9     float pri;
S-10    int i;
S-11  #pragma omp simd lastprivate(pri)
S-12    for (i = 0; i < 999; ++i) {
S-13      int j = P[i];
S-14
S-15      pri = 0.5f;
S-16      if (j % 2 == 0) {
S-17        pri = A[j+1] + arr[i];
S-18      }
S-19      A[j] = pri * 1.5f;
S-20      pri = pri + A[j];
S-21    }
S-22    return pri;
S-23  }
S-24
S-25  int main(void)
S-26  {
S-27    float pri, arr[1000];
S-28    int i;
S-29
S-30    for (i = 0; i < 1000; ++i) {
```

```
      P[i]   = i;
      A[i]   = i * 1.5f;
      arr[i] = i * 1.8f;
    }
    pri = do_work(&arr[0]);
    if (pri == 8237.25) {
      printf("passed: result pri = %7.2f (8237.25) \n", pri);
    }
    else {
      printf("failed: result pri = %7.2f (8237.25) \n", pri);
    }
    return 0;
}
```

C / C++

Fortran

Example SIMD.8.f90 (`omp_4.0`)

```
module work

integer :: P(1000)
real    :: A(1000)

contains
function do_work(arr) result(pri)
  implicit none
  real, dimension(*) :: arr

  real :: pri
  integer :: i, j

  !$omp simd private(j) lastprivate(pri)
  do i = 1, 999
    j = P(i)

    pri = 0.5
    if (mod(j-1, 2) == 0) then
      pri = A(j+1) + arr(i)
    endif
    A(j) = pri * 1.5
    pri = pri + A(j)
  end do

end function do_work

end module work
```

```
S-30    program simd_8f
S-31      use work
S-32      implicit none
S-33      real :: pri, arr(1000)
S-34      integer :: i
S-35
S-36      do i = 1, 1000
S-37         P(i)   = i
S-38         A(i)   = (i-1) * 1.5
S-39         arr(i) = (i-1) * 1.8
S-40      end do
S-41      pri = do_work(arr)
S-42      if (pri == 8237.25) then
S-43        print 2, "passed", pri
S-44      else
S-45        print 2, "failed", pri
S-46      endif
S-47    2 format(a, ": result pri = ", f7.2, " (8237.25)")
S-48
S-49    end program
```

Fortran

5.4 `ref`, `val`, `uval` Modifiers for `linear` Clause

When generating vector functions from **`declare simd`** directives, it is important for a compiler to know the proper types of function arguments in order to generate efficient codes. This is especially true for C++ reference types and Fortran arguments.

In the following example, the function *add_one2* has a C++ reference parameter (or Fortran argument) *p*. Variable *p* gets incremented by 1 in the function. The caller loop *i* in the main program passes a variable *k* as a reference to the function *add_one2* call. The **`ref`** modifier for the **`linear`** clause on the **`declare simd`** directive is used to annotate the reference-type parameter *p* to match the property of the variable *k* in the loop. This use of reference type is equivalent to the second call to *add_one2* with a direct passing of the array element *a[i]*. In the example, the preferred vector length 8 is specified for both the caller loop and the callee function.

When **`linear(ref(p))`** is applied to an argument passed by reference, it tells the compiler that the addresses in its vector argument are consecutive, and so the compiler can generate a single vector load or store instead of a gather or scatter. This allows more efficient SIMD code to be generated with less source changes.

C++

Example linear_modifier.1.cpp (**omp_4.5**)

```
S-1   #include <stdio.h>
S-2
S-3   #define NN 1023
S-4   int a[NN];
S-5
S-6   #pragma omp declare simd linear(ref(p)) simdlen(8)
S-7   void add_one2(int& p)
S-8   {
S-9      p += 1;
S-10  }
S-11
S-12  int main(void)
S-13  {
S-14     int i;
S-15     int* p = a;
S-16
S-17     for (i = 0; i < NN; i++) {
S-18        a[i] = i;
S-19     }
S-20
S-21  #pragma omp simd linear(p) simdlen(8)
S-22     for (i = 0; i < NN; i++) {
S-23        int& k = *p;
S-24        add_one2(k);
S-25        add_one2(a[i]);
S-26        p++;
S-27     }
S-28
S-29     for (i = 0; i < NN; i++) {
S-30        if (a[i] != i+2) {
S-31           printf("failed\n");
S-32           return 1;
S-33        }
S-34     }
S-35     printf("passed\n");
S-36     return 0;
S-37  }
```

C++

Fortran

Example linear_modifier.1.f90 (**omp_4.5**)

```
S-1   module m
S-2      integer, parameter :: NN = 1023
S-3      integer :: a(NN)
S-4
S-5    contains
S-6      subroutine add_one2(p)
S-7      !$omp declare simd(add_one2) linear(ref(p)) simdlen(8)
S-8      implicit none
S-9      integer :: p
S-10
S-11     p = p + 1
S-12     end subroutine
S-13  end module
S-14
S-15  program main
S-16     use m
S-17     implicit none
S-18     integer :: i, p
S-19
S-20     do i = 1, NN
S-21        a(i) = i
S-22     end do
S-23
S-24     p = 1
S-25     !$omp simd linear(p) simdlen(8)
S-26     do i = 1, NN
S-27        associate(k => a(p))
S-28           call add_one2(k)
S-29        end associate
S-30        call add_one2(a(i))
S-31        p = p + 1
S-32     end do
S-33
S-34     do i = 1, NN
S-35        if (a(i) /= i+2) then
S-36           print *, "failed"
S-37           stop
S-38        endif
S-39     end do
S-40     print *, "passed"
S-41  end program
S-42
```

Fortran

The following example is a variant of the above example. The function *add_one2* in the C++ code includes an additional C++ reference parameter *i*. The loop index *i* of the caller loop *i* in the main program is passed as a reference to the function *add_one2* call. The loop index *i* has a uniform address with linear value of step 1 across SIMD lanes. Thus, the **uval** modifier is used for the **linear** clause to annotate the C++ reference-type parameter *i* to match the property of loop index *i*.

In the correponding Fortran code the arguments *p* and *i* in the routine *add_on2* are passed by references. Similar modifiers are used for these variables in the **linear** clauses to match with the property at the caller loop in the main program.

When **linear(uval(i))** is applied to an argument passed by reference, it tells the compiler that its addresses in the vector argument are uniform so that the compiler can generate a scalar load or scalar store and create linear values. This allows more efficient SIMD code to be generated with less source changes.

C++

Example linear_modifier.2.cpp (**omp_4.5**)

```
S-1   #include <stdio.h>
S-2
S-3   #define NN 1023
S-4   int a[NN];
S-5
S-6   #pragma omp declare simd linear(ref(p)) linear(uval(i))
S-7   void add_one2(int& p, const int& i)
S-8   {
S-9      p += i;
S-10  }
S-11
S-12  int main(void)
S-13  {
S-14     int i;
S-15     int* p = a;
S-16
S-17     for (i = 0; i < NN; i++) {
S-18        a[i] = i;
S-19     }
S-20
S-21     #pragma omp simd linear(p)
S-22     for (i = 0; i < NN; i++) {
S-23        int& k = *p;
S-24        add_one2(k, i);
S-25        p++;
S-26     }
S-27
S-28     for (i = 0; i < NN; i++) {
```

```
S-29          if (a[i] != i*2) {
S-30             printf("failed\n");
S-31             return 1;
S-32          }
S-33       }
S-34       printf("passed\n");
S-35       return 0;
S-36    }
```

C++

Fortran

1 *Example linear_modifier.2.f90* (**omp_4.5**)

```
S-1     module m
S-2        integer, parameter :: NN = 1023
S-3        integer :: a(NN)
S-4
S-5     contains
S-6        subroutine add_one2(p, i)
S-7        !$omp declare simd(add_one2) linear(ref(p)) linear(uval(i))
S-8        implicit none
S-9        integer :: p
S-10       integer, intent(in) :: i
S-11
S-12       p = p + i
S-13       end subroutine
S-14    end module
S-15
S-16    program main
S-17       use m
S-18       implicit none
S-19       integer :: i, p
S-20
S-21       do i = 1, NN
S-22          a(i) = i
S-23       end do
S-24
S-25       p = 1
S-26       !$omp simd linear(p)
S-27       do i = 1, NN
S-28          call add_one2(a(p), i)
S-29          p = p + 1
S-30       end do
S-31
S-32       do i = 1, NN
S-33          if (a(i) /= i*2) then
```

```
S-34            print *, "failed"
S-35            stop
S-36         endif
S-37      end do
S-38      print *, "passed"
S-39   end program
```

Fortran

In the following example, the function *func* takes arrays *x* and *y* as arguments, and accesses the array elements referenced by the index *i*. The caller loop *i* in the main program passes a linear copy of the variable *k* to the function *func*. The **val** modifier is used for the **linear** clause in the **declare simd** directive for the function *func* to annotate argument *i* to match the property of the actual argument *k* passed in the SIMD loop. Arrays *x* and *y* have uniform addresses across SIMD lanes.

When **linear(val(i):1)** is applied to an argument, it tells the compiler that its addresses in the vector argument may not be consecutive, however, their values are linear (with stride 1 here). When the value of *i* is used in subscript of array references (e.g., *x[i]*), the compiler can generate a vector load or store instead of a gather or scatter. This allows more efficient SIMD code to be generated with less source changes.

C / C++

Example linear_modifier.3.c (**omp_4.5**)

```
S-1   #include <stdio.h>
S-2
S-3   #define N 128
S-4
S-5   #pragma omp declare simd simdlen(4) uniform(x, y) linear(val(i):1)
S-6   double func(double x[], double y[], int i)
S-7   {
S-8      return (x[i] + y[i]);
S-9   }
S-10
S-11  int main(void)
S-12  {
S-13     double x[N], y[N], z1[N], z2;
S-14     int i, k;
S-15
S-16     for (i = 0; i < N; i++) {
S-17        x[i] = (double)i;
S-18        y[i] = (double)i*2;
S-19     }
S-20
S-21     k = 0;
S-22  #pragma omp simd linear(k)
S-23     for (i = 0; i < N; i++) {
```

```
S-24        z1[i] = func(x, y, k);
S-25        k++;
S-26     }
S-27
S-28     for (i = 0; i < N; i++) {
S-29        z2 = (double)(i + i*2);
S-30        if (z1[i] != z2) {
S-31           printf("failed\n");
S-32           return 1;
S-33        }
S-34     }
S-35     printf("passed\n");
S-36     return 0;
S-37  }
```

C / C++

Fortran

1 *Example linear_modifier.3.f90* (`omp_4.5`)

```
S-1   module func_mod
S-2   contains
S-3      real(8) function func(x, y, i)
S-4   !$omp declare simd(func) simdlen(4) uniform(x, y) linear(val(i):1)
S-5         implicit none
S-6         real(8), intent(in) :: x(*), y(*)
S-7         integer, intent(in) :: i
S-8
S-9         func = x(i) + y(i)
S-10
S-11     end function func
S-12  end module func_mod
S-13
S-14  program main
S-15     use func_mod
S-16     implicit none
S-17     integer, parameter :: n = 128
S-18     real(8) :: x(n), y(n), z1(n), z2
S-19     integer :: i, k
S-20
S-21     do i=1, n
S-22        x(i) = real(i, kind=8)
S-23        y(i) = real(i*2, kind=8)
S-24     enddo
S-25
S-26     k = 1
S-27  !$omp simd linear(k)
S-28     do i=1, n
```

```
S-29          z1(i) = func(x, y, k)
S-30          k = k + 1
S-31       enddo
S-32
S-33       do i=1, n
S-34          z2 = real(i+i*2, kind=8)
S-35          if (z1(i) /= z2) then
S-36             print *, 'failed'
S-37             stop
S-38          endif
S-39       enddo
S-40       print *, 'passed'
S-41    end program main
```

Fortran

CHAPTER 6

Synchronization

The **barrier** construct is a stand-alone directive that requires all threads of a team (within a contention group) to execute the barrier and complete execution of all tasks within the region, before continuing past the barrier.

The **critical** construct is a directive that contains a structured block. The construct allows only a single thread at a time to execute the structured block (region). Multiple critical regions may exist in a parallel region, and may act cooperatively (only one thread at a time in all **critical** regions), or separately (only one thread at a time in each **critical** regions when a unique name is supplied on each **critical** construct). An optional (lock) **hint** clause may be specified on a named **critical** construct to provide the OpenMP runtime guidance in selection a locking mechanism.

On a finer scale the **atomic** construct allows only a single thread at a time to have atomic access to a storage location involving a single read, write, update or capture statement, and a limited number of combinations when specifying the **capture** *atomic-clause* clause. The *atomic-clause* clause is required for some expression statements, but is not required for **update** statements. The *memory-order* clause can be used to specify the degree of memory ordering enforced by an **atomic** construct. From weakest to strongest, they are **relaxed** (the default), acquire and/or release clauses (specified with **acquire**, **release**, or **acq_rel**), and **seq_cst**. Please see the details in the *atomic Construct* subsection of the *Directives* chapter in the OpenMP Specifications document.

The **ordered** construct either specifies a structured block in a loop, simd, or loop SIMD region that will be executed in the order of the loop iterations. The ordered construct sequentializes and orders the execution of ordered regions while allowing code outside the region to run in parallel.

Since OpenMP 4.5 the **ordered** construct can also be a stand-alone directive that specifies cross-iteration dependences in a doacross loop nest. The **depend** clause uses a **sink** *dependence-type*, along with a iteration vector argument (vec) to indicate the iteration that satisfies the dependence. The **depend** clause with a **source** *dependence-type* specifies dependence satisfaction.

The **flush** directive is a stand-alone construct for enforcing consistency between a thread's view of memory and the view of memory for other threads (see the Memory Model chapter of this document for more details). When the construct is used with an explicit variable list, a *strong flush* that forces a thread's temporary view of memory to be consistent with the actual memory is applied to all listed variables. When the construct is used without an explicit variable list and without a *memory-order* clause, a strong flush is applied to all locally thread-visible data as defined by the base language, and additionally the construct provides both acquire and release memory ordering semantics. When an explicit variable list is not present and a *memory-order* clause is present, the construct provides acquire and/or release memory ordering semantics according to the *memory-order* clause, but no strong flush is performed. A resulting strong flush that applies to a set of variables effectively ensures that no memory (load or store) operation for the affected variables may be reordered across the **flush** directive.

General-purpose routines provide mutual exclusion semantics through locks, represented by lock variables. The semantics allows a task to *set*, and hence *own* a lock, until it is *unset* by the task that set it. A *nestable* lock can be set multiple times by a task, and is used when in code requires nested control of locks. A *simple lock* can only be set once by the owning task. There are specific calls for the two types of locks, and the variable of a specific lock type cannot be used by the other lock type.

Any explicit task will observe the synchronization prescribed in a **barrier** construct and an implied barrier. Also, additional synchronizations are available for tasks. All children of a task will wait at a **taskwait** (for their siblings to complete). A **taskgroup** construct creates a region in which the current task is suspended at the end of the region until all sibling tasks, and their descendants, have completed. Scheduling constraints on task execution can be prescribed by the **depend** clause to enforce dependence on previously generated tasks. More details on controlling task executions can be found in the *Tasking* Chapter in the OpenMP Specifications document.

6.1 The `critical` Construct

The following example includes several **`critical`** constructs. 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`** region. Because the two queues in this example are independent, they are protected by **`critical`** constructs with different names, *xaxis* and *yaxis*.

C / C++

Example critical.1.c

```
S-1   int dequeue(float *a);
S-2   void work(int i, float *a);
S-3
S-4   void critical_example(float *x, float *y)
S-5   {
S-6     int ix_next, iy_next;
S-7
S-8     #pragma omp parallel shared(x, y) private(ix_next, iy_next)
S-9     {
S-10      #pragma omp critical (xaxis)
S-11        ix_next = dequeue(x);
S-12      work(ix_next, x);
S-13
S-14      #pragma omp critical (yaxis)
S-15        iy_next = dequeue(y);
S-16      work(iy_next, y);
S-17    }
S-18
S-19  }
```

C / C++

Fortran

Example critical.1.f

```
S-1         SUBROUTINE CRITICAL_EXAMPLE(X, Y)
S-2
S-3           REAL X(*), Y(*)
S-4           INTEGER IX_NEXT, IY_NEXT
S-5
S-6   !$OMP PARALLEL SHARED(X, Y) PRIVATE(IX_NEXT, IY_NEXT)
S-7
S-8   !$OMP CRITICAL(XAXIS)
S-9           CALL DEQUEUE(IX_NEXT, X)
S-10  !$OMP END CRITICAL(XAXIS)
S-11          CALL WORK(IX_NEXT, X)
S-12
```

```
!$OMP CRITICAL(YAXIS)
        CALL DEQUEUE(IY_NEXT,Y)
!$OMP END CRITICAL(YAXIS)
        CALL WORK(IY_NEXT, Y)

!$OMP END PARALLEL

      END SUBROUTINE CRITICAL_EXAMPLE
```

Fortran

The following example extends the previous example by adding the **hint** clause to the **critical** constructs.

C / C++

Example critical.2.c

```
#include <omp.h>

int dequeue(float *a);
void work(int i, float *a);

void critical_example(float *x, float *y)
{
  int ix_next, iy_next;

  #pragma omp parallel shared(x, y) private(ix_next, iy_next)
  {
    #pragma omp critical (xaxis) hint(omp_lock_hint_contended)
      ix_next = dequeue(x);
    work(ix_next, x);

    #pragma omp critical (yaxis) hint(omp_lock_hint_contended)
      iy_next = dequeue(y);
    work(iy_next, y);
  }

}
```

C / C++

Fortran

Example critical.2.f (`omp_4.5`)

```
S-1         SUBROUTINE CRITICAL_EXAMPLE(X, Y)
S-2           USE OMP_LIB        ! or INCLUDE "omp_lib.h"
S-3
S-4           REAL X(*), Y(*)
S-5           INTEGER IX_NEXT, IY_NEXT
S-6
S-7     !$OMP PARALLEL SHARED(X, Y) PRIVATE(IX_NEXT, IY_NEXT)
S-8
S-9     !$OMP CRITICAL(XAXIS) HINT(OMP_LOCK_HINT_CONTENDED)
S-10          CALL DEQUEUE(IX_NEXT, X)
S-11    !$OMP END CRITICAL(XAXIS)
S-12          CALL WORK(IX_NEXT, X)
S-13
S-14    !$OMP CRITICAL(YAXIS) HINT(OMP_LOCK_HINT_CONTENDED)
S-15          CALL DEQUEUE(IY_NEXT,Y)
S-16    !$OMP END CRITICAL(YAXIS)
S-17          CALL WORK(IY_NEXT, Y)
S-18
S-19    !$OMP END PARALLEL
S-20
S-21        END SUBROUTINE CRITICAL_EXAMPLE
```

Fortran

6.2 Worksharing Constructs Inside a `critical` Construct

The following example demonstrates using a worksharing construct inside a **`critical`** construct. This example is conforming because the worksharing **`single`** region is not closely nested inside the **`critical`** region. A single thread executes the one and only section in the **`sections`** region, and executes the **`critical`** region. The same thread encounters the nested **`parallel`** region, creates a new team of threads, and becomes the master of the new team. One of the threads in the new team enters the **`single`** region and increments **`i`** by **`1`**. At the end of this example **`i`** is equal to **`2`**.

C / C++

Example worksharing_critical.1.c

```
S-1   void critical_work()
S-2   {
S-3     int i = 1;
S-4     #pragma omp parallel sections
S-5     {
S-6       #pragma omp section
S-7       {
S-8         #pragma omp critical (name)
S-9         {
S-10          #pragma omp parallel
S-11          {
S-12            #pragma omp single
S-13            {
S-14              i++;
S-15            }
S-16          }
S-17        }
S-18      }
S-19    }
S-20  }
```

C / C++

Fortran

Example worksharing_critical.1.f

```
      SUBROUTINE CRITICAL_WORK()

        INTEGER I
        I = 1

!$OMP   PARALLEL SECTIONS
!$OMP     SECTION
!$OMP       CRITICAL (NAME)
!$OMP         PARALLEL
!$OMP           SINGLE
                  I = I + 1
!$OMP           END SINGLE
!$OMP         END PARALLEL
!$OMP       END CRITICAL (NAME)
!$OMP   END PARALLEL SECTIONS
      END SUBROUTINE CRITICAL_WORK
```

Fortran

6.3 Binding of `barrier` Regions

The binding rules call for a **`barrier`** region to bind to the closest enclosing **`parallel`** region.

In the following example, the call from the main program to *sub2* is conforming because the **`barrier`** region (in *sub3*) binds to the **`parallel`** region in *sub2*. The call from the main program to *sub1* is conforming because the **`barrier`** region binds to the **`parallel`** region in subroutine *sub2*.

The call from the main program to *sub3* is conforming because the **`barrier`** region binds to the implicit inactive **`parallel`** region enclosing the sequential part. Also note that the **`barrier`** region in *sub3* when called from *sub2* only synchronizes the team of threads in the enclosing **`parallel`** region and not all the threads created in *sub1*.

C / C++

Example barrier_regions.1.c

```
S-1   void work(int n) {}
S-2
S-3   void sub3(int n)
S-4   {
S-5     work(n);
S-6     #pragma omp barrier
S-7     work(n);
S-8   }
S-9
S-10  void sub2(int k)
S-11  {
S-12    #pragma omp parallel shared(k)
S-13      sub3(k);
S-14  }
S-15
S-16  void sub1(int n)
S-17  {
S-18    int i;
S-19    #pragma omp parallel private(i) shared(n)
S-20    {
S-21      #pragma omp for
S-22      for (i=0; i<n; i++)
S-23        sub2(i);
S-24    }
S-25  }
S-26
S-27  int main()
S-28  {
S-29    sub1(2);
S-30    sub2(2);
```

```
S-31      sub3(2);
S-32      return 0;
S-33    }
```

C / C++

Fortran

Example barrier_regions.1.f

```
S-1           SUBROUTINE WORK(N)
S-2             INTEGER N
S-3           END SUBROUTINE WORK
S-4
S-5           SUBROUTINE SUB3(N)
S-6           INTEGER N
S-7             CALL WORK(N)
S-8     !$OMP   BARRIER
S-9             CALL WORK(N)
S-10          END SUBROUTINE SUB3
S-11
S-12          SUBROUTINE SUB2(K)
S-13          INTEGER K
S-14    !$OMP   PARALLEL SHARED(K)
S-15              CALL SUB3(K)
S-16    !$OMP   END PARALLEL
S-17          END SUBROUTINE SUB2
S-18
S-19
S-20          SUBROUTINE SUB1(N)
S-21          INTEGER N
S-22            INTEGER I
S-23    !$OMP   PARALLEL PRIVATE(I) SHARED(N)
S-24    !$OMP     DO
S-25              DO I = 1, N
S-26                CALL SUB2(I)
S-27              END DO
S-28    !$OMP   END PARALLEL
S-29          END SUBROUTINE SUB1
S-30
S-31          PROGRAM EXAMPLE
S-32            CALL SUB1(2)
S-33            CALL SUB2(2)
S-34            CALL SUB3(2)
S-35          END PROGRAM EXAMPLE
```

Fortran

6.4 The `atomic` Construct

The following example avoids race conditions (simultaneous updates of an element of *x* by multiple threads) by using the **`atomic`** construct .

The advantage of using the **`atomic`** construct in this example is that it allows updates of two different elements of *x* to occur in parallel. If a **`critical`** construct were used instead, then all updates to elements of *x* would be executed serially (though not in any guaranteed order).

Note that the **`atomic`** directive applies only to the statement immediately following it. As a result, elements of *y* are not updated atomically in this example.

C / C++

Example atomic.1.c (**`omp_3.1`**)

```
S-1   float work1(int i)
S-2   {
S-3     return 1.0 * i;
S-4   }
S-5
S-6   float work2(int i)
S-7   {
S-8      return 2.0 * i;
S-9   }
S-10
S-11  void atomic_example(float *x, float *y, int *index, int n)
S-12  {
S-13    int i;
S-14
S-15    #pragma omp parallel for shared(x, y, index, n)
S-16      for (i=0; i<n; i++) {
S-17        #pragma omp atomic update
S-18        x[index[i]] += work1(i);
S-19        y[i] += work2(i);
S-20       }
S-21  }
S-22
S-23  int main()
S-24  {
S-25    float x[1000];
S-26    float y[10000];
S-27    int index[10000];
S-28    int i;
S-29
S-30    for (i = 0; i < 10000; i++) {
S-31      index[i] = i % 1000;
S-32      y[i]=0.0;
S-33    }
```

```
S-34      for (i = 0; i < 1000; i++)
S-35        x[i] = 0.0;
S-36      atomic_example(x, y, index, 10000);
S-37      return 0;
S-38    }
```

C / C++

Fortran

1 *Example atomic.1.f* (`omp_3.1`)

```
 S-1          REAL FUNCTION WORK1(I)
 S-2            INTEGER I
 S-3            WORK1 = 1.0 * I
 S-4            RETURN
 S-5          END FUNCTION WORK1
 S-6
 S-7          REAL FUNCTION WORK2(I)
 S-8            INTEGER I
 S-9            WORK2 = 2.0 * I
S-10            RETURN
S-11          END FUNCTION WORK2
S-12
S-13          SUBROUTINE SUB(X, Y, INDEX, N)
S-14            REAL X(*), Y(*)
S-15            INTEGER INDEX(*), N
S-16
S-17            INTEGER I
S-18
S-19    !$OMP   PARALLEL DO SHARED(X, Y, INDEX, N)
S-20              DO I=1,N
S-21    !$OMP       ATOMIC UPDATE
S-22                  X(INDEX(I)) = X(INDEX(I)) + WORK1(I)
S-23              Y(I) = Y(I) + WORK2(I)
S-24              ENDDO
S-25
S-26          END SUBROUTINE SUB
S-27
S-28          PROGRAM ATOMIC_EXAMPLE
S-29            REAL X(1000), Y(10000)
S-30            INTEGER INDEX(10000)
S-31            INTEGER I
S-32
S-33            DO I=1,10000
S-34              INDEX(I) = MOD(I, 1000) + 1
S-35              Y(I) = 0.0
S-36            ENDDO
S-37
```

```
S-38            DO I = 1,1000
S-39              X(I) = 0.0
S-40            ENDDO
S-41
S-42            CALL SUB(X, Y, INDEX, 10000)
S-43
S-44          END PROGRAM ATOMIC_EXAMPLE
```

Fortran

The following example illustrates the **read** and **write** clauses for the **atomic** directive. These clauses ensure that the given variable is read or written, respectively, as a whole. Otherwise, some other thread might read or write part of the variable while the current thread was reading or writing another part of the variable. Note that most hardware provides atomic reads and writes for some set of properly aligned variables of specific sizes, but not necessarily for all the variable types supported by the OpenMP API.

C / C++

Example atomic.2.c (**omp_3.1**)

```
S-1   int atomic_read(const int *p)
S-2   {
S-3       int value;
S-4   /* Guarantee that the entire value of *p is read atomically. No part of
S-5    * *p can change during the read operation.
S-6    */
S-7   #pragma omp atomic read
S-8        value = *p;
S-9        return value;
S-10  }
S-11
S-12  void atomic_write(int *p, int value)
S-13  {
S-14  /* Guarantee that value is stored atomically into *p. No part of *p can
S-15  change
S-16   * until after the entire write operation is completed.
S-17   */
S-18  #pragma omp atomic write
S-19      *p = value;
S-20  }
```

C / C++

Fortran

Example atomic.2.f (**omp_3.1**)

```
S-1         function atomic_read(p)
S-2         integer :: atomic_read
S-3         integer, intent(in) :: p
S-4   ! Guarantee that the entire value of p is read atomically. No part of
S-5   ! p can change during the read operation.
S-6
S-7   !$omp atomic read
S-8         atomic_read = p
S-9         return
S-10        end function atomic_read
S-11
S-12        subroutine atomic_write(p, value)
S-13        integer, intent(out) :: p
S-14        integer, intent(in) :: value
S-15  ! Guarantee that value is stored atomically into p. No part of p can change
S-16  ! until after the entire write operation is completed.
S-17  !$omp atomic write
S-18        p = value
S-19        end subroutine atomic_write
```

Fortran

The following example illustrates the **capture** clause for the **atomic** directive. In this case the value of a variable is captured, and then the variable is incremented. These operations occur atomically. This particular example could be implemented using the fetch-and-add instruction available on many kinds of hardware. The example also shows a way to implement a spin lock using the **capture** and **read** clauses.

C / C++

Example atomic.3.c (**omp_3.1**)

```
S-1   int fetch_and_add(int *p)
S-2   {
S-3   /* Atomically read the value of *p and then increment it. The previous value
S-4   is
S-5    * returned. This can be used to implement a simple lock as shown below.
S-6    */
S-7       int old;
S-8   #pragma omp atomic capture
S-9       { old = *p; (*p)++; }
S-10      return old;
S-11  }
S-12
S-13  /*
```

```
 * Use fetch_and_add to implement a lock
 */
struct locktype {
    int ticketnumber;
    int turn;
};
void do_locked_work(struct locktype *lock)
{
    int atomic_read(const int *p);
    void work();

    // Obtain the lock
    int myturn = fetch_and_add(&lock->ticketnumber);
    while (atomic_read(&lock->turn) != myturn)
       ;
    // Do some work. The flush is needed to ensure visibility of
    // variables not involved in atomic directives

#pragma omp flush
    work();
#pragma omp flush
    // Release the lock
    fetch_and_add(&lock->turn);
}
```

C / C++

Fortran

Example atomic.3.f (`omp_3.1`)

```
      function fetch_and_add(p)
      integer:: fetch_and_add
      integer, intent(inout) :: p

! Atomically read the value of p and then increment it. The previous value is
! returned. This can be used to implement a simple lock as shown below.
!$omp atomic capture
      fetch_and_add = p
      p = p + 1
!$omp end atomic
      end function fetch_and_add
      module m
      interface
        function fetch_and_add(p)
          integer :: fetch_and_add
          integer, intent(inout) :: p
        end function
        function atomic_read(p)
```

```
S-19              integer :: atomic_read
S-20              integer, intent(in) :: p
S-21            end function
S-22          end interface
S-23          type locktype
S-24             integer ticketnumber
S-25             integer turn
S-26          end type
S-27          contains
S-28          subroutine do_locked_work(lock)
S-29          type(locktype), intent(inout) :: lock
S-30          integer myturn
S-31          integer junk
S-32    ! obtain the lock
S-33           myturn = fetch_and_add(lock%ticketnumber)
S-34           do while (atomic_read(lock%turn) .ne. myturn)
S-35             continue
S-36           enddo
S-37    ! Do some work. The flush is needed to ensure visibility of variables
S-38    ! not involved in atomic directives
S-39    !$omp flush
S-40          call work
S-41    !$omp flush
S-42    ! Release the lock
S-43          junk = fetch_and_add(lock%turn)
S-44          end subroutine
S-45          end module
```

Fortran

6.5 Restrictions on the `atomic` Construct

The following non-conforming examples illustrate the restrictions on the **atomic** construct.

C / C++

Example atomic_restrict.1.c (**omp_3.1**)

```
S-1   void atomic_wrong ()
S-2   {
S-3    union {int n; float x;} u;
S-4
S-5   #pragma omp parallel
S-6     {
S-7   #pragma omp atomic update
S-8       u.n++;
S-9
S-10  #pragma omp atomic update
S-11      u.x += 1.0;
S-12
S-13  /* Incorrect because the atomic constructs reference the same location
S-14     through incompatible types */
S-15    }
S-16  }
```

C / C++

Fortran

Example atomic_restrict.1.f (**omp_3.1**)

```
S-1         SUBROUTINE ATOMIC_WRONG()
S-2           INTEGER:: I
S-3           REAL:: R
S-4           EQUIVALENCE(I,R)
S-5
S-6   !$OMP   PARALLEL
S-7   !$OMP     ATOMIC UPDATE
S-8               I = I + 1
S-9   !$OMP     ATOMIC UPDATE
S-10              R = R + 1.0
S-11  ! incorrect because I and R reference the same location
S-12  ! but have different types
S-13  !$OMP   END PARALLEL
S-14        END SUBROUTINE ATOMIC_WRONG
```

Fortran

C / C++

Example atomic_restrict.2.c (**omp_3.1**)

```
S-1   void atomic_wrong2 ()
S-2   {
S-3    int  x;
S-4    int *i;
S-5    float   *r;
S-6
S-7    i = &x;
S-8    r = (float *)&x;
S-9
S-10  #pragma omp parallel
S-11    {
S-12  #pragma omp atomic update
S-13      *i += 1;
S-14
S-15  #pragma omp atomic update
S-16      *r += 1.0;
S-17
S-18  /* Incorrect because the atomic constructs reference the same location
S-19     through incompatible types */
S-20
S-21    }
S-22  }
```

C / C++

Fortran

The following example is non-conforming because **I** and **R** reference the same location but have different types.

Example atomic_restrict.2.f (**omp_3.1**)

```
S-1         SUBROUTINE SUB()
S-2           COMMON /BLK/ R
S-3           REAL R
S-4
S-5   !$OMP   ATOMIC UPDATE
S-6             R = R + 1.0
S-7         END SUBROUTINE SUB
S-8
S-9         SUBROUTINE ATOMIC_WRONG2()
S-10          COMMON /BLK/ I
S-11          INTEGER I
S-12
S-13  !$OMP   PARALLEL
S-14
```

```
!$OMP     ATOMIC UPDATE
            I = I + 1
          CALL SUB()
!$OMP   END PARALLEL
      END SUBROUTINE ATOMIC_WRONG2
```

Although the following example might work on some implementations, this is also non-conforming:

Example atomic_restrict.3.f (`omp_3.1`)

```
      SUBROUTINE ATOMIC_WRONG3
        INTEGER:: I
        REAL:: R
        EQUIVALENCE(I,R)

!$OMP   PARALLEL
!$OMP     ATOMIC UPDATE
            I = I + 1
! incorrect because I and R reference the same location
! but have different types
!$OMP   END PARALLEL

!$OMP   PARALLEL
!$OMP     ATOMIC UPDATE
            R = R + 1.0
! incorrect because I and R reference the same location
! but have different types
!$OMP   END PARALLEL

      END SUBROUTINE ATOMIC_WRONG3
```

Fortran

6.6 The `flush` Construct without a List

The following example distinguishes the shared variables affected by a **flush** construct with no list from the shared objects that are not affected:

C / C++

Example flush_nolist.1.c

```
S-1   int x, *p = &x;
S-2
S-3   void f1(int *q)
S-4   {
S-5     *q = 1;
S-6     #pragma omp flush
S-7     /* x, p, and *q are flushed */
S-8     /* because they are shared and accessible */
S-9     /* q is not flushed because it is not shared. */
S-10  }
S-11
S-12  void f2(int *q)
S-13  {
S-14    #pragma omp barrier
S-15    *q = 2;
S-16    #pragma omp barrier
S-17
S-18    /* a barrier implies a flush */
S-19    /* x, p, and *q are flushed */
S-20    /* because they are shared and accessible */
S-21    /* q is not flushed because it is not shared. */
S-22  }
S-23
S-24  int g(int n)
S-25  {
S-26    int i = 1, j, sum = 0;
S-27    *p = 1;
S-28    #pragma omp parallel reduction(+: sum) num_threads(10)
S-29    {
S-30      f1(&j);
S-31
S-32      /* i, n and sum were not flushed */
S-33      /* because they were not accessible in f1 */
S-34      /* j was flushed because it was accessible */
S-35      sum += j;
S-36
S-37      f2(&j);
S-38
S-39      /* i, n, and sum were not flushed */
```

```
S-40        /* because they were not accessible in f2 */
S-41        /* j was flushed because it was accessible */
S-42        sum += i + j + *p + n;
S-43      }
S-44      return sum;
S-45    }
S-46
S-47    int main()
S-48    {
S-49      int result = g(7);
S-50      return result;
S-51    }
```

C / C++

Fortran

1 *Example flush_nolist.1.f*

```
S-1           SUBROUTINE F1(Q)
S-2             COMMON /DATA/ X, P
S-3             INTEGER, TARGET  :: X
S-4             INTEGER, POINTER :: P
S-5             INTEGER Q
S-6
S-7             Q = 1
S-8     !$OMP   FLUSH
S-9             ! X, P and Q are flushed
S-10            ! because they are shared and accessible
S-11          END SUBROUTINE F1
S-12
S-13          SUBROUTINE F2(Q)
S-14            COMMON /DATA/ X, P
S-15            INTEGER, TARGET  :: X
S-16            INTEGER, POINTER :: P
S-17            INTEGER Q
S-18
S-19    !$OMP   BARRIER
S-20              Q = 2
S-21    !$OMP   BARRIER
S-22              ! a barrier implies a flush
S-23              ! X, P and Q are flushed
S-24              ! because they are shared and accessible
S-25          END SUBROUTINE F2
S-26
S-27          INTEGER FUNCTION G(N)
S-28            COMMON /DATA/ X, P
S-29            INTEGER, TARGET  :: X
S-30            INTEGER, POINTER :: P
```

```
S-31            INTEGER N
S-32            INTEGER I, J, SUM
S-33
S-34            I = 1
S-35            SUM = 0
S-36            P = 1
S-37    !$OMP   PARALLEL REDUCTION(+: SUM) NUM_THREADS(10)
S-38              CALL F1(J)
S-39                ! I, N and SUM were not flushed
S-40                !   because they were not accessible in F1
S-41                ! J was flushed because it was accessible
S-42              SUM = SUM + J
S-43
S-44              CALL F2(J)
S-45                ! I, N, and SUM were not flushed
S-46                !   because they were not accessible in f2
S-47                ! J was flushed because it was accessible
S-48              SUM = SUM + I + J + P + N
S-49    !$OMP   END PARALLEL
S-50
S-51            G = SUM
S-52          END FUNCTION G
S-53
S-54          PROGRAM FLUSH_NOLIST
S-55            COMMON /DATA/ X, P
S-56            INTEGER, TARGET  :: X
S-57            INTEGER, POINTER :: P
S-58            INTEGER RESULT, G
S-59
S-60            P => X
S-61            RESULT = G(7)
S-62            PRINT *, RESULT
S-63          END PROGRAM FLUSH_NOLIST
```

Fortran

6.7 Synchronization Based on Acquire/Release Semantics

As explained in the Memory Model chapter of this document, a flush operation may be an *acquire flush* and/or a *release flush*, and OpenMP 5.0 defines acquire/release semantics in terms of these fundamental flush operations. For any synchronization between two threads that is specified by OpenMP, a release flush logically occurs at the source of the synchronization and an acquire flush logically occurs at the sink of the synchronization. OpenMP 5.0 added memory ordering clauses – **`acquire`**, **`release`**, and **`acq_rel`** – to the **`flush`** and **`atomic`** constructs for explicitly requesting acquire/release semantics. Furthermore, implicit flushes for all OpenMP constructs and runtime routines that synchronize OpenMP threads in some manner were redefined in terms of synchronizing release and acquire flushes to avoid the requirement of strong memory fences (see the *Flush Synchronization and Happens Before* and *Implicit Flushes* sections of the OpenMP Specifications document).

The examples that follow in this section illustrate how acquire and release flushes may be employed, implicitly or explicitly, for synchronizing threads. A **`flush`** directive without a list and without any memory ordering clause can also function as both an acquire and release flush for facilitating thread synchronization. Flushes implied on entry to, or exit from, an atomic operation (specified by an **`atomic`** construct) may function as an acquire flush or a release flush if a memory ordering clause appears on the construct. On entry to and exit from a **`critical`** construct there is now an implicit acquire flush and release flush, respectively.

The first example illustrates how the release and acquire flushes implied by a **`critical`** region guarantee a value written by the first thread is visible to a read of the value on the second thread. Thread 0 writes to *x* and then executes a **`critical`** region in which it writes to *y*; the write to *x* happens before the execution of the **`critical`** region, consistent with the program order of the thread. Meanwhile, thread 1 executes a **`critical`** region in a loop until it reads a non-zero value from *y* in the **`critical`** region, after which it prints the value of *x*; again, the execution of the **`critical`** regions happen before the read from *x* based on the program order of the thread. The **`critical`** regions executed by the two threads execute in a serial manner, with a pair-wise synchronization from the exit of one **`critical`** region to the entry to the next **`critical`** region. These pair-wise synchronizations result from the implicit release flushes that occur on exit from **`critical`** regions and the implicit acquire flushes that occur on entry to **`critical`** regions; hence, the execution of each **`critical`** region in the sequence happens before the execution of the next **`critical`** region. A "happens before" order is therefore established between the assignment to *x* by thread 0 and the read from *x* by thread 1, and so thread 1 must see that *x* equals 10.

C / C++

Example acquire_release.1.c (`omp_5.0`)

```
#include <stdio.h>
#include <omp.h>

int main()
{
   int x = 0, y = 0;
   #pragma omp parallel num_threads(2)
   {
      int thrd = omp_get_thread_num();
      if (thrd == 0) {
         x = 10;
         #pragma omp critical
         { y = 1; }
      } else {
         int tmp = 0;
         while (tmp == 0) {
           #pragma omp critical
           { tmp = y; }
         }
         printf("x = %d\n", x); // always "x = 10"
      }
   }
   return 0;
}
```

C / C++

Fortran

Example acquire_release.1.f90 (`omp_5.0`)

```
program rel_acq_ex1
   use omp_lib
   integer :: x, y, thrd, tmp
   x = 0
   y = 0
   !$omp parallel num_threads(2) private(thrd, tmp)
      thrd = omp_get_thread_num()
      if (thrd == 0) then
         x = 10
         !$omp critical
         y = 1
         !$omp end critical
      else
```

```
            tmp = 0
            do while (tmp == 0)
               !$omp critical
               tmp = y
               !$omp end critical
            end do
            print *, "x = ", x  !! always "x = 10"
         end if
      !$omp end parallel
   end program
```

Fortran

In the second example, the **critical** constructs are exchanged with **atomic** constructs that have *explicit* memory ordering specified. When the atomic read operation on thread 1 reads a non-zero value from *y*, this results in a release/acquire synchronization that in turn implies that the assignment to *x* on thread 0 happens before the read of *x* on thread 1. Therefore, thread 1 will print "x = 10".

C / C++

Example acquire_release.2.c (**omp_5.0**)

```

#include <stdio.h>
#include <omp.h>

int main()
{
   int x = 0, y = 0;
   #pragma omp parallel num_threads(2)
   {
      int thrd = omp_get_thread_num();
       if (thrd == 0) {
          x = 10;
          #pragma omp atomic write release // or seq_cst
          y = 1;
       } else {
          int tmp = 0;
          while (tmp == 0) {
            #pragma omp atomic read acquire // or seq_cst
            tmp = y;
          }
          printf("x = %d\n", x); // always "x = 10"
       }
   }
   return 0;
}
```

C / C++

Fortran

Example acquire_release.2.f90 (`omp_5.0`)

```
S-1
S-2     program rel_acq_ex2
S-3        use omp_lib
S-4        integer :: x, y, thrd, tmp
S-5        x = 0
S-6        y = 0
S-7        !$omp parallel num_threads(2) private(thrd, tmp)
S-8           thrd = omp_get_thread_num()
S-9           if (thrd == 0) then
S-10             x = 10
S-11             !$omp atomic write release ! or seq_cst
S-12              y = 1
S-13             !$omp end atomic
S-14          else
S-15             tmp = 0
S-16             do while (tmp == 0)
S-17                !$omp atomic read acquire ! or seq_cst
S-18                tmp = y
S-19                !$omp end atomic
S-20             end do
S-21             print *, "x = ", x  !! always "x = 10"
S-22          end if
S-23       !$omp end parallel
S-24    end program
```

Fortran

In the third example, **atomic** constructs that specify relaxed atomic operations are used with explicit **flush** directives to enforce memory ordering between the two threads. The explicit **flush** directive on thread 0 must specify a release flush and the explicit **flush** directive on thread 1 must specify an acquire flush to establish a release/acquire synchronization between the two threads. The **flush** and **atomic** constructs encountered by thread 0 can be replaced by the **atomic** construct used in Example 2 for thread 0, and similarly the **flush** and **atomic** constructs encountered by thread 1 can be replaced by the **atomic** construct used in Example 2 for thread 1.

C / C++

Example acquire_release.3.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #include <omp.h>
S-4
S-5   int main()
S-6   {
S-7      int x = 0, y = 0;
S-8      #pragma omp parallel num_threads(2)
S-9      {
S-10        int thrd = omp_get_thread_num();
S-11        if (thrd == 0) {
S-12           x = 10;
S-13           #pragma omp flush // or with acq_rel or release clause
S-14           #pragma omp atomic write // or with relaxed clause
S-15           y = 1;
S-16        } else {
S-17           int tmp = 0;
S-18           while (tmp == 0) {
S-19             #pragma omp atomic read // or with relaxed clause
S-20             tmp = y;
S-21           }
S-22           #pragma omp flush // or with acq_rel or acquire clause
S-23           printf("x = %d\n", x);  // always "x = 10"
S-24        }
S-25     }
S-26     return 0;
S-27  }
```

C / C++

Fortran

Example acquire_release.3.f90 (**omp_5.0**)

```
S-1
S-2   program rel_acq_ex3
S-3      use omp_lib
S-4      integer :: x, y, thrd, tmp
S-5      x = 0
S-6      y = 0
S-7      !$omp parallel num_threads(2) private(thrd, tmp)
S-8         thrd = omp_get_thread_num()
S-9         if (thrd == 0) then
S-10           x = 10
S-11           !$omp flush ! or with acq_rel or release clause
S-12           !$omp atomic write
S-13           y = 1
S-14           !$omp end atomic
S-15        else
S-16           tmp = 0
S-17           do while (tmp == 0)
S-18              !$omp atomic read
S-19              tmp = y
S-20              !$omp end atomic
S-21           end do
S-22           !$omp flush ! or with acq_rel or acquire clause
S-23           print *, "x = ", x  !! always "x = 10"
S-24        end if
S-25     !$omp end parallel
S-26  end program
```

Fortran

Example 4 will fail to order the write to *x* on thread 0 before the read from *x* on thread 1. Importantly, the implicit release flush on exit from the **critical** region will not synchronize with the acquire flush that occurs on the atomic read operation performed by thread 1. This is because implicit release flushes that occur on a given construct may only synchronize with implicit acquire flushes on a compatible construct (and vice-versa) that internally makes use of the same synchronization variable. For a **critical** construct, this might correspond to a *lock* object that is used by a given implementation (for the synchronization semantics of other constructs due to implicit release and acquire flushes, refer to the *Implicit Flushes* section of the OpenMP Specifications document). Either an explicit **flush** directive that provides a release flush (i.e., a flush without a list that does not have the **acquire** clause) must be specified between the **critical** construct and the atomic write, or an atomic operation that modifies *y* and provides release semantics must be specified.

C / C++

Example acquire_release_broke.4.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #include <omp.h>
S-4
S-5   int main()
S-6   {
S-7
S-8   // !!! THIS CODE WILL FAIL TO PRODUCE CONSISTENT RESULTS !!!!!!!
S-9   // !!! DO NOT PROGRAM SYNCHRONIZATION THIS WAY !!!!!!!
S-10
S-11     int x = 0, y;
S-12     #pragma omp parallel num_threads(2)
S-13     {
S-14        int thrd = omp_get_thread_num();
S-15        if (thrd == 0) {
S-16           #pragma omp critical
S-17           { x = 10; }
S-18           // an explicit flush directive that provides
S-19           // release semantics is needed here
S-20           // to complete the synchronization.
S-21           #pragma omp atomic write
S-22           y = 1;
S-23        } else {
S-24           int tmp = 0;
S-25           while (tmp == 0) {
S-26              #pragma omp atomic read acquire // or seq_cst
S-27              tmp = y;
S-28           }
S-29           #pragma omp critical
S-30           { printf("x = %d\n", x); }  // !! NOT ALWAYS 10
S-31        }
S-32     }
S-33     return 0;
S-34  }
```

C / C++

Fortran

Example acquire_release_broke.4.f90 (`omp_5.0`)

```
S-1
S-2     program rel_acq_ex4
S-3        use omp_lib
S-4        integer :: x, y, thrd
S-5        integer :: tmp
S-6        x = 0
S-7
S-8     !! !!! THIS CODE WILL FAIL TO PRODUCE CONSISTENT RESULTS !!!!!!!
S-9     !! !!! DO NOT PROGRAM SYNCHRONIZATION THIS WAY !!!!!!!
S-10
S-11       !$omp parallel num_threads(2) private(thrd) private(tmp)
S-12          thrd = omp_get_thread_num()
S-13          if (thrd == 0) then
S-14             !$omp critical
S-15             x = 10
S-16             !$omp end critical
S-17             ! an explicit flush directive that provides
S-18             ! release semantics is needed here to
S-19             ! complete the synchronization.
S-20             !$omp atomic write
S-21             y = 1
S-22             !$omp end atomic
S-23          else
S-24             tmp = 0
S-25             do while(tmp == 0)
S-26               !$omp atomic read acquire ! or seq_cst
S-27               tmp = x
S-28               !$omp end atomic
S-29             end do
S-30             !$omp critical
S-31             print *, "x = ", x  !! !! NOT ALWAYS 10
S-32             !$omp end critical
S-33          end if
S-34       !$omp end parallel
S-35    end program
```

Fortran

6.8 The `ordered` Clause and the `ordered` Construct

Ordered constructs are useful for sequentially ordering the output from work that is done in parallel. The following program prints out the indices in sequential order:

C / C++

Example ordered.1.c

```
S-1   #include <stdio.h>
S-2
S-3   void work(int k)
S-4   {
S-5     #pragma omp ordered
S-6       printf(" %d\n", k);
S-7   }
S-8
S-9   void ordered_example(int lb, int ub, int stride)
S-10  {
S-11    int i;
S-12
S-13    #pragma omp parallel for ordered schedule(dynamic)
S-14    for (i=lb; i<ub; i+=stride)
S-15      work(i);
S-16  }
S-17
S-18  int main()
S-19  {
S-20    ordered_example(0, 100, 5);
S-21    return 0;
S-22  }
```

C / C++

Fortran

Example ordered.1.f

```
S-1         SUBROUTINE WORK(K)
S-2           INTEGER k
S-3
S-4   !$OMP ORDERED
S-5           WRITE(*,*) K
S-6   !$OMP END ORDERED
S-7
S-8         END SUBROUTINE WORK
S-9
S-10        SUBROUTINE SUB(LB, UB, STRIDE)
S-11          INTEGER LB, UB, STRIDE
S-12          INTEGER I
```

```
!$OMP PARALLEL DO ORDERED SCHEDULE(DYNAMIC)
        DO I=LB,UB,STRIDE
          CALL WORK(I)
        END DO
!$OMP END PARALLEL DO

      END SUBROUTINE SUB

      PROGRAM ORDERED_EXAMPLE
        CALL SUB(1,100,5)
      END PROGRAM ORDERED_EXAMPLE
```

Fortran

It is possible to have multiple **ordered** constructs within a loop region with the **ordered** clause specified. The first example is non-conforming because all iterations execute two **ordered** regions. An iteration of a loop must not execute more than one **ordered** region:

C / C++

Example ordered.2.c

```
void work(int i) {}

void ordered_wrong(int n)
{
  int i;
  #pragma omp for ordered
  for (i=0; i<n; i++) {
/* incorrect because an iteration may not execute more than one
   ordered region */
    #pragma omp ordered
      work(i);
    #pragma omp ordered
      work(i+1);
  }
}
```

C / C++

Fortran

Example ordered.2.f

```
      SUBROUTINE WORK(I)
      INTEGER I
      END SUBROUTINE WORK

      SUBROUTINE ORDERED_WRONG(N)
      INTEGER N

        INTEGER I
!$OMP   DO ORDERED
        DO I = 1, N
! incorrect because an iteration may not execute more than one
! ordered region
!$OMP     ORDERED
            CALL WORK(I)
!$OMP     END ORDERED

!$OMP     ORDERED
            CALL WORK(I+1)
!$OMP     END ORDERED
        END DO
      END SUBROUTINE ORDERED_WRONG
```

Fortran

The following is a conforming example with more than one **ordered** construct. Each iteration will execute only one **ordered** region:

C / C++

Example ordered.3.c

```
void work(int i) {}
void ordered_good(int n)
{
  int i;
#pragma omp for ordered
  for (i=0; i<n; i++) {
    if (i <= 10) {
      #pragma omp ordered
         work(i);
    }
    if (i > 10) {
      #pragma omp ordered
        work(i+1);
    }
```

```
S-15      }
S-16   }
```

C / C++

Fortran

Example ordered.3.f

```
S-1          SUBROUTINE ORDERED_GOOD(N)
S-2          INTEGER N
S-3
S-4    !$OMP   DO ORDERED
S-5            DO I = 1,N
S-6              IF (I <= 10) THEN
S-7    !$OMP       ORDERED
S-8                  CALL WORK(I)
S-9    !$OMP       END ORDERED
S-10             ENDIF
S-11
S-12             IF (I > 10) THEN
S-13   !$OMP       ORDERED
S-14                 CALL WORK(I+1)
S-15   !$OMP       END ORDERED
S-16             ENDIF
S-17           ENDDO
S-18         END SUBROUTINE ORDERED_GOOD
```

Fortran

6.9 The `depobj` Construct

The stand-alone **depobj** construct provides a mechanism to create a *depend object* that expresses a dependence to be used subsequently in the **depend** clause of another construct. The dependence is created from a dependence type and a storage location, within a **depend** clause of an **depobj** construct; and it is stored in the depend object. The depend object is represented by a variable of type **omp_depend_t** in C/C++ (by a scalar variable of integer kind **omp_depend_kind** in Fortran).

In the example below the stand-alone **depobj** construct uses the **depend**, **update** and **destroy** clauses to *initialize*, *update* and *uninitialize* a depend object (**obj**).

The first **depobj** construct initializes the **obj** depend object with an **inout** dependence type with a storage location defined by variable **a**. This dependence is passed into the *driver* routine via the **obj** depend object.

In the first *driver* routine call, *Task 1* uses the dependence of the object (**inout**), while *Task 2* uses an **in** dependence specified directly in a **depend** clause. For these task dependences *Task 1* must execute and complete before *Task 2* begins.

Before the second call to *driver*, **obj** is updated using the **depobj** construct to represent an **in** dependence. Hence, in the second call to *driver*, *Task 1* will have an **in** dependence; and *Task 1* and *Task 2* can execute simultaneously. Note: in an **update** clause, only the dependence type can be (is) updated.

The third **depobj** construct uses the **destroy** clause. It frees resources as it puts the depend object in an uninitialized state– effectively destroying the depend object. After an object has been uninitialized it can be initialized again with a new dependence type *and* a new variable.

C / C++

Example depobj.1.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #include <omp.h>
S-4
S-5   #define N 100
S-6   #define TRUE  1
S-7   #define FALSE 0
S-8
S-9   void driver(int update, float a[], float b[], int n, omp_depend_t *obj);
S-10
S-11  void update_copy(int update, float a[], float b[], int n);
S-12  void checkpoint(float a[],int n);
S-13  void init(float a[], int n);
S-14
S-15
S-16  int main(){
```

```
S-17
S-18        float a[N],b[N];
S-19        omp_depend_t obj;
S-20
S-21        init(a, N);
S-22
S-23        #pragma omp depobj(obj) depend(inout: a)
S-24
S-25        driver(TRUE,  a,b,N, &obj);  // updating a occurs
S-26
S-27        #pragma omp depobj(obj) update(in)
S-28
S-29        driver(FALSE, a,b,N, &obj);  // no updating of a
S-30
S-31        #pragma omp depobj(obj) destroy  // obj is set to uninitilized state,
S-32                                         // resources are freed
S-33        return 0;
S-34
S-35     }
S-36
S-37     void driver(int update, float a[], float b[], int n, omp_depend_t *obj)
S-38     {
S-39        #pragma omp parallel num_threads(2)
S-40        #pragma omp single
S-41        {
S-42
S-43          #pragma omp task depend(depobj: *obj) // Task 1, uses depend object
S-44            update_copy(update, a,b,n);       // update a or not, always copy a to b
S-45
S-46          #pragma omp task depend(in: a[:n])    // Task 2, only read a
S-47            checkpoint(a,n);
S-48
S-49        }
S-50     }
S-51
S-52     void update_copy(int update, float a[], float b[], int n)
S-53     {
S-54        if(update) for(int i=0;i<n;i++) a[i]+=1.0f;
S-55
S-56        for(int i=0;i<n;i++) b[i]=a[i];
S-57     }
S-58
S-59     void checkpoint(float a[], int n)
S-60     {
S-61        for(int i=0;i<n;i++) printf(" %f ",a[i]);
S-62        printf("\n");
S-63     }
```

```
S-64
S-65    void init(float a[], int n)
S-66    {
S-67       for(int i=0;i<n;i++) a[i]=i;
S-68    }
S-69
```

C / C++

Fortran

1 *Example depobj.1.f90* (`omp_5.0`)

```
 S-1
 S-2    program main
 S-3        use omp_lib
 S-4        implicit none
 S-5
 S-6        integer,parameter         :: N=100
 S-7        real                      :: a(N),b(N)
 S-8        integer(omp_depend_kind)  :: obj
 S-9
S-10        call init(a, N)
S-11
S-12        !$omp depobj(obj) depend(inout: a)
S-13
S-14        call driver(.true.,  a,b,N, obj)  !! updating occurs
S-15
S-16        !$omp depobj(obj) update(in)
S-17
S-18        call driver(.false., a,b,N, obj)  !! no updating
S-19
S-20        !$omp depobj(obj) destroy    !! obj is set to uninitilized state,
S-21                                     !! resources are freed
S-22
S-23    end program
S-24
S-25    subroutine driver(update, a, b, n, obj)
S-26       use omp_lib
S-27       implicit none
S-28       logical :: update
S-29       real    :: a(n), b(n)
S-30       integer :: n
S-31       integer(omp_depend_kind) :: obj
S-32
S-33       !$omp parallel num_threads(2)
S-34
S-35         !$omp single
S-36
```

```
S-37          !$omp task depend(depobj: obj)       !! Task 1, uses depend object
S-38            call update_copy(update, a,b,n)  !! update a or not, always copy a to b
S-39          !$omp end task
S-40
S-41          !$omp task depend(in: a)             !! Task 2, only read a
S-42            call checkpoint(a,n)
S-43          !$omp end task
S-44
S-45        !$omp end single
S-46
S-47      !$omp end parallel
S-48
S-49    end subroutine
S-50
S-51    subroutine update_copy(update, a, b, n)
S-52       implicit none
S-53       logical :: update
S-54       real    :: a(n), b(n)
S-55       integer :: n
S-56
S-57       if (update) a = a + 1.0
S-58
S-59       b = a
S-60
S-61    end subroutine
S-62
S-63    subroutine checkpoint( a, n)
S-64       implicit none
S-65       integer :: n
S-66       real    :: a(n)
S-67       integer :: i
S-68
S-69       write(*,'( *(f5.0) )') (a(i), i=1,n)
S-70    end subroutine
S-71
S-72    subroutine init(a,n)
S-73       implicit none
S-74       integer :: n
S-75       real    :: a(n)
S-76       integer :: i
S-77
S-78       a=[ (i, i=1,n) ]
S-79    end subroutine
```

Fortran

6.10 Doacross Loop Nest

An **ordered** clause can be used on a loop construct with an integer parameter argument to define the number of associated loops within a *doacross loop nest* where cross-iteration dependences exist. A **depend** clause on an **ordered** construct within an ordered loop describes the dependences of the *doacross* loops.

In the code below, the **depend(sink:i-1)** clause defines an *i-1* to *i* cross-iteration dependence that specifies a wait point for the completion of computation from iteration *i-1* before proceeding to the subsequent statements. The **depend(source)** clause indicates the completion of computation from the current iteration (*i*) to satisfy the cross-iteration dependence that arises from the iteration. For this example the same sequential ordering could have been achieved with an **ordered** clause without a parameter, on the loop directive, and a single **ordered** directive without the **depend** clause specified for the statement executing the *bar* function.

C / C++

Example doacross.1.c (**omp_4.5**)

```
S-1   float foo(int i);
S-2   float bar(float a, float b);
S-3   float baz(float b);
S-4
S-5   void work( int N, float *A, float *B, float *C )
S-6   {
S-7     int i;
S-8
S-9     #pragma omp for ordered(1)
S-10    for (i=1; i<N; i++)
S-11    {
S-12      A[i] = foo(i);
S-13
S-14    #pragma omp ordered depend(sink: i-1)
S-15      B[i] = bar(A[i], B[i-1]);
S-16    #pragma omp ordered depend(source)
S-17
S-18      C[i] = baz(B[i]);
S-19    }
S-20  }
```

C / C++

Fortran

Example doacross.1.f90 (**omp_4.5**)

```
S-1   subroutine work( N, A, B, C )
S-2     integer :: N, i
S-3     real, dimension(N) :: A, B, C
S-4     real, external :: foo, bar, baz
S-5
S-6     !$omp do ordered(1)
S-7     do i=2, N
S-8       A(i) = foo(i)
S-9
S-10    !$omp ordered depend(sink: i-1)
S-11      B(i) = bar(A(i), B(i-1))
S-12    !$omp ordered depend(source)
S-13
S-14      C(i) = baz(B(i))
S-15    end do
S-16  end subroutine
```

Fortran

The following code is similar to the previous example but with *doacross loop nest* extended to two nested loops, *i* and *j*, as specified by the **ordered(2)** clause on the loop directive. In the C/C++ code, the *i* and *j* loops are the first and second associated loops, respectively, whereas in the Fortran code, the *j* and *i* loops are the first and second associated loops, respectively. The **depend(sink:i-1,j)** and **depend(sink:i,j-1)** clauses in the C/C++ code define cross-iteration dependences in two dimensions from iterations (*i-1, j*) and (*i, j-1*) to iteration (*i, j*). Likewise, the **depend(sink:j-1,i)** and **depend(sink:j,i-1)** clauses in the Fortran code define cross-iteration dependences from iterations (*j-1, i*) and (*j, i-1*) to iteration (*j, i*).

C / C++

Example doacross.2.c (**omp_4.5**)

```
S-1   float foo(int i, int j);
S-2   float bar(float a, float b, float c);
S-3   float baz(float b);
S-4
S-5   void work( int N, int M, float **A, float **B, float **C )
S-6   {
S-7     int i, j;
S-8
S-9     #pragma omp for ordered(2)
S-10    for (i=1; i<N; i++)
S-11    {
S-12      for (j=1; j<M; j++)
S-13      {
```

```
S-14          A[i][j] = foo(i, j);
S-15
S-16        #pragma omp ordered depend(sink: i-1,j) depend(sink: i,j-1)
S-17          B[i][j] = bar(A[i][j], B[i-1][j], B[i][j-1]);
S-18        #pragma omp ordered depend(source)
S-19
S-20          C[i][j] = baz(B[i][j]);
S-21        }
S-22      }
S-23    }
```

C / C++

Fortran

Example doacross.2.f90 (**omp_4.5**)

```
S-1     subroutine work( N, M, A, B, C )
S-2       integer :: N, M, i, j
S-3       real, dimension(M,N) :: A, B, C
S-4       real, external :: foo, bar, baz
S-5
S-6       !$omp do ordered(2)
S-7       do j=2, N
S-8         do i=2, M
S-9           A(i,j) = foo(i, j)
S-10
S-11        !$omp ordered depend(sink: j-1,i) depend(sink: j,i-1)
S-12          B(i,j) = bar(A(i,j), B(i-1,j), B(i,j-1))
S-13        !$omp ordered depend(source)
S-14
S-15          C(i,j) = baz(B(i,j))
S-16        end do
S-17      end do
S-18    end subroutine
```

Fortran

The following example shows the incorrect use of the **ordered** directive with a **depend** clause. There are two issues with the code. The first issue is a missing **ordered depend(source)** directive, which could cause a deadlock. The second issue is the **depend(sink:i+1,j)** and **depend(sink:i,j+1)** clauses define dependences on lexicographically later source iterations (*i+1, j*) and (*i, j+1*), which could cause a deadlock as well since they may not start to execute until the current iteration completes.

C / C++

1 *Example doacross.3.c* (**omp_4.5**)

```
S-1   #define N 100
S-2
S-3   void work_wrong(double p[][N][N])
S-4   {
S-5     int i, j, k;
S-6
S-7   #pragma omp parallel for ordered(2) private(i,j,k)
S-8     for (i=1; i<N-1; i++)
S-9     {
S-10      for (j=1; j<N-1; j++)
S-11      {
S-12  #pragma omp ordered depend(sink: i-1,j) depend(sink: i+1,j) \
S-13                     depend(sink: i,j-1) depend(sink: i,j+1)
S-14        for (k=1; k<N-1; k++)
S-15        {
S-16          double tmp1 = p[i-1][j][k] + p[i+1][j][k];
S-17          double tmp2 = p[i][j-1][k] + p[i][j+1][k];
S-18          double tmp3 = p[i][j][k-1] + p[i][j][k+1];
S-19          p[i][j][k] = (tmp1 + tmp2 + tmp3) / 6.0;
S-20        }
S-21  /* missing #pragma omp ordered depend(source) */
S-22      }
S-23    }
S-24  }
```

C / C++

Fortran

2 *Example doacross.3.f90* (**omp_4.5**)

```
S-1   subroutine work_wrong(N, p)
S-2     integer :: N
S-3     real(8), dimension(N,N,N) :: p
S-4     integer :: i, j, k
S-5     real(8) :: tmp1, tmp2, tmp3
S-6
S-7   !$omp parallel do ordered(2) private(i,j,k,tmp1,tmp2,tmp3)
S-8     do i=2, N-1
S-9       do j=2, N-1
S-10      !$omp ordered depend(sink: i-1,j) depend(sink: i+1,j) &
S-11      !$omp&        depend(sink: i,j-1) depend(sink: i,j+1)
S-12        do k=2, N-1
S-13          tmp1 = p(k-1,j,i) + p(k+1,j,i)
S-14          tmp2 = p(k,j-1,i) + p(k,j+1,i)
S-15          tmp3 = p(k,j,i-1) + p(k,j,i+1)
```

```
S-16            p(k,j,i) = (tmp1 + tmp2 + tmp3) / 6.0
S-17          end do
S-18    ! missing !$omp ordered depend(source)
S-19        end do
S-20      end do
S-21    end subroutine
```

Fortran

The following example illustrates the use of the **collapse** clause for a *doacross loop nest*. The *i* and *j* loops are the associated loops for the collapsed loop as well as for the *doacross loop nest*. The example also shows a compliant usage of the dependence source directive placed before the corresponding sink directive. Checking the completion of computation from previous iterations at the sink point can occur after the source statement.

C / C++

Example doacross.4.c (**omp_4.5**)

```
S-1     double foo(int i, int j);
S-2
S-3     void work( int N, int M, double **A, double **B, double **C )
S-4     {
S-5       int i, j;
S-6       double alpha = 1.2;
S-7
S-8       #pragma omp for collapse(2) ordered(2)
S-9       for (i = 1; i < N-1; i++)
S-10      {
S-11        for (j = 1; j < M-1; j++)
S-12        {
S-13          A[i][j] = foo(i, j);
S-14      #pragma omp ordered depend(source)
S-15
S-16          B[i][j] = alpha * A[i][j];
S-17
S-18      #pragma omp ordered depend(sink: i-1,j) depend(sink: i,j-1)
S-19          C[i][j] = 0.2 * (A[i-1][j] + A[i+1][j] +
S-20                           A[i][j-1] + A[i][j+1] + A[i][j]);
S-21        }
S-22      }
S-23    }
```

C / C++

Fortran

Example doacross.4.f90 (**omp_4.5**)

```
S-1   subroutine work( N, M, A, B, C )
S-2     integer :: N, M
S-3     real(8), dimension(M, N) :: A, B, C
S-4     real(8), external :: foo
S-5     integer :: i, j
S-6     real(8) :: alpha = 1.2
S-7
S-8     !$omp do collapse(2) ordered(2)
S-9     do j=2, N-1
S-10      do i=2, M-1
S-11        A(i,j) = foo(i, j)
S-12      !$omp ordered depend(source)
S-13
S-14        B(i,j) = alpha * A(i,j)
S-15
S-16      !$omp ordered depend(sink: j,i-1) depend(sink: j-1,i)
S-17        C(i,j) = 0.2 * (A(i-1,j) + A(i+1,j) +  &
S-18                  A(i,j-1) + A(i,j+1) + A(i,j))
S-19      end do
S-20    end do
S-21  end subroutine
```

Fortran

6.11 Lock Routines

This section is about the use of lock routines for synchronization.

6.11.1 The `omp_init_lock` Routine

The following example demonstrates how to initialize an array of locks in a **parallel** region by using **omp_init_lock**.

C++

Example init_lock.1.cpp

```
S-1   #include <omp.h>
S-2
S-3   omp_lock_t *new_locks() {
S-4     int i;
S-5     omp_lock_t *lock = new omp_lock_t[1000];
S-6
S-7     #pragma omp parallel for private(i)
S-8       for (i=0; i<1000; i++)
S-9       { omp_init_lock(&lock[i]); }
S-10
S-11      return lock;
S-12  }
```

C++

Fortran

Example init_lock.1.f

```
S-1         FUNCTION NEW_LOCKS()
S-2           USE OMP_LIB        ! or INCLUDE "omp_lib.h"
S-3           INTEGER(OMP_LOCK_KIND), DIMENSION(1000) :: NEW_LOCKS
S-4           INTEGER I
S-5
S-6   !$OMP   PARALLEL DO PRIVATE(I)
S-7             DO I=1,1000
S-8               CALL OMP_INIT_LOCK(NEW_LOCKS(I))
S-9             END DO
S-10  !$OMP   END PARALLEL DO
S-11
S-12        END FUNCTION NEW_LOCKS
```

Fortran

6.11.2 The `omp_init_lock_with_hint` Routine

The following example demonstrates how to initialize an array of locks in a **parallel** region by using **omp_init_lock_with_hint**. Note, hints are combined with an **|** or **+** operator in C/C++ and a **+** operator in Fortran.

C++

Example init_lock_with_hint.1.cpp (**omp_4.5**)

```
#include <omp.h>

omp_lock_t *new_locks()
{
  int i;
  omp_lock_t *lock = new omp_lock_t[1000];

  #pragma omp parallel for private(i)
    for (i=0; i<1000; i++)
    {
      omp_init_lock_with_hint(&lock[i],
        static_cast<omp_lock_hint_t>(omp_lock_hint_contended |
                                     omp_lock_hint_speculative));
    }
    return lock;
}
```

C++

Fortran

Example init_lock_with_hint.1.f (**omp_4.5**)

```
      FUNCTION NEW_LOCKS()
        USE OMP_LIB        ! or INCLUDE "omp_lib.h"
        INTEGER(OMP_LOCK_KIND), DIMENSION(1000) :: NEW_LOCKS

        INTEGER I

!$OMP   PARALLEL DO PRIVATE(I)
          DO I=1,1000
            CALL OMP_INIT_LOCK_WITH_HINT(NEW_LOCKS(I),
     &             OMP_LOCK_HINT_CONTENDED + OMP_LOCK_HINT_SPECULATIVE)
          END DO
!$OMP   END PARALLEL DO

      END FUNCTION NEW_LOCKS
```

Fortran

6.11.3 Ownership of Locks

Ownership of locks has changed since OpenMP 2.5. In OpenMP 2.5, locks are owned by threads; so a lock released by the **omp_unset_lock** routine must be owned by the same thread executing the routine. Beginning with OpenMP 3.0, locks are owned by task regions; so a lock released by the **omp_unset_lock** routine in a task region must be owned by the same task region.

This change in ownership requires extra care when using locks. The following program is conforming in OpenMP 2.5 because the thread that releases the lock **lck** in the parallel region is the same thread that acquired the lock in the sequential part of the program (master thread of parallel region and the initial thread are the same). However, it is not conforming beginning with OpenMP 3.0, because the task region that releases the lock **lck** is different from the task region that acquires the lock.

C / C++

Example lock_owner.1.c

```
S-1   #include <stdlib.h>
S-2   #include <stdio.h>
S-3   #include <omp.h>
S-4
S-5   int main()
S-6   {
S-7     int x;
S-8     omp_lock_t lck;
S-9
S-10    omp_init_lock (&lck);
S-11    omp_set_lock (&lck);
S-12    x = 0;
S-13
S-14  #pragma omp parallel shared (x)
S-15    {
S-16      #pragma omp master
S-17        {
S-18          x = x + 1;
S-19          omp_unset_lock (&lck);
S-20        }
S-21
S-22      /* Some more stuff. */
S-23    }
S-24    omp_destroy_lock (&lck);
S-25    return 0;
S-26  }
```

C / C++

Fortran

Example lock_owner.1.f

```
S-1         program lock
S-2         use omp_lib
S-3         integer :: x
S-4         integer (kind=omp_lock_kind) :: lck
S-5
S-6         call omp_init_lock (lck)
S-7         call omp_set_lock(lck)
S-8         x = 0
S-9
S-10    !$omp parallel shared (x)
S-11    !$omp master
S-12        x = x + 1
S-13        call omp_unset_lock(lck)
S-14    !$omp end master
S-15
S-16    !       Some more stuff.
S-17    !$omp end parallel
S-18
S-19        call omp_destroy_lock(lck)
S-20        end
```

Fortran

6.11.4 Simple Lock Routines

In the following example, the lock routines cause the threads to be idle while waiting for entry to the first critical section, but to do other work while waiting for entry to the second. The **omp_set_lock** function blocks, but the **omp_test_lock** function does not, allowing the work in **skip** to be done.

Note that the argument to the lock routines should have type **omp_lock_t**, and that there is no need to flush it.

C / C++

Example simple_lock.1.c

```
#include <stdio.h>
#include <omp.h>
void skip(int i) {}
void work(int i) {}
int main()
{
  omp_lock_t lck;
  int id;
  omp_init_lock(&lck);

  #pragma omp parallel shared(lck) private(id)
  {
    id = omp_get_thread_num();

    omp_set_lock(&lck);
    /*  only one thread at a time can execute this printf */
    printf("My thread id is %d.\n", id);
    omp_unset_lock(&lck);

    while (! omp_test_lock(&lck)) {
      skip(id);   /* we do not yet have the lock,
                     so we must do something else */
    }

    work(id);      /* we now have the lock
                      and can do the work */

    omp_unset_lock(&lck);
  }
  omp_destroy_lock(&lck);

  return 0;
}
```

C / C++

Note that there is no need to flush the lock variable.

Fortran

Example simple_lock.1.f

```
S-1         SUBROUTINE SKIP(ID)
S-2         END SUBROUTINE SKIP
S-3
S-4         SUBROUTINE WORK(ID)
S-5         END SUBROUTINE WORK
S-6
S-7         PROGRAM SIMPLELOCK
S-8
S-9           INCLUDE "omp_lib.h"     ! or USE OMP_LIB
S-10
S-11          INTEGER(OMP_LOCK_KIND) LCK
S-12          INTEGER ID
S-13
S-14          CALL OMP_INIT_LOCK(LCK)
S-15
S-16    !$OMP   PARALLEL SHARED(LCK) PRIVATE(ID)
S-17              ID = OMP_GET_THREAD_NUM()
S-18              CALL OMP_SET_LOCK(LCK)
S-19              PRINT *, 'My thread id is ', ID
S-20              CALL OMP_UNSET_LOCK(LCK)
S-21
S-22              DO WHILE (.NOT. OMP_TEST_LOCK(LCK))
S-23                CALL SKIP(ID)      ! We do not yet have the lock
S-24                                   ! so we must do something else
S-25              END DO
S-26
S-27              CALL WORK(ID)        ! We now have the lock
S-28                                   ! and can do the work
S-29
S-30              CALL OMP_UNSET_LOCK( LCK )
S-31
S-32    !$OMP   END PARALLEL
S-33
S-34          CALL OMP_DESTROY_LOCK( LCK )
S-35
S-36        END PROGRAM SIMPLELOCK
```

Fortran

6.11.5 Nestable Lock Routines

The following example demonstrates how a nestable lock can be used to synchronize updates both to a whole structure and to one of its members.

C / C++

Example nestable_lock.1.c

```
S-1
S-2   #include <omp.h>
S-3
S-4   typedef struct {
S-5      int a,b;
S-6      omp_nest_lock_t lck;
S-7   } pair;
S-8
S-9   int work1();
S-10  int work2();
S-11  int work3();
S-12
S-13  void incr_a(pair *p, int a)
S-14  {
S-15
S-16    /* Called only from incr_pair, no need to lock. */
S-17    p->a += a;
S-18
S-19  }
S-20
S-21  void incr_b(pair *p, int b)
S-22  {
S-23
S-24    /* Called both from incr_pair and elsewhere, */
S-25    /* so need a nestable lock. */
S-26
S-27    omp_set_nest_lock(&p->lck);
S-28    p->b += b;
S-29    omp_unset_nest_lock(&p->lck);
S-30
S-31  }
S-32
S-33  void incr_pair(pair *p, int a, int b)
S-34  {
S-35
S-36    omp_set_nest_lock(&p->lck);
S-37    incr_a(p, a);
S-38    incr_b(p, b);
S-39    omp_unset_nest_lock(&p->lck);
S-40
```

```
S-41  }
S-42
S-43  void nestlock(pair *p)
S-44  {
S-45
S-46    #pragma omp parallel sections
S-47    {
S-48      #pragma omp section
S-49        incr_pair(p, work1(), work2());
S-50      #pragma omp section
S-51        incr_b(p, work3());
S-52    }
S-53
S-54  }
```

C / C++

Fortran

1 *Example nestable_lock.1.f*

```
S-1         MODULE DATA
S-2           USE OMP_LIB, ONLY: OMP_NEST_LOCK_KIND
S-3           TYPE LOCKED_PAIR
S-4             INTEGER A
S-5             INTEGER B
S-6             INTEGER (OMP_NEST_LOCK_KIND) LCK
S-7           END TYPE
S-8         END MODULE DATA
S-9
S-10        SUBROUTINE INCR_A(P, A)
S-11          ! called only from INCR_PAIR, no need to lock
S-12          USE DATA
S-13          TYPE(LOCKED_PAIR) :: P
S-14          INTEGER A
S-15          P%A = P%A + A
S-16        END SUBROUTINE INCR_A
S-17
S-18        SUBROUTINE INCR_B(P, B)
S-19          ! called from both INCR_PAIR and elsewhere,
S-20          ! so we need a nestable lock
S-21          USE OMP_LIB       ! or INCLUDE "omp_lib.h"
S-22          USE DATA
S-23          TYPE(LOCKED_PAIR) :: P
S-24          INTEGER B
S-25          CALL OMP_SET_NEST_LOCK(P%LCK)
S-26          P%B = P%B + B
S-27          CALL OMP_UNSET_NEST_LOCK(P%LCK)
S-28        END SUBROUTINE INCR_B
```

```
S-29
S-30          SUBROUTINE INCR_PAIR(P, A, B)
S-31            USE OMP_LIB        ! or INCLUDE "omp_lib.h"
S-32            USE DATA
S-33            TYPE(LOCKED_PAIR) :: P
S-34            INTEGER A
S-35            INTEGER B
S-36
S-37            CALL OMP_SET_NEST_LOCK(P%LCK)
S-38            CALL INCR_A(P, A)
S-39            CALL INCR_B(P, B)
S-40            CALL OMP_UNSET_NEST_LOCK(P%LCK)
S-41          END SUBROUTINE INCR_PAIR
S-42
S-43          SUBROUTINE NESTLOCK(P)
S-44            USE OMP_LIB        ! or INCLUDE "omp_lib.h"
S-45            USE DATA
S-46            TYPE(LOCKED_PAIR) :: P
S-47            INTEGER WORK1, WORK2, WORK3
S-48            EXTERNAL WORK1, WORK2, WORK3
S-49
S-50      !$OMP   PARALLEL SECTIONS
S-51
S-52      !$OMP   SECTION
S-53              CALL INCR_PAIR(P, WORK1(), WORK2())
S-54      !$OMP   SECTION
S-55              CALL INCR_B(P, WORK3())
S-56      !$OMP   END PARALLEL SECTIONS
S-57
S-58          END SUBROUTINE NESTLOCK
```

Fortran

CHAPTER 7

Data Environment

The OpenMP *data environment* contains data attributes of variables and objects. Many constructs (such as **parallel**, **simd**, **task**) accept clauses to control *data-sharing* attributes of referenced variables in the construct, where *data-sharing* applies to whether the attribute of the variable is *shared*, is *private* storage, or has special operational characteristics (as found in the **firstprivate**, **lastprivate**, **linear**, or **reduction** clause).

The data environment for a device (distinguished as a *device data environment*) is controlled on the host by *data-mapping* attributes, which determine the relationship of the data on the host, the *original* data, and the data on the device, the *corresponding* data.

DATA-SHARING ATTRIBUTES

Data-sharing attributes of variables can be classified as being *predetermined*, *explicitly determined* or *implicitly determined*.

Certain variables and objects have predetermined attributes. A commonly found case is the loop iteration variable in associated loops of a **for** or **do** construct. It has a private data-sharing attribute. Variables with predetermined data-sharing attributes can not be listed in a data-sharing clause; but there are some exceptions (mainly concerning loop iteration variables).

Variables with explicitly determined data-sharing attributes are those that are referenced in a given construct and are listed in a data-sharing attribute clause on the construct. Some of the common data-sharing clauses are: **shared**, **private**, **firstprivate**, **lastprivate**, **linear**, and **reduction**.

Variables with implicitly determined data-sharing attributes are those that are referenced in a given construct, do not have predetermined data-sharing attributes, and are not listed in a data-sharing attribute clause of an enclosing construct. For a complete list of variables and objects with predetermined and implicitly determined attributes, please refer to the *Data-sharing Attribute Rules for Variables Referenced in a Construct* subsection of the OpenMP Specifications document.

DATA-MAPPING ATTRIBUTES

The **map** clause on a device construct explicitly specifies how the list items in the clause are mapped from the encountering task's data environment (on the host) to the corresponding item in the device data environment (on the device). The common *list items* are arrays, array sections, scalars, pointers, and structure elements (members).

Procedures and global variables have predetermined data mapping if they appear within the list or block of a **declare target** directive. Also, a C/C++ pointer is mapped as a zero-length array section, as is a C++ variable that is a reference to a pointer.

Without explicit mapping, non-scalar and non-pointer variables within the scope of the **target** construct are implicitly mapped with a *map-type* of **tofrom**. Without explicit mapping, scalar variables within the scope of the **target** construct are not mapped, but have an implicit firstprivate data-sharing attribute. (That is, the value of the original variable is given to a private variable of the same name on the device.) This behavior can be changed with the **defaultmap** clause.

The **map** clause can appear on **target**, **target data** and **target enter/exit data** constructs. The operations of creation and removal of device storage as well as assignment of the original list item values to the corresponding list items may be complicated when the list item appears on multiple constructs or when the host and device storage is shared. In these cases the item's reference count, the number of times it has been referenced (+1 on entry and -1 on exited) in nested (structured) map regions and/or accumulative (unstructured) mappings, determines the operation. Details of the **map** clause and reference count operation are specified in the *map Clause* subsection of the OpenMP Specifications document.

7.1 The `threadprivate` Directive

The following examples demonstrate how to use the **`threadprivate`** directive to give each thread a separate counter.

C / C++

Example threadprivate.1.c

```
S-1  int counter = 0;
S-2  #pragma omp threadprivate(counter)
S-3
S-4  int increment_counter()
S-5  {
S-6    counter++;
S-7    return(counter);
S-8  }
```

C / C++

Fortran

Example threadprivate.1.f

```
S-1          INTEGER FUNCTION INCREMENT_COUNTER()
S-2            COMMON/INC_COMMON/COUNTER
S-3    !$OMP   THREADPRIVATE(/INC_COMMON/)
S-4
S-5            COUNTER = COUNTER +1
S-6            INCREMENT_COUNTER = COUNTER
S-7            RETURN
S-8          END FUNCTION INCREMENT_COUNTER
```

Fortran

C / C++

The following example uses **`threadprivate`** on a static variable:

Example threadprivate.2.c

```
S-1  int increment_counter_2()
S-2  {
S-3    static int counter = 0;
S-4    #pragma omp threadprivate(counter)
S-5    counter++;
S-6    return(counter);
S-7  }
```

The following example demonstrates unspecified behavior for the initialization of a **threadprivate** variable. A **threadprivate** variable is initialized once at an unspecified point before its first reference. Because **a** is constructed using the value of **x** (which is modified by the statement **x++**), the value of **a.val** at the start of the **parallel** region could be either 1 or 2. This problem is avoided for **b**, which uses an auxiliary **const** variable and a copy-constructor.

Example threadprivate.3.cpp

```
S-1   class T {
S-2     public:
S-3       int val;
S-4       T (int);
S-5       T (const T&);
S-6   };
S-7
S-8   T :: T (int v){
S-9      val = v;
S-10  }
S-11
S-12  T :: T (const T& t) {
S-13     val = t.val;
S-14  }
S-15
S-16  void g(T a, T b){
S-17     a.val += b.val;
S-18  }
S-19
S-20  int x = 1;
S-21  T a(x);
S-22  const T b_aux(x); /* Capture value of x = 1 */
S-23  T b(b_aux);
S-24  #pragma omp threadprivate(a, b)
S-25
S-26  void f(int n) {
S-27     x++;
S-28     #pragma omp parallel for
S-29     /* In each thread:
S-30      * a is constructed from x (with value 1 or 2?)
S-31      * b is copy-constructed from b_aux
S-32      */
S-33
S-34     for (int i=0; i<n; i++) {
S-35         g(a, b); /* Value of a is unspecified. */
S-36     }
S-37  }
```

C / C++

The following examples show non-conforming uses and correct uses of the **threadprivate** directive.

Fortran

The following example is non-conforming because the common block is not declared local to the subroutine that refers to it:

Example threadprivate.2.f

```
      MODULE INC_MODULE
        COMMON /T/ A
      END MODULE INC_MODULE

      SUBROUTINE INC_MODULE_WRONG()
        USE INC_MODULE
!$OMP   THREADPRIVATE(/T/)
      !non-conforming because /T/ not declared in INC_MODULE_WRONG
      END SUBROUTINE INC_MODULE_WRONG
```

The following example is also non-conforming because the common block is not declared local to the subroutine that refers to it:

Example threadprivate.3.f

```
      SUBROUTINE INC_WRONG()
        COMMON /T/ A
!$OMP   THREADPRIVATE(/T/)

        CONTAINS
          SUBROUTINE INC_WRONG_SUB()
!$OMP       PARALLEL COPYIN(/T/)
      !non-conforming because /T/ not declared in INC_WRONG_SUB
!$OMP       END PARALLEL
          END SUBROUTINE INC_WRONG_SUB
      END SUBROUTINE INC_WRONG
```

The following example is a correct rewrite of the previous example:

Example threadprivate.4.f

```
       SUBROUTINE INC_GOOD()
        COMMON /T/ A
!$OMP   THREADPRIVATE(/T/)

        CONTAINS
          SUBROUTINE INC_GOOD_SUB()
            COMMON /T/ A
!$OMP       THREADPRIVATE(/T/)

!$OMP       PARALLEL COPYIN(/T/)
```

```
!$OMP       END PARALLEL
          END SUBROUTINE INC_GOOD_SUB
        END SUBROUTINE INC_GOOD
```

The following is an example of the use of **threadprivate** for local variables:

Example threadprivate.5.f

```
      PROGRAM INC_GOOD2
        INTEGER, ALLOCATABLE, SAVE :: A(:)
        INTEGER, POINTER, SAVE :: PTR
        INTEGER, SAVE :: I
        INTEGER, TARGET :: TARG
        LOGICAL :: FIRSTIN = .TRUE.
!$OMP   THREADPRIVATE(A, I, PTR)

        ALLOCATE (A(3))
        A = (/1,2,3/)
        PTR => TARG
        I = 5

!$OMP   PARALLEL COPYIN(I, PTR)
!$OMP     CRITICAL
            IF (FIRSTIN) THEN
              TARG = 4           ! Update target of ptr
              I = I + 10
              IF (ALLOCATED(A)) A = A + 10
              FIRSTIN = .FALSE.
            END IF

            IF (ALLOCATED(A)) THEN
              PRINT *, 'a = ', A
            ELSE
              PRINT *, 'A is not allocated'
            END IF

            PRINT *, 'ptr = ', PTR
            PRINT *, 'i = ', I
            PRINT *

!$OMP     END CRITICAL
!$OMP   END PARALLEL
      END PROGRAM INC_GOOD2
```

The above program, if executed by two threads, will print one of the following two sets of output:

```
a = 11 12 13
ptr = 4
i = 15

A is not allocated
ptr = 4
i = 5
```

or

```
A is not allocated
ptr = 4
i = 15

a = 1 2 3
ptr = 4
i = 5
```

The following is an example of the use of **threadprivate** for module variables:

Example threadprivate.6.f

```
      MODULE INC_MODULE_GOOD3
        REAL, POINTER :: WORK(:)
        SAVE WORK
!$OMP   THREADPRIVATE(WORK)
      END MODULE INC_MODULE_GOOD3

      SUBROUTINE SUB1(N)
      USE INC_MODULE_GOOD3
!$OMP   PARALLEL PRIVATE(THE_SUM)
        ALLOCATE(WORK(N))
        CALL SUB2(THE_SUM)
       WRITE(*,*)THE_SUM
!$OMP   END PARALLEL
      END SUBROUTINE SUB1

      SUBROUTINE SUB2(THE_SUM)
        USE INC_MODULE_GOOD3
        WORK(:) = 10
        THE_SUM=SUM(WORK)
      END SUBROUTINE SUB2

      PROGRAM INC_GOOD3
        N = 10
        CALL SUB1(N)
```

```
S-25        END PROGRAM INC_GOOD3
```

Fortran

C++

The following example illustrates initialization of **threadprivate** variables for class-type **T**. **t1** is default constructed, **t2** is constructed taking a constructor accepting one argument of integer type, **t3** is copy constructed with argument **f()**:

Example threadprivate.4.cpp

```
S-1   struct T { T (); T (int); ~T (); int t; };
S-2   int f();
S-3   static T t1;
S-4   #pragma omp threadprivate(t1)
S-5   static T t2( 23 );
S-6   #pragma omp threadprivate(t2)
S-7   static T t3 = f();
S-8   #pragma omp threadprivate(t3)
```

The following example illustrates the use of **threadprivate** for static class members. The **threadprivate** directive for a static class member must be placed inside the class definition.

Example threadprivate.5.cpp

```
S-1   class T {
S-2    public:
S-3     static int i;
S-4   #pragma omp threadprivate(i)
S-5   };
```

C++

7.2 The `default(none)` Clause

The following example distinguishes the variables that are affected by the **`default(none)`** clause from those that are not.

C / C++

Beginning with OpenMP 4.0, variables with **`const`**-qualified type and no mutable member are no longer predetermined shared. Thus, these variables (variable *c* in the example) need to be explicitly listed in data-sharing attribute clauses when the **`default(none)`** clause is specified.

Example default_none.1.c

```
#include <omp.h>
int x, y, z[1000];
#pragma omp threadprivate(x)

void default_none(int a) {
  const int c = 1;
  int i = 0;

  #pragma omp parallel default(none) private(a) shared(z, c)
  {
     int j = omp_get_num_threads();
          /* O.K.  - j is declared within parallel region */
     a = z[j];   /* O.K.  - a is listed in private clause */
                 /*       - z is listed in shared clause */
     x = c;      /* O.K.  - x is threadprivate */
                 /*       - c has const-qualified type and
                                 is listed in shared clause */
     z[i] = y;   /* Error - cannot reference i or y here */

  #pragma omp for firstprivate(y)
          /* Error - Cannot reference y in the firstprivate clause */
     for (i=0; i<10 ; i++) {
        z[i] = i; /* O.K. - i is the loop iteration variable */
     }

     z[i] = y;   /* Error - cannot reference i or y here */
  }
}
```

C / C++

Fortran

Example default_none.1.f

```
      SUBROUTINE DEFAULT_NONE(A)
      INCLUDE "omp_lib.h"     ! or USE OMP_LIB

      INTEGER A

      INTEGER X, Y, Z(1000)
      COMMON/BLOCKX/X
      COMMON/BLOCKY/Y
      COMMON/BLOCKZ/Z
!$OMP THREADPRIVATE(/BLOCKX/)

        INTEGER I, J
        i = 1

!$OMP   PARALLEL DEFAULT(NONE) PRIVATE(A) SHARED(Z) PRIVATE(J)
          J = OMP_GET_NUM_THREADS();
                   ! O.K.  - J is listed in PRIVATE clause
          A = Z(J) ! O.K.  - A is listed in PRIVATE clause
                   !       - Z is listed in SHARED clause
          X = 1    ! O.K.  - X is THREADPRIVATE
          Z(I) = Y ! Error - cannot reference I or Y here

!$OMP DO firstprivate(y)
    ! Error - Cannot reference y in the firstprivate clause
          DO I = 1,10
             Z(I) = I ! O.K. - I is the loop iteration variable
          END DO


          Z(I) = Y    ! Error - cannot reference I or Y here
!$OMP   END PARALLEL
      END SUBROUTINE DEFAULT_NONE
```

Fortran

7.3 The `private` Clause

In the following example, the values of original list items *i* and *j* are retained on exit from the **`parallel`** region, while the private list items *i* and *j* are modified within the **`parallel`** construct.

C / C++

Example private.1.c

```
S-1   #include <stdio.h>
S-2   #include <assert.h>
S-3
S-4   int main()
S-5   {
S-6     int i, j;
S-7     int *ptr_i, *ptr_j;
S-8
S-9     i = 1;
S-10    j = 2;
S-11
S-12    ptr_i = &i;
S-13    ptr_j = &j;
S-14
S-15    #pragma omp parallel private(i) firstprivate(j)
S-16    {
S-17      i = 3;
S-18      j = j + 2;
S-19      assert (*ptr_i == 1 && *ptr_j == 2);
S-20    }
S-21
S-22    assert(i == 1 && j == 2);
S-23
S-24    return 0;
S-25  }
```

C / C++

Fortran

Example private.1.f

```
      PROGRAM PRIV_EXAMPLE
        INTEGER I, J

        I = 1
        J = 2

!$OMP   PARALLEL PRIVATE(I) FIRSTPRIVATE(J)
          I = 3
          J = J + 2
!$OMP   END PARALLEL

        PRINT *, I, J  ! I .eq. 1 .and. J .eq. 2
      END PROGRAM PRIV_EXAMPLE
```

Fortran

In the following example, all uses of the variable *a* within the loop construct in the routine *f* refer to a private list item *a*, while it is unspecified whether references to *a* in the routine *g* are to a private list item or the original list item.

C / C++

Example private.2.c

```
int a;

void g(int k) {
  a = k; /* Accessed in the region but outside of the construct;
          * therefore unspecified whether original or private list
          * item is modified. */
}


void f(int n) {
  int a = 0;

  #pragma omp parallel for private(a)
   for (int i=1; i<n; i++) {
       a = i;
       g(a*2);     /* Private copy of "a" */
    }
}
```

C / C++

Fortran

Example private.2.f

```
      MODULE PRIV_EXAMPLE2
        REAL A

        CONTAINS

          SUBROUTINE G(K)
            REAL K
            A = K  ! Accessed in the region but outside of the
                   ! construct; therefore unspecified whether
                   ! original or private list item is modified.
          END SUBROUTINE G

          SUBROUTINE F(N)
          INTEGER N
          REAL A

            INTEGER I
!$OMP       PARALLEL DO PRIVATE(A)
              DO I = 1,N
                A = I
                CALL G(A*2)
              ENDDO
!$OMP       END PARALLEL DO
          END SUBROUTINE F

      END MODULE PRIV_EXAMPLE2
```

Fortran

The following example demonstrates that a list item that appears in a **private** clause in a **parallel** construct may also appear in a **private** clause in an enclosed worksharing construct, which results in an additional private copy.

C / C++

Example private.3.c

```
#include <assert.h>
void priv_example3()
{
  int i, a;

  #pragma omp parallel private(a)
  {
     a = 1;
    #pragma omp parallel for private(a)
```

```
S-10          for (i=0; i<10; i++)
S-11         {
S-12           a = 2;
S-13         }
S-14       assert(a == 1);
S-15     }
S-16   }
```

C / C++

Fortran

1 *Example private.3.f*

```
S-1          SUBROUTINE PRIV_EXAMPLE3()
S-2            INTEGER I, A
S-3
S-4    !$OMP   PARALLEL PRIVATE(A)
S-5             A = 1
S-6    !$OMP     PARALLEL DO PRIVATE(A)
S-7              DO I = 1, 10
S-8                A = 2
S-9              END DO
S-10   !$OMP     END PARALLEL DO
S-11           PRINT *, A ! Outer A still has value 1
S-12   !$OMP   END PARALLEL
S-13         END SUBROUTINE PRIV_EXAMPLE3
```

Fortran

7.4 Fortran Private Loop Iteration Variables

Fortran

In general loop iteration variables will be private, when used in the *do-loop* of a **do** and **parallel do** construct or in sequential loops in a **parallel** construct (see Section 2.7.1 and Section 2.14.1 of the OpenMP 4.0 specification). In the following example of a sequential loop in a **parallel** construct the loop iteration variable *I* will be private.

Example fort_loopvar.1.f90

```
S-1     SUBROUTINE PLOOP_1(A,N)
S-2     INCLUDE "omp_lib.h"      ! or USE OMP_LIB
S-3
S-4     REAL A(*)
S-5     INTEGER I, MYOFFSET, N
S-6
S-7     !$OMP PARALLEL PRIVATE(MYOFFSET)
S-8           MYOFFSET = OMP_GET_THREAD_NUM()*N
S-9           DO I = 1, N
S-10            A(MYOFFSET+I) = FLOAT(I)
S-11          ENDDO
S-12    !$OMP END PARALLEL
S-13
S-14    END SUBROUTINE PLOOP_1
```

In exceptional cases, loop iteration variables can be made shared, as in the following example:

Example fort_loopvar.2.f90

```
S-1     SUBROUTINE PLOOP_2(A,B,N,I1,I2)
S-2     REAL A(*), B(*)
S-3     INTEGER I1, I2, N
S-4
S-5     !$OMP PARALLEL SHARED(A,B,I1,I2)
S-6     !$OMP SECTIONS
S-7     !$OMP SECTION
S-8          DO I1 = I1, N
S-9            IF (A(I1).NE.0.0) EXIT
S-10         ENDDO
S-11    !$OMP SECTION
S-12         DO I2 = I2, N
S-13           IF (B(I2).NE.0.0) EXIT
S-14         ENDDO
S-15    !$OMP END SECTIONS
S-16    !$OMP SINGLE
S-17        IF (I1.LE.N) PRINT *, 'ITEMS IN A UP TO ', I1, 'ARE ALL ZERO.'
S-18        IF (I2.LE.N) PRINT *, 'ITEMS IN B UP TO ', I2, 'ARE ALL ZERO.'
S-19    !$OMP END SINGLE
```

```
S-20    !$OMP END PARALLEL
S-21
S-22    END SUBROUTINE PLOOP_2
```

Note however that the use of shared loop iteration variables can easily lead to race conditions.

Fortran

7.5 Fortran Restrictions on `shared` and `private` Clauses with Common Blocks

Fortran

When a named common block is specified in a **`private`**, **`firstprivate`**, or **`lastprivate`** clause of a construct, none of its members may be declared in another data-sharing attribute clause on that construct. The following examples illustrate this point.

The following example is conforming:

Example fort_sp_common.1.f

```
S-1         SUBROUTINE COMMON_GOOD()
S-2           COMMON /C/ X,Y
S-3           REAL X, Y
S-4
S-5     !$OMP   PARALLEL PRIVATE (/C/)
S-6             ! do work here
S-7     !$OMP   END PARALLEL
S-8     !$OMP   PARALLEL SHARED (X,Y)
S-9             ! do work here
S-10    !$OMP   END PARALLEL
S-11        END SUBROUTINE COMMON_GOOD
```

The following example is also conforming:

Example fort_sp_common.2.f

```
S-1         SUBROUTINE COMMON_GOOD2()
S-2           COMMON /C/ X,Y
S-3           REAL X, Y
S-4           INTEGER I
S-5     !$OMP   PARALLEL
S-6     !$OMP     DO PRIVATE(/C/)
S-7               DO I=1,1000
S-8                 ! do work here
S-9               ENDDO
S-10    !$OMP     END DO
S-11    !$OMP     DO PRIVATE(X)
S-12              DO I=1,1000
S-13                ! do work here
S-14              ENDDO
S-15    !$OMP     END DO
S-16    !$OMP   END PARALLEL
S-17        END SUBROUTINE COMMON_GOOD2
```

The following example is conforming:

Example fort_sp_common.3.f

```
S-1         SUBROUTINE COMMON_GOOD3()
S-2           COMMON /C/ X,Y
S-3     !$OMP   PARALLEL PRIVATE (/C/)
S-4             ! do work here
S-5     !$OMP   END PARALLEL
S-6     !$OMP   PARALLEL SHARED (/C/)
S-7             ! do work here
S-8     !$OMP   END PARALLEL
S-9         END SUBROUTINE COMMON_GOOD3
```

The following example is non-conforming because **x** is a constituent element of **c**:

Example fort_sp_common.4.f

```
S-1         SUBROUTINE COMMON_WRONG()
S-2           COMMON /C/ X,Y
S-3     ! Incorrect because X is a constituent element of C
S-4     !$OMP   PARALLEL PRIVATE(/C/), SHARED(X)
S-5             ! do work here
S-6     !$OMP   END PARALLEL
S-7         END SUBROUTINE COMMON_WRONG
```

The following example is non-conforming because a common block may not be declared both shared and private:

Example fort_sp_common.5.f

```
S-1         SUBROUTINE COMMON_WRONG2()
S-2           COMMON /C/ X,Y
S-3     ! Incorrect: common block C cannot be declared both
S-4     ! shared and private
S-5     !$OMP   PARALLEL PRIVATE (/C/), SHARED(/C/)
S-6             ! do work here
S-7     !$OMP   END PARALLEL
S-8
S-9         END SUBROUTINE COMMON_WRONG2
```

Fortran

7.6 Fortran Restrictions on Storage Association with the `private` Clause

Fortran

The following non-conforming examples illustrate the implications of the **`private`** clause rules with regard to storage association.

Example fort_sa_private.1.f

```
S-1           SUBROUTINE SUB()
S-2           COMMON /BLOCK/ X
S-3           PRINT *,X              ! X is undefined
S-4           END SUBROUTINE SUB
S-5
S-6           PROGRAM PRIV_RESTRICT
S-7             COMMON /BLOCK/ X
S-8             X = 1.0
S-9     !$OMP   PARALLEL PRIVATE (X)
S-10            X = 2.0
S-11            CALL SUB()
S-12    !$OMP   END PARALLEL
S-13          END PROGRAM PRIV_RESTRICT
```

Example fort_sa_private.2.f

```
S-1           PROGRAM PRIV_RESTRICT2
S-2             COMMON /BLOCK2/ X
S-3             X = 1.0
S-4
S-5     !$OMP   PARALLEL PRIVATE (X)
S-6               X = 2.0
S-7               CALL SUB()
S-8     !$OMP   END PARALLEL
S-9
S-10          CONTAINS
S-11
S-12             SUBROUTINE SUB()
S-13             COMMON /BLOCK2/ Y
S-14
S-15             PRINT *,X                ! X is undefined
S-16             PRINT *,Y                ! Y is undefined
S-17             END SUBROUTINE SUB
S-18
S-19          END PROGRAM PRIV_RESTRICT2
```

Example fort_sa_private.3.f

```
S-1         PROGRAM PRIV_RESTRICT3
S-2           EQUIVALENCE (X,Y)
S-3           X = 1.0
S-4
S-5     !$OMP   PARALLEL PRIVATE(X)
S-6             PRINT *,Y                  ! Y is undefined
S-7             Y = 10
S-8             PRINT *,X                  ! X is undefined
S-9     !$OMP   END PARALLEL
S-10        END PROGRAM PRIV_RESTRICT3
```

Example fort_sa_private.4.f

```
S-1         PROGRAM PRIV_RESTRICT4
S-2           INTEGER I, J
S-3           INTEGER A(100), B(100)
S-4           EQUIVALENCE (A(51), B(1))
S-5
S-6     !$OMP PARALLEL DO DEFAULT(PRIVATE) PRIVATE(I,J) LASTPRIVATE(A)
S-7             DO I=1,100
S-8                DO J=1,100
S-9                  B(J) = J - 1
S-10               ENDDO
S-11
S-12               DO J=1,100
S-13                 A(J) = J   ! B becomes undefined at this point
S-14               ENDDO
S-15
S-16               DO J=1,50
S-17                 B(J) = B(J) + 1  ! B is undefined
S-18                           ! A becomes undefined at this point
S-19               ENDDO
S-20            ENDDO
S-21    !$OMP END PARALLEL DO       ! The LASTPRIVATE write for A has
S-22                                ! undefined results
S-23
S-24          PRINT *, B    ! B is undefined since the LASTPRIVATE
S-25                        ! write of A was not defined
S-26        END PROGRAM PRIV_RESTRICT4
```

Example fort_sa_private.5.f

```
S-1         SUBROUTINE SUB1(X)
S-2           DIMENSION X(10)
S-3
S-4           ! This use of X does not conform to the
```

```
S-5           ! specification. It would be legal Fortran 90,
S-6           ! but the OpenMP private directive allows the
S-7           ! compiler to break the sequence association that
S-8           ! A had with the rest of the common block.
S-9
S-10            FORALL (I = 1:10) X(I) = I
S-11          END SUBROUTINE SUB1
S-12
S-13          PROGRAM PRIV_RESTRICT5
S-14            COMMON /BLOCK5/ A
S-15
S-16            DIMENSION B(10)
S-17            EQUIVALENCE (A,B(1))
S-18
S-19            ! the common block has to be at least 10 words
S-20            A = 0
S-21
S-22      !$OMP   PARALLEL PRIVATE(/BLOCK5/)
S-23
S-24              ! Without the private clause,
S-25              ! we would be passing a member of a sequence
S-26              ! that is at least ten elements long.
S-27              ! With the private clause, A may no longer be
S-28              ! sequence-associated.
S-29
S-30              CALL SUB1(A)
S-31      !$OMP     MASTER
S-32                PRINT *, A
S-33      !$OMP     END MASTER
S-34
S-35      !$OMP   END PARALLEL
S-36          END PROGRAM PRIV_RESTRICT5
```

Fortran

7.7 C/C++ Arrays in a `firstprivate` Clause

C / C++

The following example illustrates the size and value of list items of array or pointer type in a **`firstprivate`** clause . The size of new list items is based on the type of the corresponding original list item, as determined by the base language.

In this example:

- The type of **`A`** is array of two arrays of two ints.
- The type of **`B`** is adjusted to pointer to array of **`n`** ints, because it is a function parameter.
- The type of **`C`** is adjusted to pointer to int, because it is a function parameter.
- The type of **`D`** is array of two arrays of two ints.
- The type of **`E`** is array of **`n`** arrays of **`n`** ints.

Note that **`B`** and **`E`** involve variable length array types.

The new items of array type are initialized as if each integer element of the original array is assigned to the corresponding element of the new array. Those of pointer type are initialized as if by assignment from the original item to the new item.

Example carrays_fpriv.1.c

```
S-1   #include <assert.h>
S-2
S-3   int A[2][2] = {1, 2, 3, 4};
S-4
S-5   void f(int n, int B[n][n], int C[])
S-6   {
S-7     int D[2][2] = {1, 2, 3, 4};
S-8     int E[n][n];
S-9
S-10    assert(n >= 2);
S-11    E[1][1] = 4;
S-12
S-13    #pragma omp parallel firstprivate(B, C, D, E)
S-14    {
S-15      assert(sizeof(B) == sizeof(int (*)[n]));
S-16      assert(sizeof(C) == sizeof(int*));
S-17      assert(sizeof(D) == 4 * sizeof(int));
S-18      assert(sizeof(E) == n * n * sizeof(int));
S-19
S-20      /* Private B and C have values of original B and C. */
S-21      assert(&B[1][1] == &A[1][1]);
S-22      assert(&C[3] == &A[1][1]);
S-23      assert(D[1][1] == 4);
S-24      assert(E[1][1] == 4);
S-25    }
```

```
S-26    }
S-27
S-28    int main() {
S-29      f(2, A, A[0]);
S-30      return 0;
S-31    }
```

C / C++

7.8 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 in a **`lastprivate`** clause so that the values of the variables are the same as when the loop is executed sequentially.

C / C++

Example lastprivate.1.c

```
S-1   void lastpriv (int n, float *a, float *b)
S-2   {
S-3     int i;
S-4
S-5     #pragma omp parallel
S-6     {
S-7       #pragma omp for lastprivate(i)
S-8       for (i=0; i<n-1; i++)
S-9         a[i] = b[i] + b[i+1];
S-10    }
S-11
S-12    a[i]=b[i];       /* i == n-1 here */
S-13  }
```

C / C++

Fortran

Example lastprivate.1.f

```
S-1         SUBROUTINE LASTPRIV(N, A, B)
S-2
S-3           INTEGER N
S-4           REAL A(*), B(*)
S-5           INTEGER I
S-6   !$OMP PARALLEL
S-7   !$OMP DO LASTPRIVATE(I)
S-8
S-9           DO I=1,N-1
S-10            A(I) = B(I) + B(I+1)
S-11          ENDDO
S-12
S-13  !$OMP END PARALLEL
S-14          A(I) = B(I)      ! I has the value of N here
S-15
S-16        END SUBROUTINE LASTPRIV
```

Fortran

The next example illustrates the use of the **conditional** modifier in a **lastprivate** clause to return the last value when it may not come from the last iteration of a loop. That is, users can preserve the serial equivalence semantics of the loop. The conditional lastprivate ensures the final value of the variable after the loop is as if the loop iterations were executed in a sequential order.

C / C++

Example lastprivate.2.c (**omp_5.0**)

```
S-1   #include <math.h>
S-2
S-3   float condlastprivate(float *a, int n)
S-4   {
S-5      float x = 0.0f;
S-6
S-7      #pragma omp parallel for simd lastprivate(conditional: x)
S-8      for (int k = 0; k < n; k++) {
S-9         if (a[k] < 108.5 || a[k] > 208.5) {
S-10           x = sinf(a[k]);
S-11        }
S-12     }
S-13
S-14     return x;
S-15  }
```

C / C++

Fortran

Example lastprivate.2.f90 (**omp_5.0**)

```
S-1   function condlastprivate(a, n) result(x)
S-2      implicit none
S-3      real a(*), x
S-4      integer n, k
S-5
S-6      x = 0.0
S-7
S-8      !$omp parallel do simd lastprivate(conditional: x)
S-9      do k = 1, n
S-10        if (a(k) < 108.5 .or. a(k) > 208.5) then
S-11           x = sin(a(k))
S-12        endif
S-13     end do
S-14
S-15  end function condlastprivate
```

Fortran

7.9 Reduction

This section covers ways to perform reductions in parallel, task, taskloop, and SIMD regions.

7.9.1 The `reduction` Clause

The following example demonstrates the **`reduction`** clause; note that some reductions can be expressed in the loop in several ways, as shown for the **`max`** and **`min`** reductions below:

C / C++

Example reduction.1.c (**`omp_3.1`**)

```
S-1   #include <math.h>
S-2   void reduction1(float *x, int *y, int n)
S-3   {
S-4     int i, b, c;
S-5     float a, d;
S-6     a = 0.0;
S-7     b = 0;
S-8     c = y[0];
S-9     d = x[0];
S-10    #pragma omp parallel for private(i) shared(x, y, n) \
S-11                             reduction(+:a) reduction(^:b) \
S-12                             reduction(min:c) reduction(max:d)
S-13      for (i=0; i<n; i++) {
S-14        a += x[i];
S-15        b ^= y[i];
S-16        if (c > y[i]) c = y[i];
S-17        d = fmaxf(d,x[i]);
S-18      }
S-19  }
```

C / C++

Fortran

Example reduction.1.f90

```
S-1   SUBROUTINE REDUCTION1(A, B, C, D, X, Y, N)
S-2       REAL :: X(*), A, D
S-3       INTEGER :: Y(*), N, B, C
S-4       INTEGER :: I
S-5       A = 0
S-6       B = 0
S-7       C = Y(1)
S-8       D = X(1)
S-9       !$OMP PARALLEL DO PRIVATE(I) SHARED(X, Y, N) REDUCTION(+:A) &
S-10      !$OMP& REDUCTION(IEOR:B) REDUCTION(MIN:C)  REDUCTION(MAX:D)
S-11        DO I=1,N
S-12          A = A + X(I)
S-13          B = IEOR(B, Y(I))
S-14          C = MIN(C, Y(I))
S-15          IF (D < X(I)) D = X(I)
S-16        END DO
S-17
S-18  END SUBROUTINE REDUCTION1
```

Fortran

A common implementation of the preceding example is to treat it as if it had been written as follows:

C / C++

Example reduction.2.c

```
S-1   #include <limits.h>
S-2   #include <math.h>
S-3   void reduction2(float *x, int *y, int n)
S-4   {
S-5     int i, b, b_p, c, c_p;
S-6     float a, a_p, d, d_p;
S-7     a = 0.0f;
S-8     b = 0;
S-9     c = y[0];
S-10    d = x[0];
S-11    #pragma omp parallel shared(a, b, c, d, x, y, n) \
S-12                         private(a_p, b_p, c_p, d_p)
S-13    {
S-14      a_p = 0.0f;
S-15      b_p = 0;
S-16      c_p = INT_MAX;
S-17      d_p = -HUGE_VALF;
S-18      #pragma omp for private(i)
```

```
      for (i=0; i<n; i++) {
        a_p += x[i];
        b_p ^= y[i];
        if (c_p > y[i]) c_p = y[i];
        d_p = fmaxf(d_p,x[i]);
      }
      #pragma omp critical
      {
        a += a_p;
        b ^= b_p;
        if( c > c_p ) c = c_p;
        d = fmaxf(d,d_p);
      }
    }
}
```

C / C++

Fortran

Example reduction.2.f90

```
  SUBROUTINE REDUCTION2(A, B, C, D, X, Y, N)
    REAL :: X(*), A, D
    INTEGER :: Y(*), N, B, C
    REAL :: A_P, D_P
    INTEGER :: I, B_P, C_P
    A = 0
    B = 0
    C = Y(1)
    D = X(1)
    !$OMP PARALLEL SHARED(X, Y, A, B, C, D, N) &
    !$OMP&         PRIVATE(A_P, B_P, C_P, D_P)
      A_P = 0.0
      B_P = 0
      C_P = HUGE(C_P)
      D_P = -HUGE(D_P)
      !$OMP DO PRIVATE(I)
      DO I=1,N
        A_P = A_P + X(I)
        B_P = IEOR(B_P, Y(I))
        C_P = MIN(C_P, Y(I))
        IF (D_P < X(I)) D_P = X(I)
      END DO
      !$OMP CRITICAL
        A = A + A_P
        B = IEOR(B, B_P)
        C = MIN(C, C_P)
        D = MAX(D, D_P)
```

```
S-28          !$OMP END CRITICAL
S-29        !$OMP END PARALLEL
S-30      END SUBROUTINE REDUCTION2
```

The following program is non-conforming because the reduction is on the *intrinsic procedure name* **MAX** but that name has been redefined to be the variable named **MAX**.

Example reduction.3.f90

```
S-1     PROGRAM REDUCTION_WRONG
S-2     MAX = HUGE(0)
S-3     M = 0
S-4
S-5     !$OMP PARALLEL DO REDUCTION(MAX: M)
S-6    ! MAX is no longer the intrinsic so this is non-conforming
S-7     DO I = 1, 100
S-8        CALL SUB(M,I)
S-9     END DO
S-10
S-11    END PROGRAM REDUCTION_WRONG
S-12
S-13    SUBROUTINE SUB(M,I)
S-14       M = MAX(M,I)
S-15    END SUBROUTINE SUB
```

The following conforming program performs the reduction using the *intrinsic procedure name* **MAX** even though the intrinsic **MAX** has been renamed to **REN**.

Example reduction.4.f90

```
S-1     MODULE M
S-2        INTRINSIC MAX
S-3     END MODULE M
S-4
S-5     PROGRAM REDUCTION3
S-6        USE M, REN => MAX
S-7        N = 0
S-8     !$OMP PARALLEL DO REDUCTION(REN: N)      ! still does MAX
S-9        DO I = 1, 100
S-10          N = MAX(N,I)
S-11       END DO
S-12    END PROGRAM REDUCTION3
```

The following conforming program performs the reduction using *intrinsic procedure name* **MAX** even though the intrinsic **MAX** has been renamed to **MIN**.

Example reduction.5.f90

```
S-1     MODULE MOD
S-2        INTRINSIC MAX, MIN
S-3     END MODULE MOD
S-4
S-5     PROGRAM REDUCTION4
S-6        USE MOD, MIN=>MAX, MAX=>MIN
S-7        REAL :: R
S-8        R = -HUGE(0.0)
S-9
S-10    !$OMP PARALLEL DO REDUCTION(MIN: R)      ! still does MAX
S-11       DO I = 1, 1000
S-12          R = MIN(R, SIN(REAL(I)))
S-13       END DO
S-14       PRINT *, R
S-15    END PROGRAM REDUCTION4
```

Fortran

The following example is non-conforming because the initialization (**a = 0**) of the original list item **a** is not synchronized with the update of **a** as a result of the reduction computation in the **for** loop. Therefore, the example may print an incorrect value for **a**.

To avoid this problem, the initialization of the original list item **a** should complete before any update of **a** as a result of the **reduction** clause. This can be achieved by adding an explicit barrier after the assignment **a = 0**, or by enclosing the assignment **a = 0** in a **single** directive (which has an implied barrier), or by initializing **a** before the start of the **parallel** region.

C / C++

Example reduction.6.c

```
S-1   #include <stdio.h>
S-2
S-3   int main (void)
S-4   {
S-5     int a, i;
S-6
S-7     #pragma omp parallel shared(a) private(i)
S-8     {
S-9       #pragma omp master
S-10      a = 0;
S-11
S-12      // To avoid race conditions, add a barrier here.
S-13
S-14      #pragma omp for reduction(+:a)
S-15      for (i = 0; i < 10; i++) {
S-16          a += i;
S-17      }
S-18
S-19      #pragma omp single
S-20      printf ("Sum is %d\n", a);
S-21    }
S-22    return 0;
S-23  }
```

C / C++

Fortran

Example reduction.6.f

```
      INTEGER A, I

!$OMP PARALLEL SHARED(A) PRIVATE(I)

!$OMP MASTER
      A = 0
!$OMP END MASTER

      ! To avoid race conditions, add a barrier here.

!$OMP DO REDUCTION(+:A)
      DO I= 0, 9
         A = A + I
      END DO

!$OMP SINGLE
      PRINT *, "Sum is ", A
!$OMP END SINGLE

!$OMP END PARALLEL
      END
```

Fortran

The following example demonstrates the reduction of array *a*. In C/C++ this is illustrated by the explicit use of an array section *a[0:N]* in the **reduction** clause. The corresponding Fortran example uses array syntax supported in the base language. As of the OpenMP 4.5 specification the explicit use of array section in the **reduction** clause in Fortran is not permitted. But this oversight has been fixed in the OpenMP 5.0 specification.

C / C++

Example reduction.7.c (**omp_4.5**)

```
#include <stdio.h>

#define N 100
void init(int n, float (*b)[N]);

int main(){

  int i,j;
  float a[N], b[N][N];

  init(N,b);
```

```
S-12
S-13      for(i=0; i<N; i++) a[i]=0.0e0;
S-14
S-15      #pragma omp parallel for reduction(+:a[0:N]) private(j)
S-16      for(i=0; i<N; i++){
S-17        for(j=0; j<N; j++){
S-18           a[j] +=  b[i][j];
S-19        }
S-20      }
S-21      printf(" a[0] a[N-1]: %f %f\n", a[0], a[N-1]);
S-22
S-23      return 0;
S-24    }
```

C / C++

Fortran

1 *Example reduction.7.f90*

```
S-1     program array_red
S-2
S-3       integer,parameter :: n=100
S-4       integer           :: j
S-5       real              :: a(n), b(n,n)
S-6
S-7       call init(n,b)
S-8
S-9       a(:) = 0.0e0
S-10
S-11      !$omp parallel do reduction(+:a)
S-12      do j = 1, n
S-13         a(:) = a(:) + b(:,j)
S-14      end do
S-15
S-16      print*, " a(1) a(n): ", a(1), a(n)
S-17
S-18    end program
```

Fortran

7.9.2 Task Reduction

In OpenMP 5.0 the **task_reduction** clause was created for the **taskgroup** construct, to allow reductions among explicit tasks that have an **in_reduction** clause.

In the *task_reduction.1* example below a reduction is performed as the algorithm traverses a linked list. The reduction statement is assigned to be an explicit task using a **task** construct and is specified to be a reduction participant with the **in_reduction** clause. A **taskgroup** construct encloses the tasks participating in the reduction, and specifies, with the **task_reduction** clause, that the taskgroup has tasks participating in a reduction. After the **taskgroup** region the original variable will contain the final value of the reduction.

Note: The *res* variable is private in the *linked_list_sum* routine and is not required to be shared (as in the case of a **parallel** construct reduction).

C / C++

Example task_reduction.1.c (**omp_5.0**)

```
S-1   #include<stdlib.h>
S-2   #include<stdio.h>
S-3   #define N 10
S-4
S-5   typedef struct node_tag {
S-6       int val;
S-7       struct node_tag *next;
S-8   } node_t;
S-9
S-10  int linked_list_sum(node_t *p)
S-11  {
S-12      int res = 0;
S-13
S-14      #pragma omp taskgroup task_reduction(+: res)
S-15      {
S-16          node_t* aux = p;
S-17          while(aux != 0)
S-18          {
S-19              #pragma omp task in_reduction(+: res)
S-20              res += aux->val;
S-21
S-22              aux = aux->next;
S-23          }
S-24      }
S-25      return res;
S-26  }
S-27
S-28
S-29  int main(int argc, char *argv[]) {
S-30      int i;
```

```
S-31    //                              Create the root node.
S-32        node_t* root = (node_t*) malloc(sizeof(node_t));
S-33        root->val = 1;
S-34
S-35        node_t* aux = root;
S-36
S-37    //                              Create N-1 more nodes.
S-38        for(i=2;i<=N;++i){
S-39            aux->next = (node_t*) malloc(sizeof(node_t));
S-40            aux = aux->next;
S-41            aux->val = i;
S-42        }
S-43
S-44        aux->next = 0;
S-45
S-46        #pragma omp parallel
S-47        #pragma omp single
S-48        {
S-49            int result = linked_list_sum(root);
S-50            printf( "Calculated: %d  Analytic:%d\n", result, (N*(N+1)/2) );
S-51        }
S-52
S-53        return 0;
S-54    }
S-55
```

C / C++

Fortran

1 *Example task_reduction.1.f90* (`omp_5.0`)

```
S-1     module m
S-2         type node_t
S-3             integer :: val
S-4             type(node_t), pointer :: next
S-5         end type
S-6     end module m
S-7
S-8     function linked_list_sum(p) result(res)
S-9         use m
S-10        implicit none
S-11        type(node_t), pointer :: p
S-12        type(node_t), pointer :: aux
S-13        integer :: res
S-14
S-15        res = 0
S-16
S-17        !$omp taskgroup task_reduction(+: res)
```

```
S-18            aux => p
S-19            do while (associated(aux))
S-20                !$omp task in_reduction(+: res)
S-21                    res = res + aux%val
S-22                !$omp end task
S-23                aux => aux%next
S-24            end do
S-25        !$omp end taskgroup
S-26    end function linked_list_sum
S-27
S-28
S-29    program main
S-30        use m
S-31        implicit none
S-32        type(node_t), pointer :: root, aux
S-33        integer :: res, i
S-34        integer, parameter :: N=10
S-35
S-36        interface
S-37            function linked_list_sum(p) result(res)
S-38                use m
S-39                implicit none
S-40                type(node_t), pointer :: p
S-41                integer :: res
S-42            end function
S-43        end interface
S-44    !                       Create the root node.
S-45        allocate(root)
S-46        root%val = 1
S-47        aux => root
S-48
S-49    !                       Create N-1 more nodes.
S-50        do i = 2,N
S-51            allocate(aux%next)
S-52            aux => aux%next
S-53            aux%val = i
S-54        end do
S-55
S-56        aux%next => null()
S-57
S-58        !$omp parallel
S-59        !$omp single
S-60            res = linked_list_sum(root)
S-61            print *, "Calculated:", res, " Analytic:", (N*(N+1))/2
S-62        !$omp end single
S-63        !$omp end parallel
S-64
```

```
S-65    end program main
S-66
```

Fortran

In OpenMP 5.0 the **task** *reduction-modifier* for the **reduction** clause was introduced to provide a means of performing reductions among implicit and explicit tasks.

The **reduction** clause of a **parallel** or worksharing construct may specify the **task** *reduction-modifier* to include explicit task reductions within their region, provided the reduction operators (*reduction-identifiers*) and variables (*list items*) of the participating tasks match those of the implicit tasks.

There are 2 reduction use cases (identified by USE CASE #) in the *task_reduction.2* example below.

In USE CASE 1 a **task** modifier in the **reduction** clause of the **parallel** construct is used to include the reductions of any participating tasks, those with an **in_reduction** clause and matching *reduction-identifiers* (**+**) and list items (**x**).

Note, a **taskgroup** construct (with a **task_reduction** clause) in not necessary to scope the explicit task reduction (as seen in the example above). Hence, even without the implicit task reduction statement (without the C **x++** and Fortran **x=x+1** statements), the **task** *reduction-modifier* in a **reduction** clause of the **parallel** construct can be used to avoid having to create a **taskgroup** construct (and its **task_reduction** clause) around the task generating structure.

In USE CASE 2 tasks participating in the reduction are within a worksharing region (a parallel worksharing-loop construct). Here, too, no **taskgroup** is required, and the *reduction-identifier* (**+**) and list item (variable **x**) match as required.

C / C++

Example task_reduction.2.c (**omp_5.0**)

```
S-1     #include <stdio.h>
S-2     int main(void){
S-3        int N=100, M=10;
S-4        int i, x;
S-5
S-6     // USE CASE 1  explicit-task reduction + parallel reduction clause
S-7        x=0;
S-8        #pragma omp parallel num_threads(M) reduction(task,+:x)
S-9        {
S-10
S-11         x++;                 // implicit task reduction statement
S-12
S-13         #pragma omp single
S-14         for(i=0;i<N;i++)
S-15           #pragma omp task in_reduction(+:x)
S-16           x++;
```

```
   }
   printf("x=%d  =M+N\n",x);  // x= 110  =M+N


// USE CASE 2  task reduction +  worksharing reduction clause
   x=0;
   #pragma omp parallel for num_threads(M) reduction(task,+:x)
   for(i=0; i< N; i++){

      x++;

      if( i%2 == 0){
       #pragma omp task in_reduction(+:x)
       x--;
      }
   }
   printf("x=%d  =N-N/2\n",x);  // x= 50  =N-N/2

   return 0;
}
```

C / C++

Fortran

Example task_reduction.2.f90 (`omp_5.0`)

```
program task_modifier

   integer :: N=100, M=10
   integer :: i, x

! USE CASE 1  explicit-task reduction + parallel reduction clause
   x=0
   !$omp parallel num_threads(M) reduction(task,+:x)

     x=x+1                    !! implicit task reduction statement

     !$omp single
       do i = 1,N
         !$omp task in_reduction(+:x)
           x=x+1
         !$omp end task
       end do
     !$omp end single

   !$omp end parallel
```

```
S-22        write(*,'("x=",I0," =M+N")') x    ! x= 110 =M+N
S-23
S-24
S-25     ! USE CASE 2  task reduction +  worksharing reduction clause
S-26        x=0
S-27        !$omp parallel do num_threads(M) reduction(task,+:x)
S-28          do i = 1,N
S-29
S-30             x=x+1
S-31
S-32             if( mod(i,2) == 0) then
S-33                !$omp task in_reduction(+:x)
S-34                  x=x-1
S-35                !$omp end task
S-36             endif
S-37
S-38          end do
S-39        write(*,'("x=",I0,"  =N-N/2")') x    ! x= 50 =N-N/2
S-40
S-41     end program
```

Fortran

7.9.3 Reduction on Combined Target Constructs

When a **reduction** clause appears on a combined construct that combines a **target** construct with another construct, there is an implicit map of the list items with a **tofrom** map type for the **target** construct. Otherwise, the list items (if they are scalar variables) would be treated as firstprivate by default in the **target** construct, which is unlikely to provide the intended behavior since the result of the reduction that is in the firstprivate variable would be discarded at the end of the **target** region.

In the following example, the use of the **reduction** clause on **sum1** or **sum2** should, by default, result in an implicit **tofrom** map for that variable. So long as neither **sum1** nor **sum2** were already present on the device, the mapping behavior ensures the value for **sum1** computed in the first **target** construct is used in the second **target** construct.

C / C++

Example target_reduction.1.c (`omp_5.0`)

```
S-1   #include <stdio.h>
S-2   int f(int);
S-3   int g(int);
S-4   int main()
S-5   {
S-6      int sum1=0, sum2=0;
S-7      int i;
S-8      const int n = 100;
S-9
S-10     #pragma omp target teams distribute reduction(+:sum1)
S-11     for (int i = 0; i < n; i++) {
S-12        sum1 += f(i);
S-13     }
S-14
S-15     #pragma omp target teams distribute reduction(+:sum2)
S-16     for (int i = 0; i < n; i++) {
S-17        sum2 += g(i) * sum1;
S-18     }
S-19
S-20     printf(  "sum1 = %d, sum2 = %d\n", sum1, sum2);
S-21     //OUTPUT: sum1 = 9900, sum2 = 147015000
S-22     return 0;
S-23  }
S-24
S-25  int f(int res){ return res*2; }
S-26  int g(int res){ return res*3; }
```

C / C++

Fortran

Example target_reduction.1.f90 (`omp_5.0`)

```
S-1   program target_reduction_ex1
S-2      interface
S-3         function f(res)
S-4                integer :: f, res
S-5             end function
S-6         function g(res)
S-7                integer :: g, res
S-8             end function
S-9      end interface
S-10     integer :: sum1, sum2, i
S-11     integer, parameter :: n = 100
S-12     sum1 = 0
S-13     sum2 = 0
```

```
S-14        !$omp target teams distribute reduction(+:sum1)
S-15            do i=1,n
S-16               sum1 = sum1 + f(i)
S-17            end do
S-18        !$omp target teams distribute reduction(+:sum2)
S-19            do i=1,n
S-20               sum2 = sum2 + g(i)*sum1
S-21            end do
S-22        print *, "sum1 = ", sum1, ", sum2 = ", sum2
S-23        !!OUTPUT: sum1 =     10100 , sum2 = 153015000
S-24     end program
S-25
S-26
S-27     integer function f(res)
S-28        integer :: res
S-29        f = res*2
S-30     end function
S-31     integer function g(res)
S-32        integer :: res
S-33        g = res*3
S-34     end function
```

Fortran

In next example, the variables **sum1** and **sum2** remain on the device for the duration of the **target data** region so that it is their device copies that are updated by the reductions. Note the significance of mapping **sum1** on the second **target** construct; otherwise, it would be treated by default as firstprivate and the result computed for **sum1** in the prior **target** region may not be used. Alternatively, a **target update** construct could be used between the two **target** constructs to update the host version of **sum1** with the value that is in the corresponding device version after the completion of the first construct.

C / C++

Example target_reduction.2.c (**omp_5.0**)

```
#include <stdio.h>
int f(int);
int g(int);
int main()
{
   int sum1=0, sum2=0;
   int i;
   const int n = 100;

   #pragma omp target data map(sum1,sum2)
   {
      #pragma omp target teams distribute reduction(+:sum1)
      for (int i = 0; i < n; i++) {
         sum1 += f(i);
      }

      #pragma omp target teams distribute map(sum1) reduction(+:sum2)
      for (int i = 0; i < n; i++) {
         sum2 += g(i) * sum1;
      }
   }
   printf(  "sum1 = %d, sum2 = %d\n", sum1, sum2);
   //OUTPUT: sum1 = 9900, sum2 = 147015000
   return 0;
}

int f(int res){ return res*2; }
int g(int res){ return res*3; }
```

C / C++

Fortran

Example target_reduction.2.f90 (`omp_5.0`)

```
S-1
S-2     program target_reduction_ex2
S-3        interface
S-4           function f(res)
S-5                  integer :: f, res
S-6              end function
S-7           function g(res)
S-8                  integer :: g, res
S-9              end function
S-10       end interface
S-11       integer :: sum1, sum2, i
S-12       integer, parameter :: n = 100
S-13       sum1 = 0
S-14       sum2 = 0
S-15       !$omp target data map(sum1, sum2)
S-16           !$omp target teams distribute reduction(+:sum1)
S-17               do i=1,n
S-18                  sum1 = sum1 + f(i)
S-19               end do
S-20           !$omp target teams distribute map(sum1) reduction(+:sum2)
S-21               do i=1,n
S-22                  sum2 = sum2 + g(i)*sum1
S-23               end do
S-24       !$omp end target data
S-25       print *, "sum1 = ", sum1, ", sum2 = ", sum2
S-26       !!OUTPUT: sum1 =     10100 , sum2 = 153015000
S-27    end program
S-28
S-29
S-30    integer function f(res)
S-31       integer :: res
S-32       f = res*2
S-33    end function
S-34    integer function g(res)
S-35       integer :: res
S-36       g = res*3
S-37    end function
```

Fortran

7.9.4 Task Reduction with Target Constructs

The following examples illustrate how task reductions can apply to target tasks that result from a **target** construct with the **in_reduction** clause. Here, the **in_reduction** clause specifies that the target task participates in the task reduction defined in the scope of the enclosing **taskgroup** construct. Partial results from all tasks participating in the task reduction will be combined (in some order) into the original variable listed in the **task_reduction** clause before exiting the **taskgroup** region.

C / C++

Example target_task_reduction.1.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   #pragma omp declare target to(device_compute)
S-3   void device_compute(int *);
S-4   void host_compute(int *);
S-5   int main()
S-6   {
S-7      int sum = 0;
S-8
S-9      #pragma omp parallel master
S-10     #pragma omp taskgroup task_reduction(+:sum)
S-11     {
S-12        #pragma omp target in_reduction(+:sum) nowait
S-13            device_compute(&sum);
S-14
S-15        #pragma omp task in_reduction(+:sum)
S-16            host_compute(&sum);
S-17     }
S-18     printf(  "sum = %d\n", sum);
S-19     //OUTPUT: sum = 2
S-20     return 0;
S-21  }
S-22
S-23  void device_compute(int *sum){ *sum = 1; }
S-24  void   host_compute(int *sum){ *sum = 1; }
```

C / C++

Fortran

Example target_task_reduction.1.f90 (`omp_5.0`)

```
S-1   program target_task_reduction_ex1
S-2      interface
S-3         subroutine device_compute(res)
S-4         !$omp declare target to(device_compute)
S-5           integer :: res
S-6         end subroutine device_compute
S-7         subroutine host_compute(res)
S-8           integer :: res
S-9         end subroutine host_compute
S-10     end interface
S-11     integer :: sum
S-12     sum = 0
S-13     !$omp parallel master
S-14        !$omp taskgroup task_reduction(+:sum)
S-15           !$omp target in_reduction(+:sum) nowait
S-16              call device_compute(sum)
S-17           !$omp end target
S-18           !$omp task in_reduction(+:sum)
S-19              call host_compute(sum)
S-20           !$omp end task
S-21        !$omp end taskgroup
S-22     !$omp end parallel master
S-23     print *, "sum = ", sum
S-24     !!OUTPUT: sum = 2
S-25  end program
S-26
S-27  subroutine device_compute(sum)
S-28     integer :: sum
S-29     sum = 1
S-30  end subroutine
S-31  subroutine host_compute(sum)
S-32     integer :: sum
S-33     sum = 1
S-34  end subroutine
```

Fortran

In the next pair of examples, the task reduction is defined by a **`reduction`** clause with the **`task`** modifier, rather than a **`task_reduction`** clause on a **`taskgroup`** construct. Again, the partial results from the participating tasks will be combined in some order into the original reduction variable, **`sum`**.

C / C++

Example target_task_reduction.2a.c (`omp_5.0`)

```
S-1   #include <stdio.h>
S-2   #pragma omp declare target to(device_compute)
S-3   extern void device_compute(int *);
S-4   extern void host_compute(int *);
S-5   int main()
S-6   {
S-7      int sum = 0;
S-8
S-9      #pragma omp parallel sections reduction(task, +:sum)
S-10     {
S-11        #pragma omp section
S-12            {
S-13               #pragma omp target in_reduction(+:sum)
S-14               device_compute(&sum);
S-15            }
S-16        #pragma omp section
S-17            {
S-18               host_compute(&sum);
S-19            }
S-20     }
S-21     printf(  "sum = %d\n", sum);
S-22     //OUTPUT: sum = 2
S-23     return 0;
S-24  }
S-25
S-26  void device_compute(int *sum){ *sum = 1; }
S-27  void   host_compute(int *sum){ *sum = 1; }
```

C / C++

Fortran

Example target_task_reduction.2a.f90 (`omp_5.0`)

```
S-1   program target_task_reduction_ex2
S-2      interface
S-3         subroutine device_compute(res)
S-4         !$omp declare target to(device_compute)
S-5           integer :: res
S-6         end subroutine device_compute
S-7         subroutine host_compute(res)
S-8           integer :: res
S-9         end subroutine host_compute
S-10     end interface
S-11     integer :: sum
S-12     sum = 0
```

```
S-13        !$omp parallel sections reduction(task,+:sum)
S-14           !$omp section
S-15              !$omp target in_reduction(+:sum) nowait
S-16                call device_compute(sum)
S-17              !$omp end target
S-18           !$omp section
S-19              call host_compute(sum)
S-20        !$omp end parallel sections
S-21        print *, "sum = ", sum
S-22        !!OUTPUT: sum = 2
S-23     end program
S-24
S-25     subroutine device_compute(sum)
S-26        integer :: sum
S-27        sum = 1
S-28     end subroutine
S-29     subroutine host_compute(sum)
S-30        integer :: sum
S-31        sum = 1
S-32     end subroutine
```

Fortran

Next, the **task** modifier is again used to define a task reduction over participating tasks. This time, the participating tasks are a target task resulting from a **target** construct with the **in_reduction** clause, and the implicit task (executing on the master thread) that calls **host_compute**. As before, the partial results from these paricipating tasks are combined in some order into the original reduction variable.

C / C++

Example target_task_reduction.2b.c (**omp_5.0**)

```
S-1     #include <stdio.h>
S-2     #pragma omp declare target to(device_compute)
S-3     extern void device_compute(int *);
S-4     extern void host_compute(int *);
S-5     int main()
S-6     {
S-7        int sum = 0;
S-8
S-9        #pragma omp parallel master reduction(task, +:sum)
S-10       {
S-11           #pragma omp target in_reduction(+:sum) nowait
S-12           device_compute(&sum);
S-13
S-14           host_compute(&sum);
S-15       }
S-16       printf(  "sum = %d\n", sum);
```

```
S-17        //OUTPUT: sum = 2
S-18        return 0;
S-19     }
S-20
S-21     void device_compute(int *sum){ *sum = 1; }
S-22     void   host_compute(int *sum){ *sum = 1; }
```

C / C++

Fortran

Example target_task_reduction.2b.f90 (`omp_5.0`)

```
S-1      program target_task_reduction_ex2b
S-2         interface
S-3            subroutine device_compute(res)
S-4            !$omp declare target to(device_compute)
S-5              integer :: res
S-6            end subroutine device_compute
S-7            subroutine host_compute(res)
S-8              integer :: res
S-9            end subroutine host_compute
S-10        end interface
S-11        integer :: sum
S-12        sum = 0
S-13        !$omp parallel master reduction(task,+:sum)
S-14              !$omp target in_reduction(+:sum) nowait
S-15                call device_compute(sum)
S-16              !$omp end target
S-17              call host_compute(sum)
S-18        !$omp end parallel sections
S-19        print *, "sum = ", sum
S-20        !!OUTPUT: sum = 2
S-21     end program
S-22
S-23
S-24     subroutine device_compute(sum)
S-25        integer :: sum
S-26        sum = 1
S-27     end subroutine
S-28     subroutine host_compute(sum)
S-29        integer :: sum
S-30        sum = 1
S-31     end subroutine
S-32
```

Fortran

7.9.5 Taskloop Reduction

In the OpenMP 5.0 Specification the **taskloop** construct was extended to include the reductions.

The following two examples show how to implement a reduction over an array using taskloop reduction in two different ways. In the first example we apply the **reduction** clause to the **taskloop** construct. As it was explained above in the task reduction examples, a reduction over tasks is divided in two components: the scope of the reduction, which is defined by a **taskgroup** region, and the tasks that participate in the reduction. In this example, the **reduction** clause defines both semantics. First, it specifies that the implicit **taskgroup** region associated with the **taskloop** construct is the scope of the reduction, and second, it defines all tasks created by the **taskloop** construct as participants of the reduction. About the first property, it is important to note that if we add the **nogroup** clause to the **taskloop** construct the code will be nonconforming, basically because we have a set of tasks that participate in a reduction that has not been defined.

C / C++

Example taskloop_reduction.1.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2
S-3   int array_sum(int n, int *v) {
S-4       int i;
S-5       int res = 0;
S-6
S-7       #pragma omp taskloop reduction(+: res)
S-8       for(i = 0; i < n; ++i)
S-9           res += v[i];
S-10
S-11      return res;
S-12  }
S-13
S-14  int main(int argc, char *argv[]) {
S-15      int n = 10;
S-16      int v[10] = {1,2,3,4,5,6,7,8,9,10};
S-17
S-18      #pragma omp parallel
S-19      #pragma omp single
S-20      {
S-21          int res = array_sum(n, v);
S-22          printf("The result is %d\n", res);
S-23      }
S-24      return 0;
S-25  }
```

C / C++

Fortran

Example taskloop_reduction.1.f90 (**omp_5.0**)

```
S-1
S-2     function array_sum(n, v) result(res)
S-3         implicit none
S-4         integer :: n, v(n), res
S-5         integer :: i
S-6
S-7         res = 0
S-8         !$omp taskloop reduction(+: res)
S-9         do i=1, n
S-10            res = res + v(i)
S-11        end do
S-12        !$omp end taskoop
S-13
S-14    end function array_sum
S-15
S-16    program main
S-17        implicit none
S-18        integer :: n, v(10), res
S-19        integer :: i
S-20
S-21        integer, external :: array_sum
S-22
S-23        n = 10
S-24        do i=1, n
S-25            v(i) = i
S-26        end do
S-27
S-28        !$omp parallel
S-29        !$omp single
S-30        res = array_sum(n, v)
S-31        print *, "The result is", res
S-32        !$omp end single
S-33        !$omp end parallel
S-34    end program main
```

Fortran

The second example computes exactly the same value as in the preceding*taskloop_reduction.1* code section, but in a very different way. First, in the *array_sum* function a **taskgroup** region is created that defines the scope of a new reduction using the **task_reduction** clause. After that, a task and also the tasks generated by a taskloop participate in that reduction by using the **in_reduction** clause on the **task** and **taskloop** constructs, respectively. Note that the **nogroup** clause was added to the **taskloop** construct. This is allowed because what is expressed with the **in_reduction** clause is different from what is expressed with the

reduction clause. In one case the generated tasks are specified to participate in a previously declared reduction (**in_reduction** clause) whereas in the other case creation of a new reduction is specified and also that all tasks generated by the taskloop will participate on it.

C / C++

Example taskloop_reduction.2.c (**omp_5.0**)

```
#include <stdio.h>

int array_sum(int n, int *v) {
    int i;
    int res = 0;

    #pragma omp taskgroup task_reduction(+: res)
    {
        if (n > 0) {
            #pragma omp task in_reduction(+: res)
            res = res + v[0];

            #pragma omp taskloop in_reduction(+: res) nogroup
            for(i = 1; i < n; ++i)
                res += v[i];
        }
    }

    return res;
}

int main(int argc, char *argv[]) {
    int n = 10;
    int v[10] = {1,2,3,4,5,6,7,8,9,10};

    #pragma omp parallel
    #pragma omp single
    {
        int res = array_sum(n, v);
        printf("The result is %d\n", res);
    }
    return 0;
}
```

C / C++

Fortran

Example taskloop_reduction.2.f90 (`omp_5.0`)

```
S-1
S-2     function array_sum(n, v) result(res)
S-3         implicit none
S-4         integer :: n, v(n), res
S-5         integer :: i
S-6
S-7         res = 0
S-8         !$omp taskgroup task_reduction(+: res)
S-9         if (n > 0) then
S-10            !$omp task in_reduction(+: res)
S-11            res = res + v(1)
S-12            !$omp end task
S-13
S-14            !$omp taskloop in_reduction(+: res) nogroup
S-15            do i=2, n
S-16                res = res + v(i)
S-17            end do
S-18            !$omp end taskoop
S-19        endif
S-20        !$omp end taskgroup
S-21
S-22    end function array_sum
S-23
S-24    program main
S-25        implicit none
S-26        integer :: n, v(10), res
S-27        integer :: i
S-28
S-29        integer, external :: array_sum
S-30
S-31        n = 10
S-32        do i=1, n
S-33            v(i) = i
S-34        end do
S-35
S-36        !$omp parallel
S-37        !$omp single
S-38        res = array_sum(n, v)
S-39        print *, "The result is", res
S-40        !$omp end single
S-41        !$omp end parallel
S-42    end program main
```

Fortran

In the OpenMP 5.0 Specification, **reduction** clauses for the **taskloop simd** construct were also added.

The examples below compare reductions for the **taskloop** and the **taskloop simd** constructs. These examples illustrate the use of **reduction** clauses within "stand-alone" **taskloop** constructs, and the use of **in_reduction** clauses for tasks of taskloops to participate with other reductions within the scope of a parallel region.

taskloop reductions:

In the *taskloop reductions* section of the example below, *taskloop 1* uses the **reduction** clause in a **taskloop** construct for a sum reduction, accumulated in *asum*. The behavior is as though a **taskgroup** construct encloses the taskloop region with a **task_reduction** clause, and each taskloop task has an **in_reduction** clause with the specifications of the **reduction** clause. At the end of the taskloop region *asum* contains the result of the reduction.

The next taskloop, *taskloop 2*, illustrates the use of the **in_reduction** clause to participate in a previously defined reduction scope of a **parallel** construct.

The task reductions of *task 2* and *taskloop 2* are combined across the **taskloop** construct and the single **task** construct, as specified in the **reduction(task, +:asum)** clause of the **parallel** construct. At the end of the parallel region *asum* contains the combined result of all reductions.

taskloop simd reductions:

Reductions for the **taskloop simd** construct are shown in the second half of the code. Since each component construct, **taskloop** and **simd**, can accept a reduction-type clause, the **taskloop simd** construct is a composite construct, and the specific application of the reduction clause is defined within the **taskloop simd** construct section of the OpenMP 5.0 Specification. The code below illustrates use cases for these reductions.

In the *taskloop simd reduction* section of the example below, *taskloop simd 3* uses the **reduction** clause in a **taskloop simd** construct for a sum reduction within a loop. For this case a **reduction** clause is used, as one would use for a **simd** construct. The SIMD reductions of each task are combined, and the results of these tasks are further combined just as in the **taskloop** construct with the **reduction** clause for *taskloop 1*. At the end of the taskloop region *asum* contains the combined result of all reductions.

If a **taskloop simd** construct is to participate in a previously defined reduction scope, the reduction participation should be specified with a **in_reduction** clause, as shown in the **parallel** region enclosing *task 4* and *taskloop simd 4* code sections.

Here the **taskloop simd** construct's **in_reduction** clause specifies participation of the construct's tasks as a task reduction within the scope of the parallel region. That is, the results of each task of the **taskloop** construct component contribute to the reduction in a broader level, just as in *parallel reduction a* code section above. Also, each **simd**-component construct occurs as if it has a **reduction** clause, and the SIMD results of each task are combined as though to form a

single result for each task (that participates in the **in_reduction** clause). At the end of the parallel region *asum* contains the combined result of all reductions.

C / C++

Example taskloop_simd_reduction.1.c (**omp_5.0**)

```
S-1
S-2   #include <stdio.h>
S-3   #define N 100
S-4
S-5   int main(){
S-6     int i, a[N], asum=0;
S-7
S-8     for(i=0;i<N;i++) a[i]=i;
S-9
S-10  // taskloop reductions
S-11
S-12    #pragma omp parallel master
S-13    #pragma omp taskloop reduction(+:asum)                //taskloop 1
S-14      for(i=0;i<N;i++){ asum += a[i]; }
S-15
S-16
S-17    #pragma omp parallel reduction(task, +:asum)          // parallel reduction a
S-18    {
S-19       #pragma omp master
S-20       #pragma omp task            in_reduction(+:asum)    //task 2
S-21         for(i=0;i<N;i++){ asum += a[i]; }
S-22
S-23       #pragma omp master taskloop in_reduction(+:asum)    //taskloop 2
S-24         for(i=0;i<N;i++){ asum += a[i]; }
S-25    }
S-26
S-27  // taskloop simd reductions
S-28
S-29    #pragma omp parallel master
S-30    #pragma omp taskloop simd reduction(+:asum)            //taskloop simd 3
S-31      for(i=0;i<N;i++){ asum += a[i]; }
S-32
S-33
S-34    #pragma omp parallel reduction(task, +:asum)          // parallel reduction b
S-35    {
S-36       #pragma omp master
S-37       #pragma omp task                 in_reduction(+:asum) //task 4
S-38         for(i=0;i<N;i++){ asum += a[i]; }
S-39
S-40       #pragma omp master taskloop simd in_reduction(+:asum) //taskloop simd 4
S-41         for(i=0;i<N;i++){ asum += a[i]; }
```

```
S-42
S-43      }
S-44
S-45      printf("asum=%d \n",asum); // output: asum=29700
S-46    }
```

C / C++

Fortran

Example taskloop_simd_reduction.1.f90 (`omp_5.0`)

```
 S-1
 S-2    program main
 S-3
 S-4      use omp_lib
 S-5      integer, parameter ::  N=100
 S-6      integer            :: i, a(N), asum=0
 S-7
 S-8      a = [( i, i=1,N )]    !! initialize
 S-9
S-10    !! taskloop reductions
S-11
S-12      !$omp parallel master
S-13      !$omp taskloop reduction(+:asum)                    !! taskloop 1
S-14        do i=1,N;  asum = asum + a(i);  enddo
S-15      !$omp end taskloop
S-16      !$omp end parallel master
S-17
S-18
S-19      !$omp parallel reduction(task, +:asum)              !! parallel reduction a
S-20
S-21         !$omp master
S-22         !$omp task            in_reduction(+:asum)       !! task 2
S-23           do i=1,N;  asum = asum + a(i);  enddo
S-24         !$omp end task
S-25         !$omp end master
S-26
S-27         !$omp master taskloop in_reduction(+:asum)       !! taskloop 2
S-28           do i=1,N;  asum = asum + a(i);  enddo
S-29         !$omp end master taskloop
S-30
S-31      !$omp end parallel
S-32
S-33    !! taskloop simd reductions
S-34
S-35      !$omp parallel master
S-36      !$omp taskloop simd reduction(+:asum)               !! taskloop simd 3
S-37        do i=1,N;  asum = asum + a(i);  enddo
```

```
S-38      !$omp end taskloop simd
S-39      !$omp end parallel master
S-40
S-41
S-42      !$omp parallel reduction(task, +:asum)                !! parallel reduction b
S-43
S-44        !$omp master
S-45        !$omp task                  in_reduction(+:asum)    !! task 4
S-46           do i=1,N;  asum = asum + a(i);  enddo
S-47        !$omp end task
S-48        !$omp end master
S-49
S-50        !$omp master taskloop simd in_reduction(+:asum)    !! taskloop simd 4
S-51           do i=1,N;  asum = asum + a(i);  enddo
S-52        !$omp end master taskloop simd
S-53
S-54      !$omp end parallel
S-55
S-56      print*,"asum=",asum   !! output: asum=30300
S-57
S-58    end program
```

Fortran

7.9.6 User-Defined Reduction

The **declare reduction** directive can be used to specify user-defined reductions (UDR) for user data types.

In the following example, **declare reduction** directives are used to define *min* and *max* operations for the *point* data structure for computing the rectangle that encloses a set of 2-D points.

Each **declare reduction** directive defines new reduction identifiers, *min* and *max*, to be used in a **reduction** clause. The next item in the declaration list is the data type (*struct point*) used in the reduction, followed by the combiner, here the functions *minproc* and *maxproc* perform the min and max operations, respectively, on the user data (of type *struct point*). In the function argument list are two special OpenMP variable identifiers, **omp_in** and **omp_out**, that denote the two values to be combined in the "real" function; the **omp_out** identifier indicates which one is to hold the result.

The initializer of the **declare reduction** directive specifies the initial value for the private variable of each implicit task. The **omp_priv** identifier is used to denote the private variable.

C / C++

Example udr.1.c (`omp_4.0`)

```
S-1   #include <stdio.h>
S-2   #include <limits.h>
S-3
S-4   struct point {
S-5     int x;
S-6     int y;
S-7   };
S-8
S-9   void minproc ( struct point *out, struct point *in )
S-10  {
S-11    if ( in->x < out->x ) out->x = in->x;
S-12    if ( in->y < out->y ) out->y = in->y;
S-13  }
S-14
S-15  void maxproc ( struct point *out, struct point *in )
S-16  {
S-17    if ( in->x > out->x ) out->x = in->x;
S-18    if ( in->y > out->y ) out->y = in->y;
S-19  }
S-20
S-21  #pragma omp declare reduction(min : struct point : \
S-22          minproc(&omp_out, &omp_in)) \
S-23          initializer( omp_priv = { INT_MAX, INT_MAX } )
S-24
S-25  #pragma omp declare reduction(max : struct point : \
S-26          maxproc(&omp_out, &omp_in)) \
S-27          initializer( omp_priv = { 0, 0 } )
S-28
S-29  void find_enclosing_rectangle ( int n, struct point points[] )
S-30  {
S-31    struct point minp = { INT_MAX, INT_MAX }, maxp = {0,0};
S-32    int i;
S-33
S-34  #pragma omp parallel for reduction(min:minp) reduction(max:maxp)
S-35    for ( i = 0; i < n; i++ ) {
S-36       minproc(&minp, &points[i]);
S-37       maxproc(&maxp, &points[i]);
S-38    }
S-39    printf("min = (%d, %d)\n", minp.x, minp.y);
S-40    printf("max = (%d, %d)\n", maxp.x, maxp.y);
S-41  }
```

C / C++

The following example shows the corresponding code in Fortran. The **declare reduction** directives are specified as part of the declaration in subroutine *find_enclosing_rectangle* and the procedures that perform the min and max operations are specified as subprograms.

Fortran

Example udr.1.f90 (**omp_4.0**)

```
module data_type

  type :: point
    integer :: x
    integer :: y
  end type

end module data_type

subroutine find_enclosing_rectangle ( n, points )
  use data_type
  implicit none
  integer :: n
  type(point) :: points(*)

  !$omp declare reduction(min : point : minproc(omp_out, omp_in)) &
  !$omp&  initializer( omp_priv = point( HUGE(0), HUGE(0) ) )

  !$omp declare reduction(max : point : maxproc(omp_out, omp_in)) &
  !$omp&  initializer( omp_priv = point( 0, 0 ) )

  type(point) :: minp = point( HUGE(0), HUGE(0) ), maxp = point( 0, 0 )
  integer :: i

  !$omp parallel do reduction(min:minp) reduction(max:maxp)
  do i = 1, n
     call minproc(minp, points(i))
     call maxproc(maxp, points(i))
  end do
  print *, "min = (", minp%x, minp%y, ")"
  print *, "max = (", maxp%x, maxp%y, ")"

 contains
  subroutine minproc ( out, in )
    implicit none
    type(point), intent(inout) :: out
    type(point), intent(in) :: in

    out%x = min( out%x, in%x )
    out%y = min( out%y, in%y )
```

```
S-41      end subroutine minproc
S-42
S-43      subroutine maxproc ( out, in )
S-44        implicit none
S-45        type(point), intent(inout) :: out
S-46        type(point), intent(in) :: in
S-47
S-48        out%x = max( out%x, in%x )
S-49        out%y = max( out%y, in%y )
S-50      end subroutine maxproc
S-51
S-52    end subroutine
```

Fortran

The following example shows the same computation as *udr.1* but it illustrates that you can craft complex expressions in the user-defined reduction declaration. In this case, instead of calling the *minproc* and *maxproc* functions we inline the code in a single expression.

C / C++

Example udr.2.c (`omp_4.0`)

```
S-1     #include <stdio.h>
S-2     #include <limits.h>
S-3
S-4     struct point {
S-5       int x;
S-6       int y;
S-7     };
S-8
S-9     #pragma omp declare reduction(min : struct point : \
S-10            omp_out.x = omp_in.x > omp_out.x  ? omp_out.x : omp_in.x, \
S-11            omp_out.y = omp_in.y > omp_out.y  ? omp_out.y : omp_in.y ) \
S-12            initializer( omp_priv = { INT_MAX, INT_MAX } )
S-13
S-14    #pragma omp declare reduction(max : struct point : \
S-15            omp_out.x = omp_in.x < omp_out.x  ? omp_out.x : omp_in.x,  \
S-16            omp_out.y = omp_in.y < omp_out.y  ? omp_out.y : omp_in.y ) \
S-17            initializer( omp_priv = { 0, 0 } )
S-18
S-19    void find_enclosing_rectangle ( int n, struct point points[] )
S-20    {
S-21      struct point minp = { INT_MAX, INT_MAX }, maxp = {0,0};
S-22      int i;
S-23
S-24    #pragma omp parallel for reduction(min:minp) reduction(max:maxp)
S-25      for ( i = 0; i < n; i++ ) {
S-26        if ( points[i].x < minp.x ) minp.x = points[i].x;
```

```
      if ( points[i].y < minp.y ) minp.y = points[i].y;
      if ( points[i].x > maxp.x ) maxp.x = points[i].x;
      if ( points[i].y > maxp.y ) maxp.y = points[i].y;
    }
    printf("min = (%d, %d)\n", minp.x, minp.y);
    printf("max = (%d, %d)\n", maxp.x, maxp.y);
}
```

C / C++

The corresponding code of the same example in Fortran is very similar except that the assignment expression in the **declare reduction** directive can only be used for a single variable, in this case through a type structure constructor *point(...)*.

Fortran

Example udr.2.f90 (**omp_4.0**)

```
module data_type

  type :: point
    integer :: x
    integer :: y
  end type

end module data_type

subroutine find_enclosing_rectangle ( n, points )
  use data_type
  implicit none
  integer :: n
  type(point) :: points(*)

  !$omp declare reduction( min : point :  &
  !$omp&   omp_out = point(min( omp_out%x, omp_in%x ), &
  !$omp&                   min( omp_out%y, omp_in%y )) ) &
  !$omp&   initializer( omp_priv = point( HUGE(0), HUGE(0) ) )

  !$omp declare reduction( max : point :  &
  !$omp&   omp_out = point(max( omp_out%x, omp_in%x ), &
  !$omp&                   max( omp_out%y, omp_in%y )) ) &
  !$omp&   initializer( omp_priv = point( 0, 0 ) )

  type(point) :: minp = point( HUGE(0), HUGE(0) ), maxp = point( 0, 0 )
  integer :: i

  !$omp parallel do reduction(min: minp) reduction(max: maxp)
  do i = 1, n
     minp%x = min(minp%x, points(i)%x)
```

```
S-32        minp%y = min(minp%y, points(i)%y)
S-33        maxp%x = max(maxp%x, points(i)%x)
S-34        maxp%y = max(maxp%y, points(i)%y)
S-35      end do
S-36      print *, "min = (", minp%x, minp%y, ")"
S-37      print *, "max = (", maxp%x, maxp%y, ")"
S-38
S-39    end subroutine
```

Fortran

The following example shows the use of special variables in arguments for combiner (**omp_in** and **omp_out**) and initializer (**omp_priv** and **omp_orig**) routines. This example returns the maximum value of an array and the corresponding index value. The **declare reduction** directive specifies a user-defined reduction operation *maxloc* for data type *struct mx_s*. The function *mx_combine* is the combiner and the function *mx_init* is the initializer.

C / C++

Example udr.3.c (**omp_4.0**)

```
S-1
S-2     #include <stdio.h>
S-3     #define N 100
S-4
S-5     struct mx_s {
S-6        float value;
S-7        int index;
S-8     };
S-9
S-10    /* prototype functions for combiner and initializer in
S-11       the declare reduction */
S-12    void mx_combine(struct mx_s *out, struct mx_s *in);
S-13    void mx_init(struct mx_s *priv, struct mx_s *orig);
S-14
S-15    #pragma omp declare reduction(maxloc: struct mx_s: \
S-16            mx_combine(&omp_out, &omp_in)) \
S-17            initializer(mx_init(&omp_priv, &omp_orig))
S-18
S-19    void mx_combine(struct mx_s *out, struct mx_s *in)
S-20    {
S-21       if ( out->value < in->value ) {
S-22          out->value = in->value;
S-23          out->index = in->index;
S-24       }
S-25    }
S-26
S-27    void mx_init(struct mx_s *priv, struct mx_s *orig)
S-28    {
```

```
S-29        priv->value = orig->value;
S-30        priv->index = orig->index;
S-31     }
S-32
S-33     int main(void)
S-34     {
S-35        struct mx_s mx;
S-36        float val[N], d;
S-37        int i, count = N;
S-38
S-39        for (i = 0; i < count; i++) {
S-40           d = (N*0.8f - i);
S-41           val[i] = N * N - d * d;
S-42        }
S-43
S-44        mx.value = val[0];
S-45        mx.index = 0;
S-46        #pragma omp parallel for reduction(maxloc: mx)
S-47        for (i = 1; i < count; i++) {
S-48           if (mx.value < val[i])
S-49           {
S-50              mx.value = val[i];
S-51              mx.index = i;
S-52           }
S-53        }
S-54
S-55        printf("max value = %g, index = %d\n", mx.value, mx.index);
S-56        /* prints 10000, 80 */
S-57
S-58        return 0;
S-59     }
```

C / C++

Below is the corresponding Fortran version of the above example. The **declare reduction** directive specifies the user-defined operation *maxloc* for user-derived type *mx_s*. The combiner *mx_combine* and the initializer *mx_init* are specified as subprograms.

Fortran

Example udr.3.f90 (**omp_4.0**)

```
S-1      program max_loc
S-2         implicit none
S-3
S-4         type :: mx_s
S-5            real value
S-6            integer index
S-7         end type
```

```
S-8
S-9       !$omp declare reduction(maxloc: mx_s: &
S-10      !$omp&         mx_combine(omp_out, omp_in)) &
S-11      !$omp&         initializer(mx_init(omp_priv, omp_orig))
S-12
S-13      integer, parameter :: N = 100
S-14      type(mx_s) :: mx
S-15      real :: val(N), d
S-16      integer :: i, count
S-17
S-18      count = N
S-19      do i = 1, count
S-20         d = N*0.8 - i + 1
S-21         val(i) = N * N - d * d
S-22      enddo
S-23
S-24      mx%value = val(1)
S-25      mx%index = 1
S-26      !$omp parallel do reduction(maxloc: mx)
S-27      do i = 2, count
S-28         if (mx%value < val(i)) then
S-29            mx%value = val(i)
S-30            mx%index = i
S-31         endif
S-32      enddo
S-33
S-34      print *, 'max value = ', mx%value, ' index = ', mx%index
S-35      ! prints 10000, 81
S-36
S-37    contains
S-38
S-39    subroutine mx_combine(out, in)
S-40      implicit none
S-41      type(mx_s), intent(inout) :: out
S-42      type(mx_s), intent(in) :: in
S-43
S-44      if ( out%value < in%value ) then
S-45         out%value = in%value
S-46         out%index = in%index
S-47      endif
S-48    end subroutine mx_combine
S-49
S-50    subroutine mx_init(priv, orig)
S-51      implicit none
S-52      type(mx_s), intent(out) :: priv
S-53      type(mx_s), intent(in) :: orig
S-54
```

```
S-55        priv%value = orig%value
S-56        priv%index = orig%index
S-57      end subroutine mx_init
S-58
S-59    end program
```

Fortran

The following example explains a few details of the user-defined reduction in Fortran through modules. The **declare reduction** directive is declared in a module (*data_red*). The reduction-identifier *.add.* is a user-defined operator that is to allow accessibility in the scope that performs the reduction operation. The user-defined operator *.add.* and the subroutine *dt_init* specified in the **initializer** clause are defined in the same subprogram.

The reduction operation (that is, the **reduction** clause) is in the main program. The reduction identifier *.add.* is accessible by use association. Since *.add.* is a user-defined operator, the explicit interface should also be accessible by use association in the current program unit. Since the **declare reduction** associated to this **reduction** clause has the **initializer** clause, the subroutine specified on the clause must be accessible in the current scoping unit. In this case, the subroutine *dt_init* is accessible by use association.

Fortran

Example udr.4.f90 (**omp_4.0**)

```
S-1     module data_red
S-2     ! Declare data type.
S-3       type dt
S-4         real :: r1
S-5         real :: r2
S-6       end type
S-7
S-8     ! Declare the user-defined operator .add.
S-9       interface operator(.add.)
S-10        module procedure addc
S-11      end interface
S-12
S-13    ! Declare the user-defined reduction operator .add.
S-14    !$omp declare reduction(.add.:dt:omp_out=omp_out.add.omp_in) &
S-15    !$omp& initializer(dt_init(omp_priv))
S-16
S-17     contains
S-18    ! Declare the initialization routine.
S-19      subroutine dt_init(u)
S-20        type(dt) :: u
S-21        u%r1 = 0.0
S-22        u%r2 = 0.0
S-23      end subroutine
S-24
```

```
S-25    ! Declare the specific procedure for the .add. operator.
S-26      function addc(x1, x2) result(xresult)
S-27        type(dt), intent(in) :: x1, x2
S-28        type(dt) :: xresult
S-29        xresult%r1 = x1%r1 + x2%r2
S-30        xresult%r2 = x1%r2 + x2%r1
S-31      end function
S-32
S-33    end module data_red
S-34
S-35    program main
S-36      use data_red, only : dt, dt_init, operator(.add.)
S-37
S-38      type(dt) :: xdt1, xdt2
S-39      integer :: i
S-40
S-41      xdt1 = dt(1.0,2.0)
S-42      xdt2 = dt(2.0,3.0)
S-43
S-44    ! The reduction operation
S-45    !$omp parallel do reduction(.add.: xdt1)
S-46      do i = 1, 10
S-47        xdt1 = xdt1 .add. xdt2
S-48      end do
S-49    !$omp end parallel do
S-50
S-51      print *, xdt1
S-52
S-53    end program
```

Fortran

The following example uses user-defined reductions to declare a plus (+) reduction for a C++ class. As the **declare reduction** directive is inside the context of the *V* class the expressions in the **declare reduction** directive are resolved in the context of the class. Also, note that the **initializer** clause uses a copy constructor to initialize the private variables of the reduction and it uses as parameter to its original variable by using the special variable **omp_orig**.

C++

Example udr.5.cpp (**omp_4.0**)

```
S-1   class V {
S-2      float *p;
S-3      int n;
S-4
S-5   public:
S-6      V( int _n )      : n(_n)  { p = new float[n]; }
S-7      V( const V& m ) : n(m.n) { p = new float[n]; }
S-8      ~V() { delete[] p; }
S-9
S-10     V& operator+= ( const V& );
S-11
S-12     #pragma omp declare reduction( + : V : omp_out += omp_in ) \
S-13             initializer(omp_priv(omp_orig))
S-14  };
```

C++

The following examples shows how user-defined reductions can be defined for some STL containers. The first **declare reduction** defines the plus (+) operation for *std::vector<int>* by making use of the *std::transform* algorithm. The second and third define the merge (or concatenation) operation for *std::vector<int>* and *std::list<int>*. It shows how the user-defined reduction operation can be applied to specific data types of an STL.

C++

Example udr.6.cpp (**omp_4.0**)

```
S-1   #include <algorithm>
S-2   #include <list>
S-3   #include <vector>
S-4
S-5   #pragma omp declare reduction( + : std::vector<int> : \
S-6        std::transform (omp_out.begin(), omp_out.end(),  \
S-7                        omp_in.begin(), omp_in.end(),std::plus<int>()))
S-8
S-9   #pragma omp declare reduction( merge : std::vector<int> : \
S-10       omp_out.insert(omp_out.end(), omp_in.begin(), omp_in.end()))
S-11
S-12  #pragma omp declare reduction( merge : std::list<int> : \
S-13       omp_out.merge(omp_in))
```

C++

7.10 The `scan` Directive

The following examples illustrate how to parallelize a loop that saves the *prefix sum* of a reduction. This is accomplished by using the **`inscan`** modifier in the **`reduction`** clause for the input variable of the scan, and specifying with a **`scan`** directive whether the storage statement includes or excludes the scan input of the present iteration (`k`).

Basically, the **`inscan`** modifier connects a loop and/or SIMD reduction to the scan operation, and a **`scan`** construct with an **`inclusive`** or **`exclusive`** clause specifies whether the "scan phase" (lexical block before and after the directive, respectively) is to use an *inclusive* or *exclusive* scan value for the list item (`x`).

The first example uses the *inclusive* scan operation on a composite loop-SIMD construct. The **`scan`** directive separates the reduction statement on variable `x` from the use of `x` (saving to array `b`). The order of the statements in this example indicates that value `a[k]` (`a(k)` in Fortran) is included in the computation of the prefix sum `b[k]` (`b(k)` in Fortran) for iteration `k`.

C / C++

Example scan.1.c (**`omp_5.0`**)

```
S-1   #include <stdio.h>
S-2   #define N 100
S-3
S-4   int main(void)
S-5   {
S-6      int a[N], b[N];
S-7      int x = 0;
S-8
S-9      // initialization
S-10     for (int k = 0; k < N; k++)
S-11        a[k] = k + 1;
S-12
S-13     // a[k] is included in the computation of producing results in b[k]
S-14     #pragma omp parallel for simd reduction(inscan,+: x)
S-15     for (int k = 0; k < N; k++) {
S-16        x += a[k];
S-17        #pragma omp scan inclusive(x)
S-18        b[k] = x;
S-19     }
S-20
S-21     printf("x = %d, b[0:3] = %d %d %d\n", x, b[0], b[1], b[2]);
S-22     //           5050,        1  3  6
S-23
S-24     return 0;
S-25  }
```

C / C++

Fortran

Example scan.1.f90 (**omp_5.0**)

```
S-1   program inclusive_scan
S-2      implicit none
S-3      integer, parameter :: n = 100
S-4      integer a(n), b(n)
S-5      integer x, k
S-6
S-7      ! initialization
S-8      x = 0
S-9      do k = 1, n
S-10        a(k) = k
S-11     end do
S-12
S-13     ! a(k) is included in the computation of producing results in b(k)
S-14     !$omp parallel do simd reduction(inscan,+: x)
S-15     do k = 1, n
S-16        x = x + a(k)
S-17        !$omp scan inclusive(x)
S-18        b(k) = x
S-19     end do
S-20
S-21     print *,'x =', x, ', b(1:3) =', b(1:3)
S-22     !           5050,            1  3  6
S-23
S-24  end program
```

Fortran

The second example uses the *exclusive* scan operation on a composite loop-SIMD construct. The **scan** directive separates the use of `x` (saving to array `b`) from the reduction statement on variable `x`. The order of the statements in this example indicates that value `a[k]` (`a(k)` in Fortran) is excluded from the computation of the prefix sum `b[k]` (`b(k)` in Fortran) for iteration `k`.

C / C++

Example scan.2.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   #define N 100
S-3
S-4   int main(void)
S-5   {
S-6      int a[N], b[N];
S-7      int x = 0;
S-8
S-9      // initialization
```

```
S-10        for (int k = 0; k < N; k++)
S-11           a[k] = k + 1;
S-12
S-13        // a[k] is not included in the computation of producing results in b[k]
S-14        #pragma omp parallel for simd reduction(inscan,+: x)
S-15        for (int k = 0; k < N; k++) {
S-16           b[k] = x;
S-17           #pragma omp scan exclusive(x)
S-18           x += a[k];
S-19        }
S-20
S-21        printf("x = %d, b[0:3] = %d %d %d\n", x, b[0], b[1], b[2]);
S-22        //           5050,        0  1  3
S-23
S-24        return 0;
S-25     }
```

C / C++

Fortran

1 *Example scan.2.f90* (`omp_5.0`)

```
S-1      program exclusive_scan
S-2         implicit none
S-3         integer, parameter :: n = 100
S-4         integer a(n), b(n)
S-5         integer x, k
S-6
S-7         ! initialization
S-8         x = 0
S-9         do k = 1, n
S-10           a(k) = k
S-11        end do
S-12
S-13        ! a(k) is not included in the computation of producing results in b(k)
S-14        !$omp parallel do simd reduction(inscan,+: x)
S-15        do k = 1, n
S-16           b(k) = x
S-17           !$omp scan exclusive(x)
S-18           x = x + a(k)
S-19        end do
S-20
S-21        print *,'x =', x, ', b(1:3) =', b(1:3)
S-22        !           5050,            0  1  3
S-23
S-24     end program
```

Fortran

7.11 The `copyin` Clause

The **copyin** clause is used to initialize threadprivate data upon entry to a **parallel** region. The value of the threadprivate variable in the master thread is copied to the threadprivate variable of each other team member.

C / C++

Example copyin.1.c

```
S-1   #include <stdlib.h>
S-2
S-3   float* work;
S-4   int size;
S-5   float tol;
S-6
S-7   #pragma omp threadprivate(work,size,tol)
S-8
S-9   void build()
S-10  {
S-11    int i;
S-12    work = (float*)malloc( sizeof(float)*size );
S-13    for( i = 0; i < size; ++i ) work[i] = tol;
S-14  }
S-15
S-16  void copyin_example( float t, int n )
S-17  {
S-18    tol = t;
S-19    size = n;
S-20    #pragma omp parallel copyin(tol,size)
S-21    {
S-22      build();
S-23    }
S-24  }
```

C / C++

Fortran

1 *Example copyin.1.f*

```
S-1         MODULE M
S-2           REAL, POINTER, SAVE :: WORK(:)
S-3           INTEGER :: SIZE
S-4           REAL :: TOL
S-5   !$OMP   THREADPRIVATE(WORK,SIZE,TOL)
S-6         END MODULE M
S-7
S-8         SUBROUTINE COPYIN_EXAMPLE( T, N )
S-9           USE M
S-10          REAL :: T
S-11          INTEGER :: N
S-12          TOL = T
S-13          SIZE = N
S-14  !$OMP   PARALLEL COPYIN(TOL,SIZE)
S-15          CALL BUILD
S-16  !$OMP   END PARALLEL
S-17        END SUBROUTINE COPYIN_EXAMPLE
S-18
S-19        SUBROUTINE BUILD
S-20          USE M
S-21          ALLOCATE(WORK(SIZE))
S-22          WORK = TOL
S-23        END SUBROUTINE BUILD
```

Fortran

7.12 The `copyprivate` Clause

The **copyprivate** clause can be used to broadcast values acquired by a single thread directly to all instances of the private variables in the other threads. In this example, if the routine is called from the sequential part, its behavior is not affected by the presence of the directives. If it is called from a **parallel** region, then the actual arguments with which **a** and **b** are associated must be private.

The thread that executes the structured block associated with the **single** construct broadcasts the values of the private variables **a**, **b**, **x**, and **y** from its implicit task's data environment to the data environments of the other implicit tasks in the thread team. The broadcast completes before any of the threads have left the barrier at the end of the construct.

C / C++

Example copyprivate.1.c

```
S-1   #include <stdio.h>
S-2   float x, y;
S-3   #pragma omp threadprivate(x, y)
S-4
S-5   void init(float a, float b ) {
S-6       #pragma omp single copyprivate(a,b,x,y)
S-7       {
S-8           scanf("%f %f %f %f", &a, &b, &x, &y);
S-9       }
S-10  }
```

C / C++

Fortran

Example copyprivate.1.f

```
S-1         SUBROUTINE INIT(A,B)
S-2         REAL A, B
S-3           COMMON /XY/ X,Y
S-4   !$OMP   THREADPRIVATE (/XY/)
S-5
S-6   !$OMP   SINGLE
S-7             READ (11) A,B,X,Y
S-8   !$OMP   END SINGLE COPYPRIVATE (A,B,/XY/)
S-9
S-10        END SUBROUTINE INIT
```

Fortran

In this example, assume that the input must be performed by the master thread. Since the **master** construct does not support the **copyprivate** clause, it cannot broadcast the input value that is read. However, **copyprivate** is used to broadcast an address where the input value is stored.

C / C++

1 *Example copyprivate.2.c*

```
S-1   #include <stdio.h>
S-2   #include <stdlib.h>
S-3
S-4   float read_next( ) {
S-5     float * tmp;
S-6     float return_val;
S-7
S-8     #pragma omp single copyprivate(tmp)
S-9     {
S-10      tmp = (float *) malloc(sizeof(float));
S-11    }  /* copies the pointer only */
S-12
S-13
S-14    #pragma omp master
S-15    {
S-16      scanf("%f", tmp);
S-17    }
S-18
S-19    #pragma omp barrier
S-20    return_val = *tmp;
S-21    #pragma omp barrier
S-22
S-23    #pragma omp single nowait
S-24    {
S-25      free(tmp);
S-26    }
S-27
S-28    return return_val;
S-29  }
```

C / C++

Fortran

2 *Example copyprivate.2.f*

```
S-1           REAL FUNCTION READ_NEXT()
S-2           REAL, POINTER :: TMP
S-3
S-4   !$OMP   SINGLE
S-5             ALLOCATE (TMP)
S-6   !$OMP   END SINGLE COPYPRIVATE (TMP)  ! copies the pointer only
S-7
S-8   !$OMP   MASTER
S-9             READ (11) TMP
S-10  !$OMP   END MASTER
```

```
S-11
S-12    !$OMP   BARRIER
S-13              READ_NEXT = TMP
S-14    !$OMP   BARRIER
S-15
S-16    !$OMP   SINGLE
S-17              DEALLOCATE (TMP)
S-18    !$OMP   END SINGLE NOWAIT
S-19            END FUNCTION READ_NEXT
```

Fortran

Suppose that the number of lock variables required within a **parallel** region cannot easily be determined prior to entering it. The **copyprivate** clause can be used to provide access to shared lock variables that are allocated within that **parallel** region.

C / C++

Example copyprivate.3.c

```
S-1     #include <stdio.h>
S-2     #include <stdlib.h>
S-3     #include <omp.h>
S-4
S-5     omp_lock_t *new_lock()
S-6     {
S-7       omp_lock_t *lock_ptr;
S-8
S-9       #pragma omp single copyprivate(lock_ptr)
S-10      {
S-11        lock_ptr = (omp_lock_t *) malloc(sizeof(omp_lock_t));
S-12        omp_init_lock( lock_ptr );
S-13      }
S-14
S-15      return lock_ptr;
S-16    }
```

C / C++

Fortran

Example copyprivate.3.f

```
      FUNCTION NEW_LOCK()
      USE OMP_LIB       ! or INCLUDE "omp_lib.h"
        INTEGER(OMP_LOCK_KIND), POINTER :: NEW_LOCK

!$OMP   SINGLE
          ALLOCATE(NEW_LOCK)
          CALL OMP_INIT_LOCK(NEW_LOCK)
!$OMP   END SINGLE COPYPRIVATE(NEW_LOCK)
      END FUNCTION NEW_LOCK
```

Note that the effect of the **copyprivate** clause on a variable with the **allocatable** attribute is different than on a variable with the **pointer** attribute. The value of **A** is copied (as if by intrinsic assignment) and the pointer **B** is copied (as if by pointer assignment) to the corresponding list items in the other implicit tasks belonging to the **parallel** region.

Example copyprivate.4.f

```
      SUBROUTINE S(N)
      INTEGER N

        REAL, DIMENSION(:), ALLOCATABLE :: A
        REAL, DIMENSION(:), POINTER :: B

        ALLOCATE (A(N))
!$OMP   SINGLE
          ALLOCATE (B(N))
          READ (11) A,B
!$OMP   END SINGLE COPYPRIVATE(A,B)
        ! Variable A is private and is
        ! assigned the same value in each thread
        ! Variable B is shared

!$OMP   BARRIER
!$OMP   SINGLE
          DEALLOCATE (B)
!$OMP   END SINGLE NOWAIT
      END SUBROUTINE S
```

Fortran

7.13 C++ Reference in Data-Sharing Clauses

C++

C++ reference types are allowed in data-sharing attribute clauses as of OpenMP 4.5, except for the **threadprivate**, **copyin** and **copyprivate** clauses. (See the Data-Sharing Attribute Clauses Section of the 4.5 OpenMP specification.) When a variable with C++ reference type is privatized, the object the reference refers to is privatized in addition to the reference itself. The following example shows the use of reference types in data-sharing clauses in the usual way. Additionally it shows how the data-sharing of formal arguments with a C++ reference type on an orphaned task generating construct is determined implicitly. (See the Data-sharing Attribute Rules for Variables Referenced in a Construct Section of the 4.5 OpenMP specification.)

Example cpp_reference.1.cpp (**omp_4.5**)

```
S-1
S-2   void task_body (int &);
S-3   void gen_task (int &x) { // on orphaned task construct reference argument
S-4     #pragma omp task // x is implicitly determined firstprivate(x)
S-5     task_body (x);
S-6   }
S-7   void test (int &y, int &z) {
S-8     #pragma omp parallel private(y)
S-9     {
S-10      y = z + 2;
S-11      gen_task (y); // no matter if the argument is determined private
S-12      gen_task (z); // or shared in the enclosing context.
S-13
S-14      y++;          // each thread has its own int object y refers to
S-15      gen_task (y);
S-16    }
S-17  }
S-18
```

C++

7.14 Fortran ASSOCIATE Construct

Fortran

The following is an invalid example of specifying an associate name on a data-sharing attribute clause. The constraint in the Data Sharing Attribute Rules section in the OpenMP 4.0 API Specifications states that an associate name preserves the association with the selector established at the **ASSOCIATE** statement. The associate name *b* is associated with the shared variable *a*. With the predetermined data-sharing attribute rule, the associate name *b* is not allowed to be specified on the **private** clause.

Example associate.1.f (**omp_4.0**)

```
S-1         program example_broken
S-2         real :: a, c
S-3         associate (b => a)
S-4   !$omp parallel private(b, c)        ! invalid to privatize b
S-5         c = 2.0*b
S-6   !$omp end parallel
S-7         end associate
S-8         end program
```

In next example, within the **parallel** construct, the association name *thread_id* is associated with the private copy of *i*. The print statement should output the unique thread number.

Example associate.2.f (**omp_4.0**)

```
S-1         program example
S-2         use omp_lib
S-3         integer  i
S-4   !$omp parallel private(i)
S-5         i = omp_get_thread_num()
S-6         associate(thread_id => i)
S-7           print *, thread_id       ! print private i value
S-8         end associate
S-9   !$omp end parallel
S-10        end program
```

The following example illustrates the effect of specifying a selector name on a data-sharing attribute clause. The associate name *u* is associated with *v* and the variable *v* is specified on the **private** clause of the **parallel** construct. The construct association is established prior to the **parallel** region. The association between *u* and the original *v* is retained (see the Data Sharing Attribute Rules section in the OpenMP 4.0 API Specifications). Inside the **parallel** region, *v* has the value of -1 and *u* has the value of the original *v*.

Example associate.3.f90 (**omp_4.0**)

```
S-1    program example
S-2      integer :: v
S-3      v = 15
S-4    associate(u => v)
S-5    !$omp parallel private(v)
S-6      v = -1
S-7      print *, v                ! private v=-1
S-8      print *, u                ! original v=15
S-9    !$omp end parallel
S-10   end associate
S-11   end program
```

Fortran

CHAPTER 8

Memory Model

OpenMP provides a shared-memory model that allows all threads on a given device shared access to *memory*. For a given OpenMP region that may be executed by more than one thread or SIMD lane, variables in memory may be *shared* or *private* with respect to those threads or SIMD lanes. A variable's data-sharing attribute indicates whether it is shared (the *shared* attribute) or private (the *private*, *firstprivate*, *lastprivate*, *linear*, and *reduction* attributes) in the data environment of an OpenMP region. While private variables in an OpenMP region are new copies of the original variable (with same name) that may then be concurrently accessed or modified by their respective threads or SIMD lanes, a shared variable in an OpenMP region is the same as the variable of the same name in the enclosing region. Concurrent accesses or modifications to a shared variable may therefore require synchronization to avoid data races.

OpenMP's memory model also includes a *temporary view* of memory that is associated with each thread. Two different threads may see different values for a given variable in their respective temporary views. Threads may employ flush operations for the purposes of making their temporary view of a variable consistent with the value of the variable in memory. The effect of a given flush operation is characterized by its flush properties – some combination of *strong*, *release*, and *acquire* – and, for *strong* flushes, a *flush-set*.

A *strong* flush will force consistency between the temporary view and the memory for all variables in its *flush-set*. Furthermore all strong flushes in a program that have intersecting flush-sets will execute in some total order, and within a thread strong flushes may not be reordered with respect to other memory operations on variables in its flush-set. *Release* and *acquire* flushes operate in pairs. A release flush may "synchronize" with an acquire flush, and when it does so the local memory operations that precede the release flush will appear to have been completed before the local memory operations on the same variables that follow the acquire flush.

Flush operations arise from explicit **`flush`** directives, implicit **`flush`** directives, and also from the execution of **`atomic`** constructs. The **`flush`** directive forces a consistent view of local variables of the thread executing the **`flush`**. When a list is supplied on the directive, only the items (variables) in the list are guaranteed to be flushed. Implied flushes exist at prescribed locations of

certain constructs. For the complete list of these locations and associated constructs, please refer to the *flush Construct* section of the OpenMP Specifications document.

In this chapter, examples illustrate how race conditions may arise for accesses to variables with a *shared* data-sharing attribute when flush operations are not properly employed. A race condition can exist when two or more threads are involved in accessing a variable and at least one of the accesses modifies the variable. In particular, a data race will arise when conflicting accesses do not have a well-defined *completion order*. The existence of data races in OpenMP programs result in undefined behavior, and so they should generally be avoided for programs to be correct. The completion order of accesses to a shared variable is guaranteed in OpenMP through a set of memory consistency rules that are described in the *OpenMP Memory Consitency* section of the OpenMP Specifications document.

8.1 The OpenMP Memory Model

The following examples illustrate two major concerns for concurrent thread execution: ordering of thread execution and memory accesses that may or may not lead to race conditions.

In the following example, at Print 1, the value of **xval** could be either 2 or 5, depending on the timing of the threads. The **atomic** directives are necessary for the accesses to **x** by threads 1 and 2 to avoid a data race. If the atomic write completes before the atomic read, thread 1 is guaranteed to see 5 in **xval**. Otherwise, thread 1 is guaranteed to see 2 in **xval**.

The barrier after Print 1 contains implicit flushes on all threads, as well as a thread synchronization, so the programmer is guaranteed that the value 5 will be printed by both Print 2 and Print 3. Since neither Print 2 or Print 3 are modifying **x**, they may concurrently access **x** without requiring **atomic** directives to avoid a data race.

C / C++

Example mem_model.1.c (**omp_3.1**)

```
S-1   #include <stdio.h>
S-2   #include <omp.h>
S-3
S-4   int main(){
S-5     int x;
S-6
S-7     x = 2;
S-8     #pragma omp parallel num_threads(2) shared(x)
S-9     {
S-10
S-11      if (omp_get_thread_num() == 0) {
S-12         #pragma omp atomic write
S-13         x = 5;
S-14      } else {
S-15        int xval;
S-16        #pragma omp atomic read
S-17        xval = x;
S-18      /* Print 1: xval can be 2 or 5 */
S-19        printf("1: Thread# %d: x = %d\n", omp_get_thread_num(), xval);
S-20      }
S-21
S-22      #pragma omp barrier
S-23
S-24      if (omp_get_thread_num() == 0) {
S-25      /* Print 2 */
S-26        printf("2: Thread# %d: x = %d\n", omp_get_thread_num(), x);
S-27      } else {
S-28      /* Print 3 */
S-29        printf("3: Thread# %d: x = %d\n", omp_get_thread_num(), x);
S-30      }
```

```
S-31      }
S-32      return 0;
S-33    }
```

C / C++

Fortran

Example mem_model.1.f90 (`omp_3.1`)

```
 S-1    PROGRAM MEMMODEL
 S-2      INCLUDE "omp_lib.h"      ! or USE OMP_LIB
 S-3      INTEGER X, XVAL
 S-4
 S-5      X = 2
 S-6    !$OMP PARALLEL NUM_THREADS(2) SHARED(X)
 S-7
 S-8        IF (OMP_GET_THREAD_NUM() .EQ. 0) THEN
 S-9        !$OMP ATOMIC WRITE
S-10           X = 5
S-11        ELSE
S-12        !$OMP ATOMIC READ
S-13          XVAL = X
S-14        ! PRINT 1: XVAL can be 2 or 5
S-15          PRINT *,"1: THREAD# ", OMP_GET_THREAD_NUM(), "X = ", XVAL
S-16        ENDIF
S-17
S-18     !$OMP BARRIER
S-19
S-20        IF (OMP_GET_THREAD_NUM() .EQ. 0) THEN
S-21        ! PRINT 2
S-22          PRINT *,"2: THREAD# ", OMP_GET_THREAD_NUM(), "X = ", X
S-23        ELSE
S-24        ! PRINT 3
S-25          PRINT *,"3: THREAD# ", OMP_GET_THREAD_NUM(), "X = ", X
S-26        ENDIF
S-27
S-28    !$OMP END PARALLEL
S-29
S-30    END PROGRAM MEMMODEL
```

Fortran

The following example demonstrates why synchronization is difficult to perform correctly through variables. The write to **flag** on thread 0 and the read from **flag** in the loop on thread 1 must be atomic to avoid a data race. When thread 1 breaks out of the loop, **flag** will have the value of 1. However, **data** will still be undefined at the first print statement. Only after the flush of both **flag** and **data** after the first print statement will **data** have the well-defined value of 42.

C / C++

Example mem_model.2.c (**omp_3.1**)

```
#include <omp.h>
#include <stdio.h>
int main()
{
    int data;
    int flag=0;
    #pragma omp parallel num_threads(2)
    {
      if (omp_get_thread_num()==0)
      {
      /* Write to the data buffer that will be
       * read by thread */
          data = 42;
      /* Flush data to thread 1 and strictly order
       * the write to data relative to the write to the flag */
          #pragma omp flush(flag, data)
      /* Set flag to release thread 1 */
          #pragma omp atomic write
          flag = 1;
      }
      else if(omp_get_thread_num()==1)
      {
      /* Loop until we see the update to the flag */
          #pragma omp flush(flag, data)
          int flag_val = 0;
          while (flag_val < 1)
          {
             #pragma omp atomic read
             flag_val = flag;
          }
      /* Value of flag is 1; value of data is undefined */
          printf("flag=%d data=%d\n", flag, data);
          #pragma omp flush(flag, data)
      /* Value of flag is 1; value of data is 42 */
          printf("flag=%d data=%d\n", flag, data);
      }
    }
```

```
S-38        return 0;
S-39    }
```

C / C++

Fortran

Example mem_model.2.f (`omp_3.1`)

```
 S-1        PROGRAM EXAMPLE
 S-2        INCLUDE "omp_lib.h" ! or USE OMP_LIB
 S-3        INTEGER DATA
 S-4        INTEGER FLAG, FLAG_VAL
 S-5
 S-6        FLAG = 0
 S-7    !$OMP  PARALLEL NUM_THREADS(2)
 S-8          IF(OMP_GET_THREAD_NUM() .EQ. 0) THEN
 S-9          ! Write to the data buffer that will be read by thread 1
S-10             DATA = 42
S-11
S-12          ! Flush DATA to thread 1 and strictly order the write to DATA
S-13          ! relative to the write to the FLAG
S-14    !$OMP     FLUSH(FLAG, DATA)
S-15
S-16          ! Set FLAG to release thread 1
S-17    !$OMP     ATOMIC WRITE
S-18              FLAG = 1
S-19
S-20          ELSE IF(OMP_GET_THREAD_NUM() .EQ. 1) THEN
S-21          ! Loop until we see the update to the FLAG
S-22    !$OMP     FLUSH(FLAG, DATA)
S-23              FLAG_VAL = 0
S-24              DO WHILE(FLAG_VAL .LT. 1)
S-25    !$OMP        ATOMIC READ
S-26                 FLAG_VAL = FLAG
S-27              ENDDO
S-28
S-29          ! Value of FLAG is 1; value of DATA is undefined
S-30              PRINT *, 'FLAG=', FLAG, ' DATA=', DATA
S-31
S-32    !$OMP     FLUSH(FLAG, DATA)
S-33          ! Value of FLAG is 1; value of DATA is 42
S-34              PRINT *, 'FLAG=', FLAG, ' DATA=', DATA
S-35
S-36          ENDIF
S-37    !$OMP  END PARALLEL
S-38        END
```

Fortran

The next example demonstrates why synchronization is difficult to perform correctly through variables. As in the preceding example, the updates to **flag** and the reading of **flag** in the loops on threads 1 and 2 are performed atomically to avoid data races on **flag**. However, the code still contains data race due to the incorrect use of "flush with a list" after the assignment to **data1** on thread 1. By not including **flag** in the flush-set of that **flush** directive, the assignment can be reordered with respect to the subsequent atomic update to **flag**. Consequentially, **data1** is undefined at the print statement on thread 2.

C / C++

Example mem_model.3.c (**omp_3.1**)

```
S-1   #include <omp.h>
S-2   #include <stdio.h>
S-3
S-4   int data0 = 0, data1 = 0;
S-5
S-6   int main()
S-7   {
S-8       int flag=0;
S-9
S-10      #pragma omp parallel num_threads(3)
S-11      {
S-12          if(omp_get_thread_num()==0)
S-13          {
S-14              data0 = 17;
S-15              #pragma omp flush
S-16          /* Set flag to release thread 1 */
S-17              #pragma omp atomic update
S-18              flag++;
S-19          /* Flush of flag is implied by the atomic directive */
S-20          }
S-21          else if(omp_get_thread_num()==1)
S-22          {
S-23              int flag_val = 0;
S-24          /* Loop until we see that flag reaches 1*/
S-25              while(flag_val < 0)
S-26              {
S-27                  #pragma omp atomic read
S-28                  flag_val = flag;
S-29              }
S-30              #pragma omp flush(data0)
S-31          /* data0 is 17 here */
S-32              printf("Thread 1 awoken (data0 = %d)\n", data0);
S-33              data1 = 42;
S-34              #pragma omp flush(data1)
S-35          /* Set flag to release thread 2 */
S-36              #pragma omp atomic update
```

```
            flag++;
        /* Flush of flag is implied by the atomic directive */
        }
        else if(omp_get_thread_num()==2)
        {
            int flag_val = 0;
        /* Loop until we see that flag reaches 2 */
            while(flag_val < 2)
            {
                #pragma omp atomic read
                flag_val = flag;
            }
            #pragma omp flush(data0,data1)
        /* there is a data race here; data0 is 17 and data1 is undefined */
            printf("Thread 2 awoken (data0 = %d, data1 = %d)\n", data0, data1);
        }
    }
    return 0;
}
```

C / C++

Fortran

Example mem_model.3.f (`omp_3.1`)

```
        PROGRAM EXAMPLE
        INCLUDE "omp_lib.h" ! or USE OMP_LIB
        INTEGER FLAG, FLAG_VAL
        INTEGER DATA0, DATA1

        FLAG = 0
!$OMP  PARALLEL NUM_THREADS(3)
         IF(OMP_GET_THREAD_NUM() .EQ. 0) THEN
            DATA0 = 17
!$OMP       FLUSH

          ! Set flag to release thread 1
!$OMP       ATOMIC UPDATE
            FLAG = FLAG + 1
          ! Flush of FLAG is implied by the atomic directive

          ELSE IF(OMP_GET_THREAD_NUM() .EQ. 1) THEN
          ! Loop until we see that FLAG reaches 1
!$OMP       FLUSH(FLAG, DATA)
            FLAG_VAL = 0
            DO WHILE(FLAG_VAL .LT. 1)
!$OMP          ATOMIC READ
               FLAG_VAL = FLAG
```

```
S-24                ENDDO
S-25    !$OMP       FLUSH
S-26
S-27              ! DATA0 is 17 here
S-28                PRINT *, 'Thread 1 awoken. DATA0 = ', DATA0
S-29
S-30                DATA1 = 42
S-31    !$OMP       FLUSH(DATA1)
S-32
S-33              ! Set FLAG to release thread 2
S-34    !$OMP       ATOMIC UPDATE
S-35                FLAG = FLAG + 1
S-36              ! Flush of FLAG is implied by the atomic directive
S-37
S-38              ELSE IF(OMP_GET_THREAD_NUM() .EQ. 2) THEN
S-39              ! Loop until we see that FLAG reaches 2
S-40                FLAG_VAL = 0
S-41                DO WHILE(FLAG_VAL .LT. 2)
S-42    !$OMP          ATOMIC READ
S-43                   FLAG_VAL = FLAG
S-44                ENDDO
S-45    !$OMP       FLUSH(DATA0, DATA1)
S-46
S-47              ! There is a data race here; data0 is 17 and data1 is undefined
S-48                PRINT *, 'Thread 2 awoken. DATA0 = ', DATA0,
S-49         &                 ' and DATA1 = ', DATA1
S-50
S-51              ENDIF
S-52    !$OMP  END PARALLEL
S-53           END
```

Fortran

8.2 Memory Allocators

OpenMP memory allocators can be used to allocate memory with specific allocator traits. In the following example an OpenMP allocator is used to specify an alignment for arrays *x* and *y*. The general approach for attributing traits to variables allocated by OpenMP is to create or specify a pre-defined *memory space*, create an array of *traits*, and then form an *allocator* from the memory space and trait. The allocator is then specified in an OpenMP allocation (using an API *omp_alloc()* function for C/C++ code and an **allocate** directive for Fortran code in the allocators.1 example).

In the example below the *xy_memspace* variable is declared and assigned the default memory space (*omp_default_mem_space*). Next, an array for *traits* is created. Since only one trait will be used, the array size is *1*. A trait is a structure in C/C++ and a derived type in Fortran, containing 2 components: a key and a corresponding value (key-value pair). The trait key used here is *omp_atk_alignment* (an enum for C/C++ and a parameter for Fortran) and the trait value of 64 is specified in the *xy_traits* declaration. These declarations are followed by a call to the *omp_init_allocator()* function to combine the memory space (*xy_memspace*) and the traits (*xy_traits*) to form an allocator (*xy_alloc*).

In the C/C++ code the API *omp_allocate()* function is used to allocate space, similar to *malloc*, except that the allocator is specified as the second argument. In Fortran an **allocate** directive is used to specify an allocator for a following Fortran *allocate* statement. A variable list may be supplied if the allocator is to be applied to a subset of variables in the Fortran allocate statement. Specifying the complete list is optional. Here, the *xy_alloc* allocator is specified in the **allocator** clause, and the set of all variables used in the allocate statement is specified in the list.

C / C++

Example allocators.1.c (**omp_5.0**)

```
S-1   #include    <omp.h>
S-2   #include  <stdio.h>
S-3   #include <stdlib.h>
S-4   #include <stdint.h>
S-5   #define N 1000
S-6
S-7   int main()
S-8   {
S-9      float  *x, *y;
S-10     float s=2.0;
S-11
S-12  omp_memspace_handle_t  xy_memspace = omp_default_mem_space;
S-13  omp_alloctrait_t       xy_traits[1]={omp_atk_alignment, 64};
S-14  omp_allocator_handle_t xy_alloc    = omp_init_allocator(xy_memspace,1,xy_traits);
S-15
S-16
S-17     x=(float *)omp_alloc(N*sizeof(float), xy_alloc);
S-18     y=(float *)omp_alloc(N*sizeof(float), xy_alloc);
```

```
S-19
S-20        if( ((intptr_t)(y))%64 != 0 || ((intptr_t)(x))%64 != 0 )
S-21        { printf("ERROR: x|y not 64-Byte aligned\n"); exit(1); }
S-22
S-23        #pragma omp parallel
S-24        {
S-25           #pragma omp for simd simdlen(16) aligned(x,y:64)
S-26           for(int i=0; i<N; i++){ x[i]=i+1; y[i]=i+1; } // initialize
S-27
S-28           #pragma omp for simd simdlen(16) aligned(x,y:64)
S-29           for(int i=0; i<N; i++) y[i] = s*x[i] + y[i];
S-30         }
S-31
S-32
S-33        printf("y[0],y[N-1]: %5.0f %5.0f\n",y[0],y[N-1]); //output: y...   3  3000
S-34
S-35        omp_free(x, xy_alloc);
S-36        omp_free(y, xy_alloc);
S-37        omp_destroy_allocator(xy_alloc);
S-38
S-39        return 0;
S-40     }
```

C / C++

Fortran

1 *Example allocators.1.f90* (`omp_5.0`)

```
S-1
S-2     program main
S-3      use omp_lib
S-4
S-5      integer, parameter :: N=1000, align=64
S-6      real, allocatable  :: x(:),y(:)
S-7      real               :: s = 2.0e0
S-8      integer            :: i
S-9
S-10     integer(omp_memspace_handle_kind ) :: xy_memspace = omp_default_mem_space
S-11     type(   omp_alloctrait           ) :: xy_traits(1) = &
S-12                                             [omp_alloctrait(omp_atk_alignment,64)]
S-13     integer(omp_allocator_handle_kind) :: xy_alloc
S-14
S-15       xy_alloc   =    omp_init_allocator(   xy_memspace, 1, xy_traits)
S-16
S-17       !$omp allocate(x,y) allocator(xy_alloc)
S-18       allocate(x(N),y(N))
S-19                             !! loc is non-standard, but found everywhere
S-20                             !! remove these lines if not available
```

```
S-21        if(modulo(loc(x),align) /= 0 .and. modulo(loc(y),align) /=0 ) then
S-22           print*,"ERROR: x|y not 64-byte aligned"; stop
S-23        endif
S-24
S-25        !$omp parallel
S-26
S-27           !$omp do simd simdlen(16) aligned(x,y: 64) !! 64B aligned
S-28           do i=1,N  !! initialize
S-29             x(i)=i
S-30             y(i)=i
S-31           end do
S-32
S-33           !$omp do simd simdlen(16) aligned(x,y: 64) !! 64B aligned
S-34           do i = 1,N
S-35              y(i) = s*x(i) + y(i)
S-36           end do
S-37
S-38        !$omp end parallel
S-39
S-40        write(*,'("y(1),y(N):",2f6.0)') y(1),y(N) !!output: y... 3. 3000.
S-41
S-42        deallocate(x,y)
S-43        call omp_destroy_allocator(xy_alloc)
S-44
S-45     end program
S-46
```

Fortran

8.3 Race Conditions Caused by Implied Copies of Shared Variables in Fortran

Fortran

The following example contains a race condition, because the shared variable, which is an array section, is passed as an actual argument to a routine that has an assumed-size array as its dummy argument. The subroutine call passing an array section argument may cause the compiler to copy the argument into a temporary location prior to the call and copy from the temporary location into the original variable when the subroutine returns. This copying would cause races in the **parallel** region.

Example fort_race.1.f90

```
S-1     SUBROUTINE SHARED_RACE
S-2
S-3       INCLUDE "omp_lib.h"      ! or USE OMP_LIB
S-4
S-5       REAL A(20)
S-6       INTEGER MYTHREAD
S-7
S-8     !$OMP PARALLEL SHARED(A) PRIVATE(MYTHREAD)
S-9
S-10      MYTHREAD = OMP_GET_THREAD_NUM()
S-11      IF (MYTHREAD .EQ. 0) THEN
S-12         CALL SUB(A(1:10)) ! compiler may introduce writes to A(6:10)
S-13      ELSE
S-14         A(6:10) = 12
S-15      ENDIF
S-16
S-17    !$OMP END PARALLEL
S-18
S-19    END SUBROUTINE SHARED_RACE
S-20
S-21    SUBROUTINE SUB(X)
S-22      REAL X(*)
S-23      X(1:5) = 4
S-24    END SUBROUTINE SUB
```

Fortran

CHAPTER 9

Program Control

Some specific and elementary concepts of controlling program execution are illustrated in the examples of this chapter. Control can be directly managed with conditional control code (ifdef's with the **_OPENMP** macro, and the Fortran sentinel (**!$**) for conditionally compiling). The **if** clause on some constructs can direct the runtime to ignore or alter the behavior of the construct. Of course, the base-language **if** statements can be used to control the "execution" of stand-alone directives (such as **flush**, **barrier**, **taskwait**, and **taskyield**). However, the directives must appear in a block structure, and not as a substatement as shown in examples 1 and 2 of this chapter.

CANCELLATION

Cancellation (termination) of the normal sequence of execution for the threads in an OpenMP region can be accomplished with the **cancel** construct. The construct uses a *construct-type-clause* to set the region-type to activate for the cancellation. That is, inclusion of one of the *construct-type-clause* names **parallel**, **for**, **do**, **sections** or **taskgroup** on the directive line activates the corresponding region. The **cancel** construct is activated by the first encountering thread, and it continues execution at the end of the named region. The **cancel** construct is also a cancellation point for any other thread of the team to also continue execution at the end of the named region.

Also, once the specified region has been activated for cancellation any thread that encounnters a **cancellation point** construct with the same named region (*construct-type-clause*), continues execution at the end of the region.

For an activated **cancel taskgroup** construct, the tasks that belong to the taskgroup set of the innermost enclosing taskgroup region will be canceled.

A task that encounters the cancel taskgroup construct continues execution at the end of its task region. Any task of the taskgroup that has already begun execution will run to completion, unless it encounters a **cancellation point**; tasks that have not begun execution "may" be discarded as completed tasks.

CONTROL VARIABLES

Internal control variables (ICV) are used by implementations to hold values which control the execution of OpenMP regions. Control (and hence the ICVs) may be set as implementation defaults, or set and adjusted through environment variables, clauses, and API functions. Many of the ICV control values are accessible through API function calls. Also, initial ICV values are reported by the runtime if the **OMP_DISPLAY_ENV** environment variable has been set to **TRUE**.

NESTED CONSTRUCTS

Certain combinations of nested constructs are permitted, giving rise to a *combined* construct consisting of two or more constructs. These can be used when the two (or several) constructs would be used immediately in succession (closely nested). A combined construct can use the clauses of the component constructs without restrictions. A *composite* construct is a combined construct which has one or more clauses with (an often obviously) modified or restricted meaning, relative to when the constructs are uncombined.

Certain nestings are forbidden, and often the reasoning is obvious. Worksharing constructs cannot be nested, and the **barrier** construct cannot be nested inside a worksharing construct, or a **critical** construct. Also, **target** constructs cannot be nested.

The **parallel** construct can be nested, as well as the **task** construct. The parallel execution in the nested **parallel** construct(s) is control by the **OMP_NESTED** and **OMP_MAX_ACTIVE_LEVELS** environment variables, and the **omp_set_nested()** and **omp_set_max_active_levels()** functions.

More details on nesting can be found in the *Nesting of Regions* of the *Directives* chapter in the OpenMP Specifications document.

9.1 Conditional Compilation

C / C++

The following example illustrates the use of conditional compilation using the OpenMP macro **_OPENMP**. With OpenMP compilation, the **_OPENMP** macro becomes defined.

Example cond_comp.1.c

```
#include <stdio.h>

int main()
{

# ifdef _OPENMP
    printf("Compiled by an OpenMP-compliant implementation.\n");
# endif

    return 0;
}
```

C / C++

Fortran

The following example illustrates the use of the conditional compilation sentinel. With OpenMP compilation, the conditional compilation sentinel **!$** is recognized and treated as two spaces. In fixed form source, statements guarded by the sentinel must start after column 6.

Example cond_comp.1.f

```
      PROGRAM EXAMPLE

C234567890
!$    PRINT *, "Compiled by an OpenMP-compliant implementation."

      END PROGRAM EXAMPLE
```

Fortran

9.2 Internal Control Variables (ICVs)

According to Section 2.3 of the OpenMP 4.0 specification, an OpenMP implementation must act as if there are ICVs that control the behavior of the program. This example illustrates two ICVs, *nthreads-var* and *max-active-levels-var*. The *nthreads-var* ICV controls the number of threads requested for encountered parallel regions; there is one copy of this ICV per task. The *max-active-levels-var* ICV controls the maximum number of nested active parallel regions; there is one copy of this ICV for the whole program.

In the following example, the *nest-var*, *max-active-levels-var*, *dyn-var*, and *nthreads-var* ICVs are modified through calls to the runtime library routines **omp_set_nested**, **omp_set_max_active_levels**, **omp_set_dynamic**, and **omp_set_num_threads** respectively. These ICVs affect the operation of **parallel** regions. Each implicit task generated by a **parallel** region has its own copy of the *nest-var, dyn-var*, and *nthreads-var* ICVs.

In the following example, the new value of *nthreads-var* applies only to the implicit tasks that execute the call to **omp_set_num_threads**. There is one copy of the *max-active-levels-var* ICV for the whole program and its value is the same for all tasks. This example assumes that nested parallelism is supported.

The outer **parallel** region creates a team of two threads; each of the threads will execute one of the two implicit tasks generated by the outer **parallel** region.

Each implicit task generated by the outer **parallel** region calls **omp_set_num_threads(3)**, assigning the value 3 to its respective copy of *nthreads-var*. Then each implicit task encounters an inner **parallel** region that creates a team of three threads; each of the threads will execute one of the three implicit tasks generated by that inner **parallel** region.

Since the outer **parallel** region is executed by 2 threads, and the inner by 3, there will be a total of 6 implicit tasks generated by the two inner **parallel** regions.

Each implicit task generated by an inner **parallel** region will execute the call to **omp_set_num_threads(4)**, assigning the value 4 to its respective copy of *nthreads-var*.

The print statement in the outer **parallel** region is executed by only one of the threads in the team. So it will be executed only once.

The print statement in an inner **parallel** region is also executed by only one of the threads in the team. Since we have a total of two inner **parallel** regions, the print statement will be executed twice – once per inner **parallel** region.

C / C++

Example icv.1.c

```
#include <stdio.h>
#include <omp.h>

int main (void)
{
  omp_set_nested(1);
  omp_set_max_active_levels(8);
  omp_set_dynamic(0);
  omp_set_num_threads(2);
  #pragma omp parallel
    {
      omp_set_num_threads(3);

      #pragma omp parallel
        {
          omp_set_num_threads(4);
          #pragma omp single
            {
                 // The following should print:
                 // Inner: max_act_lev=8, num_thds=3, max_thds=4
                 // Inner: max_act_lev=8, num_thds=3, max_thds=4
              printf ("Inner: max_act_lev=%d, num_thds=%d, max_thds=%d\n",
              omp_get_max_active_levels(), omp_get_num_threads(),
              omp_get_max_threads());
            }
        }

      #pragma omp barrier
      #pragma omp single
        {
                 // The following should print:
                 // Outer: max_act_lev=8, num_thds=2, max_thds=3
          printf ("Outer: max_act_lev=%d, num_thds=%d, max_thds=%d\n",
                  omp_get_max_active_levels(), omp_get_num_threads(),
                  omp_get_max_threads());
        }
    }
    return 0;
}
```

C / C++

Fortran

Example icv.1.f

```
S-1       program icv
S-2       use omp_lib
S-3
S-4       call omp_set_nested(.true.)
S-5       call omp_set_max_active_levels(8)
S-6       call omp_set_dynamic(.false.)
S-7       call omp_set_num_threads(2)
S-8
S-9    !$omp parallel
S-10      call omp_set_num_threads(3)
S-11
S-12   !$omp parallel
S-13      call omp_set_num_threads(4)
S-14   !$omp single
S-15   !     The following should print:
S-16   !     Inner: max_act_lev= 8 , num_thds= 3 , max_thds= 4
S-17   !     Inner: max_act_lev= 8 , num_thds= 3 , max_thds= 4
S-18       print *, "Inner: max_act_lev=", omp_get_max_active_levels(),
S-19      &           ", num_thds=", omp_get_num_threads(),
S-20      &           ", max_thds=", omp_get_max_threads()
S-21   !$omp end single
S-22   !$omp end parallel
S-23
S-24   !$omp barrier
S-25   !$omp single
S-26   !     The following should print:
S-27   !     Outer: max_act_lev= 8 , num_thds= 2 , max_thds= 3
S-28       print *, "Outer: max_act_lev=", omp_get_max_active_levels(),
S-29      &           ", num_thds=", omp_get_num_threads(),
S-30      &           ", max_thds=", omp_get_max_threads()
S-31   !$omp end single
S-32   !$omp end parallel
S-33       end
```

Fortran

9.3 Placement of `flush`, `barrier`, `taskwait` and `taskyield` Directives

The following example is non-conforming, because the **`flush`**, **`barrier`**, **`taskwait`**, and **`taskyield`** directives are stand-alone directives and cannot be the immediate substatement of an **`if`** statement.

C / C++

Example standalone.1.c (**`omp_3.1`**)

```
S-1
S-2   void standalone_wrong()
S-3   {
S-4     int a = 1;
S-5
S-6           if (a != 0)
S-7     #pragma omp flush(a)
S-8   /* incorrect as flush cannot be immediate substatement
S-9      of if statement */
S-10
S-11          if (a != 0)
S-12    #pragma omp barrier
S-13  /* incorrect as barrier cannot be immediate substatement
S-14     of if statement */
S-15
S-16          if (a!=0)
S-17    #pragma omp taskyield
S-18  /* incorrect as taskyield cannot be immediate substatement of if statement
S-19  */
S-20
S-21          if (a != 0)
S-22    #pragma omp taskwait
S-23  /* incorrect as taskwait cannot be immediate substatement
S-24     of if statement */
S-25
S-26  }
```

C / C++

The following example is non-conforming, because the **flush**, **barrier**, **taskwait**, and **taskyield** directives are stand-alone directives and cannot be the action statement of an **if** statement or a labeled branch target.

Fortran

Example standalone.1.f90 (**omp_3.1**)

```
S-1
S-2
S-3     SUBROUTINE STANDALONE_WRONG()
S-4
S-5       INTEGER  A
S-6
S-7       A = 1
S-8
S-9       ! the FLUSH directive must not be the action statement
S-10      ! in an IF statement
S-11      IF (A .NE. 0) !$OMP FLUSH(A)
S-12
S-13      ! the BARRIER directive must not be the action statement
S-14      ! in an IF statement
S-15      IF (A .NE. 0) !$OMP BARRIER
S-16
S-17      ! the TASKWAIT directive must not be the action statement
S-18      ! in an IF statement
S-19      IF (A .NE. 0) !$OMP TASKWAIT
S-20
S-21      ! the TASKYIELD directive must not be the action statement
S-22      ! in an IF statement
S-23      IF (A .NE. 0) !$OMP TASKYIELD
S-24
S-25      GOTO 100
S-26
S-27      ! the FLUSH directive must not be a labeled branch target
S-28      ! statement
S-29      100 !$OMP FLUSH(A)
S-30      GOTO 200
S-31
S-32      ! the BARRIER directive must not be a labeled branch target
S-33      ! statement
S-34      200 !$OMP BARRIER
S-35      GOTO 300
S-36
S-37      ! the TASKWAIT directive must not be a labeled branch target
S-38      ! statement
S-39      300 !$OMP TASKWAIT
S-40      GOTO 400
```

```
S-41
S-42      ! the TASKYIELD directive must not be a labeled branch target
S-43      ! statement
S-44      400 !$OMP TASKYIELD
S-45
S-46    END SUBROUTINE
```

Fortran

The following version of the above example is conforming because the **flush**, **barrier**, **taskwait**, and **taskyield** directives are enclosed in a compound statement.

C / C++

Example standalone.2.c (**omp_3.1**)

```
S-1   void standalone_ok()
S-2   {
S-3     int a = 1;
S-4
S-5     #pragma omp parallel
S-6     {
S-7        if (a != 0) {
S-8     #pragma omp flush(a)
S-9        }
S-10       if (a != 0) {
S-11    #pragma omp barrier
S-12       }
S-13       if (a != 0) {
S-14    #pragma omp taskwait
S-15       }
S-16          if (a != 0) {
S-17    #pragma omp taskyield
S-18          }
S-19    }
S-20  }
```

C / C++

The following example is conforming because the **flush**, **barrier**, **taskwait**, and **taskyield** directives are enclosed in an **if** construct or follow the labeled branch target.

Fortran

Example standalone.2.f90 (**omp_3.1**)

```
S-1   SUBROUTINE STANDALONE_OK()
S-2     INTEGER  A
S-3     A = 1
S-4     IF (A .NE. 0) THEN
S-5       !$OMP FLUSH(A)
S-6     ENDIF
S-7     IF (A .NE. 0) THEN
S-8       !$OMP BARRIER
S-9     ENDIF
S-10    IF (A .NE. 0) THEN
S-11      !$OMP TASKWAIT
S-12    ENDIF
S-13    IF (A .NE. 0) THEN
S-14      !$OMP TASKYIELD
S-15    ENDIF
S-16    GOTO 100
S-17    100 CONTINUE
S-18    !$OMP FLUSH(A)
S-19    GOTO 200
S-20    200 CONTINUE
S-21    !$OMP BARRIER
S-22    GOTO 300
S-23    300 CONTINUE
S-24    !$OMP TASKWAIT
S-25    GOTO 400
S-26    400 CONTINUE
S-27    !$OMP TASKYIELD
S-28  END SUBROUTINE
```

Fortran

9.4 Cancellation Constructs

The following example shows how the **cancel** directive can be used to terminate an OpenMP region. Although the **cancel** construct terminates the OpenMP worksharing region, programmers must still track the exception through the pointer ex and issue a cancellation for the **parallel** region if an exception has been raised. The master thread checks the exception pointer to make sure that the exception is properly handled in the sequential part. If cancellation of the **parallel** region has been requested, some threads might have executed **phase_1()**. However, it is guaranteed that none of the threads executed **phase_2()**.

C++

Example cancellation.1.cpp (**omp_4.0**)

```
S-1   #include <iostream>
S-2   #include <exception>
S-3   #include <cstddef>
S-4
S-5   #define N 10000
S-6
S-7   extern void causes_an_exception();
S-8   extern void phase_1();
S-9   extern void phase_2();
S-10
S-11  void example() {
S-12      std::exception *ex = NULL;
S-13  #pragma omp parallel shared(ex)
S-14      {
S-15  #pragma omp for
S-16          for (int i = 0; i < N; i++) {
S-17              // no 'if' that prevents compiler optimizations
S-18              try {
S-19                  causes_an_exception();
S-20              }
S-21              catch (std::exception *e) {
S-22                  // still must remember exception for later handling
S-23  #pragma omp atomic write
S-24                  ex = e;
S-25                                  // cancel worksharing construct
S-26  #pragma omp cancel for
S-27              }
S-28          }
S-29     // if an exception has been raised, cancel parallel region
S-30          if (ex) {
S-31  #pragma omp cancel parallel
S-32          }
S-33          phase_1();
S-34  #pragma omp barrier
```

```
S-35            phase_2();
S-36        }
S-37        // continue here if an exception has been thrown in the worksharing loop
S-38        if (ex) {
S-39            // handle exception stored in ex
S-40        }
S-41    }
```

C++

The following example illustrates the use of the **cancel** construct in error handling. If there is an error condition from the **allocate** statement, the cancellation is activated. The encountering thread sets the shared variable **err** and other threads of the binding thread set proceed to the end of the worksharing construct after the cancellation has been activated.

Fortran

Example cancellation.1.f90 (**omp_4.0**)

```
S-1     subroutine example(n, dim)
S-2       integer, intent(in) :: n, dim(n)
S-3       integer :: i, s, err
S-4       real, allocatable :: B(:)
S-5       err = 0
S-6     !$omp parallel shared(err)
S-7     ! ...
S-8     !$omp do private(s, B)
S-9       do i=1, n
S-10    !$omp cancellation point do
S-11        allocate(B(dim(i)), stat=s)
S-12        if (s .gt. 0) then
S-13    !$omp atomic write
S-14          err = s
S-15    !$omp cancel do
S-16        endif
S-17    !   ...
S-18    ! deallocate private array B
S-19        if (allocated(B)) then
S-20          deallocate(B)
S-21        endif
S-22      enddo
S-23    !$omp end parallel
S-24    end subroutine
```

Fortran

The following example shows how to cancel a parallel search on a binary tree as soon as the search value has been detected. The code creates a task to descend into the child nodes of the current tree node. If the search value has been found, the code remembers the tree node with the found value through an **atomic** write to the result variable and then cancels execution of all search tasks. The function **search_tree_parallel** groups all search tasks into a single task group to control the effect of the **cancel taskgroup** directive. The *level* argument is used to create undeferred tasks after the first ten levels of the tree.

C / C++

Example cancellation.2.c (**omp_4.0**)

```
S-1   #include <stddef.h>
S-2
S-3   typedef struct binary_tree_s {
S-4      int value;
S-5      struct binary_tree_s *left, *right;
S-6   } binary_tree_t;
S-7
S-8   binary_tree_t *search_tree(binary_tree_t *tree, int value, int level) {
S-9       binary_tree_t *found = NULL;
S-10      if (tree) {
S-11          if (tree->value == value) {
S-12              found = tree;
S-13          }
S-14          else {
S-15  #pragma omp task shared(found) if(level < 10)
S-16              {
S-17                  binary_tree_t *found_left = NULL;
S-18                  found_left = search_tree(tree->left, value, level + 1);
S-19                  if (found_left) {
S-20  #pragma omp atomic write
S-21                      found = found_left;
S-22  #pragma omp cancel taskgroup
S-23                  }
S-24              }
S-25  #pragma omp task shared(found) if(level < 10)
S-26              {
S-27                  binary_tree_t *found_right = NULL;
S-28                  found_right = search_tree(tree->right, value, level + 1);
S-29                  if (found_right) {
S-30  #pragma omp atomic write
S-31                      found = found_right;
S-32  #pragma omp cancel taskgroup
S-33                  }
S-34              }
S-35  #pragma omp taskwait
S-36          }
```

```
S-37        }
S-38        return found;
S-39    }
S-40    binary_tree_t *search_tree_parallel(binary_tree_t *tree, int value) {
S-41        binary_tree_t *found = NULL;
S-42    #pragma omp parallel shared(found, tree, value)
S-43        {
S-44    #pragma omp master
S-45            {
S-46    #pragma omp taskgroup
S-47                {
S-48                    found = search_tree(tree, value, 0);
S-49                }
S-50            }
S-51        }
S-52        return found;
S-53    }
```

C / C++

1 The following is the equivalent parallel search example in Fortran.

Fortran

2 *Example cancellation.2.f90* (**omp_4.0**)

```
S-1     module parallel_search
S-2       type binary_tree
S-3         integer :: value
S-4         type(binary_tree), pointer :: right
S-5         type(binary_tree), pointer :: left
S-6       end type
S-7
S-8     contains
S-9       recursive subroutine search_tree(tree, value, level, found)
S-10        type(binary_tree), intent(in), pointer :: tree
S-11        integer, intent(in) :: value, level
S-12        type(binary_tree), pointer :: found
S-13        type(binary_tree), pointer :: found_left => NULL(), found_right => NULL()
S-14
S-15        if (associated(tree)) then
S-16          if (tree%value .eq. value) then
S-17            found => tree
S-18          else
S-19    !$omp task shared(found) if(level<10)
S-20            call search_tree(tree%left, value, level+1, found_left)
S-21            if (associated(found_left)) then
S-22    !$omp critical
S-23              found => found_left
```

```
S-24    !$omp end critical
S-25
S-26    !$omp cancel taskgroup
S-27            endif
S-28    !$omp end task
S-29
S-30    !$omp task shared(found) if(level<10)
S-31            call search_tree(tree%right, value, level+1, found_right)
S-32            if (associated(found_right)) then
S-33    !$omp critical
S-34              found => found_right
S-35    !$omp end critical
S-36
S-37    !$omp cancel taskgroup
S-38            endif
S-39    !$omp end task
S-40
S-41    !$omp taskwait
S-42          endif
S-43        endif
S-44      end subroutine
S-45
S-46      subroutine search_tree_parallel(tree, value, found)
S-47        type(binary_tree), intent(in), pointer :: tree
S-48        integer, intent(in) :: value
S-49        type(binary_tree), pointer :: found
S-50
S-51        found => NULL()
S-52    !$omp parallel shared(found, tree, value)
S-53    !$omp master
S-54    !$omp taskgroup
S-55        call search_tree(tree, value, 0, found)
S-56    !$omp end taskgroup
S-57    !$omp end master
S-58    !$omp end parallel
S-59      end subroutine
S-60
S-61    end module parallel_search
```

Fortran

9.5 The `requires` Directive

The declarative **`requires`** directive can be used to specify features that an implementation must provide to compile and execute correctly.

In the following example the **`unified_shared_memory`** clause of the **`requires`** directive ensures that the host and all devices accessible through OpenMP provide a *unified address* space for memory that is shared by all devices.

The example illustrates the use of the **`requires`** directive specifying *unified shared memory* in file scope, before any device directives or device routines. No **`map`** clause is needed for the *p* structure on the device (and its address *&p*, for the C++ code, is the same address on the host and device). However, scalar variables referenced within the **`target`** construct still have a default data-sharing attribute of firstprivate. The *q* scalar is incremented on the device, and its change is not updated on the host.

C++

Example requires.1.cpp (**`omp_5.0`**)

```
S-1
S-2   #include <iostream>
S-3   using namespace std;
S-4
S-5   #pragma omp requires unified_shared_memory
S-6
S-7   typedef struct mypoints
S-8   {
S-9      double res;
S-10     double data[500];
S-11  } mypoints_t;
S-12
S-13  void do_something_with_p(mypoints_t *p, int q);
S-14
S-15  int main()
S-16  {
S-17    mypoints_t p;
S-18    int q=0;
S-19
S-20    #pragma omp target // no map clauses needed
S-21    {                  // q is firstprivate
S-22       q++;
S-23       do_something_with_p(&p,q);
S-24    }
S-25    cout<< p.res << " " << q << endl;  // output 1 0
S-26    return 0;
S-27  }
S-28  void do_something_with_p(mypoints_t *p, int q)
S-29  {
```

```
  p->res = q;
  for(int i=0;i<sizeof(p->data)/sizeof(double);i++)
      p->data[i]=q*i;
}
```

C++

Fortran

Example requires.1.f90 (`omp_5.0`)

```
module data
!$omp requires unified_shared_memory
  type,public :: mypoints
     double precision :: res
     double precision :: data(500)
  end type
end module

program main
  use data
  type(mypoints) :: p
  integer        :: q=0

  !$omp target    !! no map clauses needed
     q = q + 1    !! q is firstprivate
     call do_something_with_p(p,q)
  !$omp end target

  write(*,'(f5.0,i5)') p%res, q    !! output 1.   0

end program

subroutine do_something_with_p(p,q)
  use data
  type(mypoints) :: p
  integer        :: q

  p%res = q;
  do i=1,size(p%data)
     p%data(i)=q*i
  enddo

end subroutine
```

Fortran

9.6 `declare variant` Directive

A **`declare variant`** directive specifies an alternate function, *function variant*, to be used in place of the *base function* when the trait within the **`match`** clause matches the OpenMP context at a given call site. The base function follows the directive in the C and C++ languages. In Fortran, either a subroutine or function may be used as the *base function*, and the **`declare variant`** directive must be in the specification part of a subroutine or function (unless a *base-proc-name* modifier is used, as in the case of a procedure declaration statement). See the OpenMP 5.0 Specification for details on the modifier.

When multiple **`declare variant`** directives are used a function variant becomes a candidate for replacing the base function if the context at the base function call matches the traits of all selectors in the **`match`** clause. If there are multiple candidates, a score is assigned with rules for each of the selector traits. The scoring algorithm can be found in the OpenMP 5.0 Specification.

In the first example the *vxv()* function is called within a **`parallel`** region, a **`target`** region, and in a sequential part of the program. Two function variants, *p_vxv()* and *t_vxv()*, are defined for the first two regions by using *parallel* and *target* selectors (within the *construct* trait set) in a **`match`** clause. The *p_vxv()* function variant includes a **`for`** construct (**`do`** construct for Fortran) for the **`parallel`** region, while *t_vxv()* includes a **`distribute simd`** construct for the **`target`** region. The *t_vxv()* function is explicitly compiled for the device using a **`declare target`** directive.

Since the two **`declare variant`** directives have no selectors that match traits for the context of the base function call in the sequential part of the program, the base *vxv()* function is used there, as expected. (The vectors in the *p_vxv* and *t_vxv* functions have been multiplied by 3 and 2, respectively, for checking the validity of the replacement. Normally the purpose of a function variant is to produce the same results by a different method.)

C / C++

Example declare_variant.1.c (**`omp_5.0`**)

```
S-1
S-2   #define N 100
S-3   #include <stdio.h>
S-4   #include <omp.h>
S-5
S-6   void p_vxv(int *v1,int *v2,int *v3,int n);
S-7   void t_vxv(int *v1,int *v2,int *v3,int n);
S-8
S-9   #pragma omp declare variant( p_vxv ) match( construct={parallel} )
S-10  #pragma omp declare variant( t_vxv ) match( construct={target}   )
S-11  void vxv(int *v1,int *v2,int *v3,int n)     // base function
S-12  {
S-13     for (int i= 0; i< n; i++)  v3[i] = v1[i] * v2[i];
S-14  }
S-15
```

```
S-16    void p_vxv(int *v1,int *v2,int *v3,int n)   // function variant
S-17    {
S-18       #pragma omp for
S-19       for (int i= 0; i< n; i++)  v3[i] = v1[i] * v2[i]*3;
S-20    }
S-21
S-22    #pragma omp declare target
S-23    void t_vxv(int *v1,int *v2,int *v3,int n)   // function variant
S-24    {
S-25       #pragma omp distribute simd
S-26       for (int i= 0; i< n; i++)  v3[i] = v1[i] * v2[i]*2;
S-27    }
S-28    #pragma omp end declare target
S-29
S-30    int main()
S-31    {
S-32       int v1[N], v2[N], v3[N];
S-33       for(int i=0; i<N; i++){ v1[i]=(i+1); v2[i]=-(i+1); v3[i]=0; }   //init
S-34
S-35       #pragma omp parallel
S-36       {
S-37          vxv(v1,v2,v3,N);
S-38       }
S-39       printf(" %d  %d\n",v3[0],v3[N-1]); //from p_vxv --  output: -3  -30000
S-40
S-41       #pragma omp target teams map(to: v1[:N],v2[:N]) map(from: v3[:N])
S-42       {
S-43          vxv(v1,v2,v3,N);
S-44       }
S-45       printf(" %d  %d\n",v3[0],v3[N-1]); //from t_vxv --  output: -2  -20000
S-46
S-47       vxv(v1,v2,v3,N);
S-48       printf(" %d  %d\n",v3[0],v3[N-1]); //from   vxv --  output: -1  -10000
S-49
S-50       return 0;
S-51    }
```

C / C++

Fortran

1 *Example declare_variant.1.f90* (`omp_5.0`)

```
S-1
S-2     module subs
S-3       use omp_lib
S-4     contains
S-5        subroutine vxv(v1, v2, v3)             !! base function
S-6           integer,intent(in)  :: v1(:),v2(:)
```

```
S-7          integer,intent(out) :: v3(:)
S-8          integer             :: i,n
S-9          !$omp  declare variant( p_vxv ) match( construct={parallel} )
S-10         !$omp  declare variant( t_vxv ) match( construct={target}   )
S-11
S-12         n=size(v1)
S-13         do i = 1,n; v3(i) = v1(i) * v2(i); enddo
S-14
S-15      end subroutine
S-16
S-17      subroutine p_vxv(v1, v2, v3)            !! function variant
S-18         integer,intent(in)  :: v1(:),v2(:)
S-19         integer,intent(out) :: v3(:)
S-20         integer             :: i,n
S-21         n=size(v1)
S-22
S-23         !$omp do
S-24         do i = 1,n; v3(i) = v1(i) * v2(i) * 3; enddo
S-25
S-26      end subroutine
S-27
S-28      subroutine t_vxv(v1, v2, v3)            !! function variant
S-29         integer,intent(in)  :: v1(:),v2(:)
S-30         integer,intent(out) :: v3(:)
S-31         integer             :: i,n
S-32         !$omp declare target
S-33         n=size(v1)
S-34
S-35         !$omp distribute simd
S-36         do i = 1,n; v3(i) = v1(i) * v2(i) * 2; enddo
S-37
S-38      end subroutine
S-39
S-40   end module subs
S-41
S-42
S-43   program main
S-44      use omp_lib
S-45      use subs
S-46      integer,parameter :: N = 100
S-47      integer           :: v1(N), v2(N), v3(N)
S-48
S-49      do i= 1,N; v1(i)= i; v2(i)= -i; v3(i)= 0;  enddo  !! init
S-50
S-51      !$omp parallel
S-52         call vxv(v1,v2,v3)
S-53      !$omp end parallel
```

```
   print *, v3(1),v3(N)    !! from p_vxv -- output: -3  -30000

   !$omp target teams map(to: v1,v2) map(from: v3)
      call vxv(v1,v2,v3)
   !$omp end target teams
   print *, v3(1),v3(N)    !! from t_vxv -- output: -2  -20000

   call vxv(v1,v2,v3)
   print *, v3(1),v3(N)    !! from   vxv -- output: -1  -10000

end program
```

Fortran

In this example, traits from the *device* set are used to select a function variant. In the **declare variant** directive, an *isa* selector specifies that if the implementation of the "*core-avx512*" instruction set is detected at compile time the *avx512_saxpy()* variant function is used for the call to *base_saxpy()*.

A compilation of *avx512_saxpy()* is aware of the AVX-512 instruction set that supports 512-bit vector extensions (for Xeon or Xeon Phi architectures). Within *avx512_saxpy()*, the **parallel for simd** construct performs parallel execution, and takes advantage of 64-byte data alignment. When the *avx512_saxpy()* function variant is not selected, the base *base_saxpy()* function variant containing only a basic **parallel for** construct is used for the call to *base_saxpy()*.

C / C++

Example declare_variant.2.c (**omp_5.0**)

```
#include <omp.h>

void   base_saxpy(int, float, float *, float *);
void avx512_saxpy(int, float, float *, float *);

#pragma omp declare variant( avx512_saxpy ) \
                        match( device={isa("core-avx512")} )
void base_saxpy(int n, float s, float *x, float *y)    // base function
{
   #pragma omp parallel for
   for(int i=0; i<n; i++) y[i] = s*x[i] + y[i];
}

void avx512_saxpy(int n, float s, float *x, float *y)    //function variant
{
                                       //assume 64-byte alignment for AVX-512
   #pragma omp parallel for simd simdlen(16) aligned(x,y:64)
   for(int i=0; i<n; i++) y[i] = s*x[i] + y[i];
}
```

```
S-20
S-21    // Above may be in another file scope.
S-22
S-23    #include <stdio.h>
S-24    #include <stdlib.h>
S-25    #include <stdint.h>
S-26    #define N 1000
S-27
S-28    int main()
S-29    {
S-30       static float x[N],y[N] __attribute__ ((aligned(64)));
S-31       float s=2.0;
S-32                              // Check for 64-byte aligned
S-33       if( ((intptr_t)y)%64 != 0 || ((intptr_t)x)%64 != 0 )
S-34       { printf("ERROR: x|y not 64-Byte aligned\n"); exit(1); }
S-35
S-36       for(int i=0; i<N; i++){ x[i]=i+1; y[i]=i+1; } // initialize
S-37
S-38       base_saxpy(N,s,x,y);
S-39
S-40       printf("y[0],y[N-1]: %5.0f %5.0f\n",y[0],y[N-1]); //output: y...   3  3000
S-41
S-42       return 0;
S-43    }
```

C / C++

Fortran

1 *Example declare_variant.2.f90* (`omp_5.0`)

```
S-1
S-2     module subs
S-3       use omp_lib
S-4     contains
S-5
S-6        subroutine base_saxpy(s,x,y)               !! base function
S-7           real,intent(inout) :: s,x(:),y(:)
S-8          !$omp  declare variant( avx512_saxpy ) &
S-9          !$omp&           match( device={isa("core-avx512")} )
S-10
S-11          y = s*x + y
S-12
S-13       end subroutine
S-14
S-15       subroutine avx512_saxpy(s,x,y)               !! function variant
S-16          real,intent(inout) :: s,x(:),y(:)
S-17          integer            :: i,n
S-18          n=size(x)
```

```
                                !!assume 64-byte alignment for AVX-512
      !$omp parallel do simd simdlen(16) aligned(x,y: 64)
      do i = 1,n
         y(i) = s*x(i) + y(i)
      end do

   end subroutine

end module subs


program main
   use omp_lib
   use subs

   integer, parameter :: N=1000, align=64
   real, allocatable  :: x(:),y(:)
   real               :: s = 2.0e0
   integer            :: i

   allocate(x(N),y(N))    !! Assumes allocation is 64-byte aligned
                          !! (using compiler options, or another
                          !! allocation method).

                          !! loc is non-standard, but found everywhere
                          !! remove these lines if not available
   if(modulo(loc(x),align) /= 0 .and. modulo(loc(y),align) /=0 ) then
      print*,"ERROR: x|y not 64-byte aligned"; stop
   endif

   do i=1,N  !! initialize
     x(i)=i
     y(i)=i
   end do

   call base_saxpy(s,x,y)

   write(*,'("y(1),y(N):",2f6.0)') y(1),y(N) !!output: y... 3. 3000.

   deallocate(x,y)

end program
```

Fortran

9.7 Metadirective Directive

A **metadirective** directive provides a mechanism to select a directive in a **when** clause to be used, depending upon one or more contexts: implementation, available devices and the present enclosing construct. The directive in a **default** clause is used when a directive of the **when** clause is not selected.

In the **when** clause the *context selector* (or just *selector*) defines traits that are evaluated for selection of the directive that follows the selector. This "selectable" directive is called a *directive variant*. Traits are grouped by *construct*, *implementation* and *device sets* to be used by a selector of the same name.

In the first example the architecture trait *arch* of the *device* selector set specifies that if an *nvptx* (NVIDIA) architecture is active in the OpenMP context, then the **teams loop** *directive variant* is selected as the directive; otherwise, the **parallel loop** *directive variant* of the **default** clause is selected as the directive. That is, if a *device* of *nvptx* architecture is supported by the implementation within the enclosing **target** construct, its *directive variant* is selected. The architecture names, such as *nvptx*, are implementation defined. Also, note that *device* as used in a **target** construct specifies a device number, while *device*, as used in the **metadirective** directive as selector set, has traits of *kind*, *isa* and *arch*.

C / C++

Example metadirective.1.c (**omp_5.0**)

```
S-1
S-2   #define N 100
S-3   #include <stdio.h>
S-4
S-5   int main()
S-6   {
S-7      int v1[N], v2[N], v3[N];
S-8      for(int i=0; i<N; i++){ v1[i]=(i+1); v2[i]=-(i+1); }
S-9
S-10     #pragma omp target map(to:v1,v2) map(from:v3) device(0)
S-11     #pragma omp metadirective \
S-12                     when(   device={arch("nvptx")}: teams loop) \
S-13                     default(                        parallel loop)
S-14       for (int i= 0; i< N; i++)  v3[i] = v1[i] * v2[i];
S-15
S-16     printf(" %d  %d\n",v3[0],v3[N-1]); //output: -1  -10000
S-17
S-18     return 0;
S-19  }
```

C / C++

Fortran

Example metadirective.1.f90 (**omp_5.0**)

```
S-1
S-2   program main
S-3      integer, parameter :: N= 100
S-4      integer ::  v1(N), v2(N), v3(N);
S-5
S-6      do i=1,N;  v1(i)=i; v2(i)=-i;  enddo   ! initialize
S-7
S-8      !$omp  target map(to:v1,v2) map(from:v3) device(0)
S-9      !$omp  metadirective &
S-10     !$omp&     when(   device={arch("nvptx")}: teams loop) &
S-11     !$omp&     default(                      parallel loop)
S-12       do i= 1,N; v3(i) = v1(i) * v2(i); enddo
S-13     !$omp  end target
S-14
S-15     print *, v3(1),v3(N) !!output: -1  -10000
S-16  end program
```

Fortran

In the second example, the *implementation* selector set is specified in the **when** clause to distinguish between AMD and NVIDIA platforms. Additionally, specific architectures are specified with the *device* selector set.

In the code, different **teams** constructs are employed as determined by the **metadirective** directive. The number of teams is restricted by a **num_teams** clause and a thread limit is also set by a **thread_limit** clause for *vendor* AMD and NVIDIA platforms and specific architecture traits. Otherwise, just the **teams** construct is used without any clauses, as prescribed by the **default** clause.

C / C++

Example metadirective.2.c (**omp_5.0**)

```
S-1
S-2   #define N 100
S-3   #include <stdio.h>
S-4   #include <omp.h>
S-5
S-6   void work_on_chunk(int idev, int i);
S-7
S-8   int main()                     //Driver
S-9   {
S-10     int i,idev;
S-11
S-12     for (idev=0; idev<omp_get_num_devices(); idev++)
S-13     {
```

```
S-14          #pragma omp target device(idev)
S-15          #pragma omp metadirective \
S-16                      when( implementation={vendor(nvidia)}, device={arch("kepler")}:
S-17                            teams num_teams(512) thread_limit(32) )
S-18                      when( implementation={vendor(amd)},    device={arch("fiji"  )}:
S-19                            teams num_teams(512) thread_limit(64) )
S-20                      default(
S-21                            teams)
S-22          #pragma omp distribute parallel for
S-23          for (i=0; i<N; i++) work_on_chunk(idev,i);
S-24       }
S-25       return 0;
S-26    }
S-27
```

C / C++

Fortran

1 *Example metadirective.2.f90* (`omp_5.0`)

```
 S-1
 S-2    program main                    !!Driver
 S-3      use omp_lib
 S-4      implicit none
 S-5      integer, parameter :: N=1000
 S-6      external           :: work_on_chunk
 S-7      integer            :: i,idev
 S-8
 S-9      do idev=0,omp_get_num_devices()-1
S-10
S-11        !$omp target device(idev)
S-12        !$omp begin metadirective &
S-13        !$omp&  when( implementation={vendor(nvidia)}, device={arch("kepler")}: &
S-14        !$omp&        teams num_teams(512) thread_limit(32) )                   &
S-15        !$omp&  when( implementation={vendor(amd)},    device={arch("fiji"  )}: &
S-16        !$omp&        teams num_teams(512) thread_limit(64) )                   &
S-17        !$omp&  default(                                                        &
S-18        !$omp&        teams)
S-19        !$omp distribute parallel for
S-20        do i=1,N
S-21           call work_on_chunk(idev,i)
S-22        end do
S-23        !$omp end metadirective
S-24        !$omp end target
S-25
S-26      end do
S-27
```

```
S-28     end program
S-29
```

Fortran

In the third example, a *construct* selector set is specified in the **when** clause. Here, a **metadirective** directive is used within a function that is also compiled as a function for a target device as directed by the **declare target** directive. The *target* directive name of the **construct** selector ensures that the **distribute parallel for/do** construct is employed for the target compilation. Otherwise, for the host-compiled version the **parallel for/do simd** construct is used.

In the first call to the *exp_pi_diff()* routine the context is a **target teams** construct and the **distribute parallel for/do** construct version of the function is invoked, while in the second call the **parallel for/do simd** construct version is used.

This case illustrates an important point for users that may want to hoist the **target** directive out of a function that contains the usual **target teams distribute parallel for/do** construct (for providing alternate constructs through the **metadirective** directive as here). While this combined construct can be decomposed into a **target** and **teams distribute parallel for/do** constructs, the OpenMP 5.0 specification has the restriction: "If a **teams** construct is nested within a **target** construct, that **target** construct must contain no statements, declarations or directives outside of the **teams** construct". So, the **teams** construct must immediately follow the **target** construct without any intervening code statements (which includes function calls). Since the **target** construct alone cannot be hoisted out of a function, the **target teams** construct has been hoisted out of the function, and the **distribute parallel for/do** construct is used as the *variant* directive of the **metadirective** directive within the function.

C / C++

Example metadirective.3.c (**omp_5.0**)

```
S-1   #include <stdio.h>
S-2   #include  <math.h>
S-3   #define      N 1000
S-4
S-5   #pragma omp declare target
S-6   void exp_pi_diff(double *d, double my_pi){
S-7      #pragma omp metadirective \
S-8                  when( construct={target}: distribute parallel for  ) \
S-9                  default(                              parallel for simd)
S-10     for(int i = 0; i<N; i++) d[i] = exp( (M_PI-my_pi)*i );
S-11  }
S-12  #pragma omp end declare target
S-13
S-14  int main()
S-15  {
S-16    //Calculates sequence of exponentials: (M_PI-my_pi) * index
S-17    //M_PI is from math.h, and my_pi is user provided.
S-18
S-19    double d[N];
S-20    double my_pi=3.14159265358979e0;
S-21
```

```
S-22        #pragma omp target teams map(tofrom: d[0:N])
S-23        exp_pi_diff(d,my_pi);
S-24                                              // value should be near 1
S-25        printf("d[N-1] = %20.14f\n",d[N-1]); // ...= 1.00000000000311
S-26
S-27        exp_pi_diff(d,my_pi);                 // value should be near 1
S-28        printf("d[N-1] = %20.14f\n",d[N-1]); // ...= 1.00000000000311
S-29    }
```

C / C++

Fortran

Example metadirective.3.f90 (`omp_5.0`)

```
S-1
S-2     module params
S-3        integer, parameter :: N=1000
S-4        DOUBLE PRECISION, PARAMETER::M_PI=4.0d0*DATAN(1.0d0) !3.1415926535897932_8
S-5     end module
S-6
S-7
S-8     subroutine exp_pi_diff(d,    my_pi)
S-9       use params
S-10      implicit none
S-11      integer          ::  i
S-12      double precision ::  d(N), my_pi
S-13      !$omp declare target
S-14
S-15      !$omp   metadirective &
S-16      !$omp&      when( construct={target}: distribute parallel do  )  &
S-17      !$omp&      default(                              parallel do simd)
S-18
S-19      do i = 1,size(d)
S-20         d(i) = exp( (M_PI-my_pi)*i )
S-21      end do
S-22
S-23    end subroutine
S-24
S-25    program main
S-26      ! Calculates sequence of exponentials: (M_PI-my_pi) * index
S-27      ! M_PI is from usual way, and my_pi is user provided.
S-28      ! Fortran Standard does not provide PI
S-29
S-30      use params
S-31      implicit none
S-32      double precision   :: d(N)
S-33      double precision   :: my_pi=3.14159265358979d0
S-34
```

```
S-35          !$omp target teams map(from: d)
S-36          call exp_pi_diff(d,my_pi)
S-37          !$omp end target teams
S-38                                        ! value should be near 1
S-39          print*, "d(N) = ",d(N)        ! 1.00000000000311
S-40
S-41          call exp_pi_diff(d,my_pi) ! value should be near 1
S-42          print*, "d(N) = ",d(N)        ! 1.00000000000311
S-43
S-44    end program
```

Fortran

9.8 Nested Loop Constructs

The following example of loop construct nesting is conforming because the inner and outer loop regions bind to different **parallel** regions:

C / C++

Example nested_loop.1.c

```
S-1   void work(int i, int j) {}
S-2
S-3   void good_nesting(int n)
S-4   {
S-5     int i, j;
S-6     #pragma omp parallel default(shared)
S-7     {
S-8       #pragma omp for
S-9       for (i=0; i<n; i++) {
S-10        #pragma omp parallel shared(i, n)
S-11        {
S-12          #pragma omp for
S-13          for (j=0; j < n; j++)
S-14            work(i, j);
S-15        }
S-16      }
S-17    }
S-18  }
```

C / C++

Fortran

Example nested_loop.1.f

```
S-1         SUBROUTINE WORK(I, J)
S-2         INTEGER I, J
S-3         END SUBROUTINE WORK
S-4
S-5         SUBROUTINE GOOD_NESTING(N)
S-6         INTEGER N
S-7
S-8           INTEGER I
S-9   !$OMP   PARALLEL DEFAULT(SHARED)
S-10  !$OMP     DO
S-11            DO I = 1, N
S-12  !$OMP       PARALLEL SHARED(I,N)
S-13  !$OMP         DO
S-14                DO J = 1, N
S-15                  CALL WORK(I,J)
S-16                END DO
```

```
!$OMP       END PARALLEL
          END DO
!$OMP   END PARALLEL
      END SUBROUTINE GOOD_NESTING
```

Fortran

The following variation of the preceding example is also conforming:

C / C++

Example nested_loop.2.c

```
void work(int i, int j) {}


void work1(int i, int n)
{
  int j;
  #pragma omp parallel default(shared)
  {
    #pragma omp for
    for (j=0; j<n; j++)
      work(i, j);
  }
}


void good_nesting2(int n)
{
  int i;
  #pragma omp parallel default(shared)
  {
    #pragma omp for
    for (i=0; i<n; i++)
      work1(i, n);
  }
}
```

C / C++

Fortran

Example nested_loop.2.f

```
      SUBROUTINE WORK(I, J)
      INTEGER I, J
      END SUBROUTINE WORK

      SUBROUTINE WORK1(I, N)
      INTEGER J
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
        DO J = 1, N
          CALL WORK(I,J)
        END DO
!$OMP END PARALLEL
      END SUBROUTINE WORK1

      SUBROUTINE GOOD_NESTING2(N)
      INTEGER N
!$OMP PARALLEL DEFAULT(SHARED)
!$OMP DO
      DO I = 1, N
         CALL WORK1(I, N)
      END DO
!$OMP END PARALLEL
      END SUBROUTINE GOOD_NESTING2
```

Fortran

9.9 Restrictions on Nesting of Regions

The examples in this section illustrate the region nesting rules.

The following example is non-conforming because the inner and outer loop regions are closely nested:

C / C++

Example nesting_restrict.1.c

```
S-1
S-2   void work(int i, int j) {}
S-3
S-4   void wrong1(int n)
S-5   {
S-6
S-7     #pragma omp parallel default(shared)
S-8     {
S-9       int i, j;
S-10      #pragma omp for
S-11      for (i=0; i<n; i++) {
S-12         /* incorrect nesting of loop regions */
S-13         #pragma omp for
S-14           for (j=0; j<n; j++)
S-15             work(i, j);
S-16      }
S-17    }
S-18
S-19  }
```

C / C++

Fortran

Example nesting_restrict.1.f

```
S-1
S-2         SUBROUTINE WORK(I, J)
S-3         INTEGER I, J
S-4
S-5         END SUBROUTINE WORK
S-6
S-7         SUBROUTINE WRONG1(N)
S-8
S-9         INTEGER N
S-10        INTEGER I,J
S-11  !$OMP   PARALLEL DEFAULT(SHARED)
S-12  !$OMP     DO
S-13            DO I = 1, N
S-14  !$OMP       DO             ! incorrect nesting of loop regions
```

```
S-15                DO J = 1, N
S-16                  CALL WORK(I,J)
S-17                END DO
S-18              END DO
S-19     !$OMP    END PARALLEL
S-20
S-21           END SUBROUTINE WRONG1
```

Fortran

1 The following orphaned version of the preceding example is also non-conforming:

C / C++

2 *Example nesting_restrict.2.c*

```
S-1   void work(int i, int j) {}
S-2   void work1(int i, int n)
S-3   {
S-4     int j;
S-5   /* incorrect nesting of loop regions */
S-6     #pragma omp for
S-7       for (j=0; j<n; j++)
S-8         work(i, j);
S-9   }
S-10
S-11  void wrong2(int n)
S-12  {
S-13    #pragma omp parallel default(shared)
S-14    {
S-15      int i;
S-16      #pragma omp for
S-17        for (i=0; i<n; i++)
S-18           work1(i, n);
S-19    }
S-20  }
```

C / C++

Fortran

Example nesting_restrict.2.f

```
S-1         SUBROUTINE WORK1(I,N)
S-2         INTEGER I, N
S-3         INTEGER J
S-4   !$OMP   DO      ! incorrect nesting of loop regions
S-5           DO J = 1, N
S-6             CALL WORK(I,J)
S-7           END DO
S-8         END SUBROUTINE WORK1
S-9         SUBROUTINE WRONG2(N)
S-10        INTEGER N
S-11        INTEGER I
S-12  !$OMP   PARALLEL DEFAULT(SHARED)
S-13  !$OMP     DO
S-14            DO I = 1, N
S-15              CALL WORK1(I,N)
S-16            END DO
S-17  !$OMP   END PARALLEL
S-18        END SUBROUTINE WRONG2
```

Fortran

The following example is non-conforming because the loop and **single** regions are closely nested:

C / C++

Example nesting_restrict.3.c

```
S-1   void work(int i, int j) {}
S-2   void wrong3(int n)
S-3   {
S-4     #pragma omp parallel default(shared)
S-5     {
S-6       int i;
S-7       #pragma omp for
S-8         for (i=0; i<n; i++) {
S-9   /* incorrect nesting of regions */
S-10          #pragma omp single
S-11            work(i, 0);
S-12        }
S-13    }
S-14  }
```

C / C++

Fortran

Example nesting_restrict.3.f

```
      SUBROUTINE WRONG3(N)
      INTEGER N

        INTEGER I
!$OMP   PARALLEL DEFAULT(SHARED)
!$OMP     DO
          DO I = 1, N
!$OMP       SINGLE            ! incorrect nesting of regions
              CALL WORK(I, 1)
!$OMP       END SINGLE
          END DO
!$OMP   END PARALLEL
      END SUBROUTINE WRONG3
```

Fortran

The following example is non-conforming because a **barrier** region cannot be closely nested inside a loop region:

C / C++

Example nesting_restrict.4.c

```
void work(int i, int j) {}
void wrong4(int n)
{

  #pragma omp parallel default(shared)
  {
    int i;
    #pragma omp for
      for (i=0; i<n; i++) {
        work(i, 0);
/* incorrect nesting of barrier region in a loop region */
        #pragma omp barrier
        work(i, 1);
      }
  }
}
```

C / C++

Fortran

Example nesting_restrict.4.f

```
      SUBROUTINE WRONG4(N)
      INTEGER N

        INTEGER I
!$OMP   PARALLEL DEFAULT(SHARED)
!$OMP     DO
          DO I = 1, N
            CALL WORK(I, 1)
! incorrect nesting of barrier region in a loop region
!$OMP       BARRIER
            CALL WORK(I, 2)
          END DO
!$OMP   END PARALLEL
      END SUBROUTINE WRONG4
```

Fortran

The following example is non-conforming because the **barrier** region cannot be closely nested inside the **critical** region. If this were permitted, it would result in deadlock due to the fact that only one thread at a time can enter the **critical** region:

C / C++

Example nesting_restrict.5.c

```
void work(int i, int j) {}
void wrong5(int n)
{
  #pragma omp parallel
  {
    #pragma omp critical
    {
       work(n, 0);
/* incorrect nesting of barrier region in a critical region */
       #pragma omp barrier
       work(n, 1);
    }
  }
}
```

C / C++

Fortran

Example nesting_restrict.5.f

```
      SUBROUTINE WRONG5(N)
      INTEGER N

!$OMP   PARALLEL DEFAULT(SHARED)
!$OMP     CRITICAL
            CALL WORK(N,1)
! incorrect nesting of barrier region in a critical region
!$OMP       BARRIER
            CALL WORK(N,2)
!$OMP     END CRITICAL
!$OMP   END PARALLEL
      END SUBROUTINE WRONG5
```

Fortran

The following example is non-conforming because the **barrier** region cannot be closely nested inside the **single** region. If this were permitted, it would result in deadlock due to the fact that only one thread executes the **single** region:

C / C++

Example nesting_restrict.6.c

```
void work(int i, int j) {}
void wrong6(int n)
{
  #pragma omp parallel
  {
    #pragma omp single
    {
      work(n, 0);
/* incorrect nesting of barrier region in a single region */
      #pragma omp barrier
      work(n, 1);
    }
  }
}
```

C / C++

Fortran

Example nesting_restrict.6.f

```
S-1         SUBROUTINE WRONG6(N)
S-2         INTEGER N
S-3
S-4   !$OMP   PARALLEL DEFAULT(SHARED)
S-5   !$OMP     SINGLE
S-6               CALL WORK(N,1)
S-7   ! incorrect nesting of barrier region in a single region
S-8   !$OMP       BARRIER
S-9               CALL WORK(N,2)
S-10  !$OMP     END SINGLE
S-11  !$OMP   END PARALLEL
S-12        END SUBROUTINE WRONG6
```

Fortran

9.10 Target Offload

In the OpenMP 5.0 implementation the **OMP_TARGET_OFFLOAD** environment variable was defined to change *default* offload behavior. By *default* the target code (region) is executed on the host if the target device does not exist or the implementation does not support the target device.

In an OpenMP 5.0 compliant implementation, setting the **OMP_TARGET_OFFLOAD** variable to **MANDATORY** will force the program to terminate execution when a **target** construct is encountered and the target device is not supported or is not available. With a value **DEFAULT** the target region will execute on a device if the device exists and is supported by the implementation, otherwise it will execute on the host. Support for the **DISABLED** value is optional; when it is supported the behavior is as if only the host device exists (other devices are considered non-existent to the runtime), and target regions are executed on the host.

The following example reports execution behavior for different values of the **OMP_TARGET_OFFLOAD** variable. A handy routine for extracting the **OMP_TARGET_OFFLOAD** environment variable value is deployed here, because the OpenMP API does not have a routine for obtaining the value.

Note: The example issues a warning when a pre-5.0 implementation is used, indicating that the **OMP_TARGET_OFFLOAD** is ignored. The value of the **OMP_TARGET_OFFLOAD** variable is reported when the **OMP_DISPLAY_ENV** environment variable is set to **TRUE** or **VERBOSE**.

C / C++

Example target_offload_control.1.c (**omp_5.0**)

```
S-1
S-2   #include    <omp.h>
S-3   #include  <stdio.h>
S-4   #include  <ctype.h>
S-5   #include <stdlib.h>
S-6   #include <string.h>
S-7
S-8   typedef enum offload_policy
S-9   {MANDATORY, DISABLED, DEFAULT, UNKNOWN, NOTSET} offload_policy_t;
S-10
S-11
S-12  offload_policy_t get_offload_policy()
S-13  {
S-14     char *env, *end;
S-15     size_t n;
S-16
S-17     env = getenv("OMP_TARGET_OFFLOAD");
S-18     if(env  == NULL) return NOTSET;
S-19
S-20     end = env + strlen(env);                //Find trimmed beginning/end
S-21     while (      *env && isspace(*(env  )) ) env++;
```

```
S-22        while (end != env && isspace(*(end-1)) ) end--;
S-23        n = (int)(end - env);
S-24
S-25                          //Find ONLY string -nothing more, case insensitive
S-26        if      (n == 9 && !strncasecmp(env, "MANDATORY",n)) return MANDATORY;
S-27        else if (n == 8 && !strncasecmp(env, "DISABLED" ,n)) return DISABLED ;
S-28        else if (n == 7 && !strncasecmp(env, "DEFAULT"  ,n)) return DEFAULT  ;
S-29        else                                                  return UNKNOWN  ;
S-30     }
S-31
S-32
S-33     int main()
S-34     {
S-35        int i;
S-36        int device_num, on_init_dev;
S-37
S-38                                    // get policy from OMP_TARGET_OFFLOAD variable
S-39        offload_policy_t policy = get_offload_policy();
S-40
S-41        if(_OPENMP< 201811)
S-42        {
S-43           printf("Warning: OMP_TARGET_OFFLOAD NOT supported by VER. %d\n",_OPENMP );
S-44           printf("         If OMP_TARGET_OFFLOAD is set, it will be ignored.\n");
S-45        }
S-46
S-47                                    // Set target device number to an unavailable
S-48                                    // device to test offload policy.
S-49        device_num = omp_get_num_devices() + 1;
S-50
S-51                                         // Policy:
S-52        printf("OMP_TARGET_OFFLOAD Policy:  ");
S-53        if     (policy==MANDATORY) printf("MANDATORY-Terminate if dev. not avail\n");
S-54        else if(policy==DISABLED ) printf("DISABLED -(if supported) Only on Host\n");
S-55        else if(policy==DEFAULT  ) printf("DEFAULT  -On host if device not avail\n");
S-56        else if(policy==UNKNOWN  ) printf("OMP_TARGET_OFFLOAD has unknown value\n" );
S-57        else if(policy==NOTSET   ) printf("OMP_TARGET_OFFLOAD not set\n" );
S-58
S-59
S-60        on_init_dev = 1;
S-61                                               // device# out of range--not supported
S-62        #pragma omp target device(device_num) map(tofrom: on_init_dev)
S-63          on_init_dev=omp_is_initial_device();
S-64
S-65        if (policy == MANDATORY && _OPENMP >= 201811)
S-66           printf("ERROR: OpenMP 5.0 implementation ignored MANDATORY policy.\n");
S-67
S-68        printf("Target region executed on init dev %s\n", on_init_dev ? "TRUE":"FALSE"
```

```
   return 0;
}
```

C / C++

Fortran

Example target_offload_control.1.f90 (`omp_5.0`)

```
module offload_policy
  implicit none
  integer, parameter :: LEN_POLICY=10
contains
  character(LEN_POLICY) function get_offload_policy()
     character(64) :: env
     integer       :: length, i
     env=repeat(' ',len(env))
                                     !policy is blank if not found *
     call get_environment_variable("OMP_TARGET_OFFLOAD",env,length)

     do i = 1,len(env)               !Makes a-z upper case
        if(iachar(env(i:i))>96) env(i:i)=achar(iachar(env(i:i))-32)
     end do

     get_offload_policy = trim(adjustl(env)) !remove peripheral spaces

     if(length==0) get_offload_policy="NOTSET"

     return

  end function

end module

program policy_test

  use omp_lib
  use offload_policy

  integer                :: i, device_num
  logical                :: on_init_dev
  character(LEN_POLICY)  :: policy

  policy = get_offload_policy() !!Get OMP_TARGET_OFFLOAD value

  if (OPENMP_VERSION < 201811) then
```

```
S-40        print*,"Warning: OMP_TARGET_OFFLOAD NOT supported by VER.",OPENMP_VERSION
S-41        print*,"         If OMP_TARGET_OFFLOAD is set, it will be ignored."
S-42     endif
S-43
S-44        !Set target device number to an unavailable device to test offload policy.
S-45     device_num = omp_get_num_devices() + 1
S-46
S-47                          !!Report OMP_TARGET_OFFOAD value
S-48     select CASE (policy)
S-49        case("MANDATORY")
S-50                         print*,"Policy:  MANDATORY-Terminate if dev. not avail."
S-51        case("DISABLED")
S-52                         print*,"Policy:  DISABLED-(if supported) Only on Host."
S-53        case("DEFAULT")
S-54                         print*,"Policy:  DEFAULT On host if device not avail."
S-55        case("NOTSET")
S-56                         print*,"         OMP_TARGET_OFFLOAD is not set."
S-57        case DEFAULT
S-58                         print*,"         OMP_TARGET_OFFLOAD has unknown value."
S-59                         print*,"         UPPER CASE VALUE=",policy
S-60     end select
S-61
S-62
S-63     on_init_dev = .FALSE.
S-64                                        !! device# out of range--not supported
S-65     !$omp target  device(device_num) map(tofrom: on_init_dev)
S-66        on_init_dev=omp_is_initial_device()
S-67     !$omp end target
S-68
S-69     if (policy=="MANDATORY" .and. OPENMP_VERSION>=201811) then
S-70        print*,"OMP ERROR: OpenMP 5.0 implementation ignored MANDATORY policy."
S-71        print*,"           Termination should have occurred at target directive."
S-72     endif
S-73
S-74     print*, "Target executed on init dev (T|F): ", on_init_dev
S-75
S-76  end program policy_test
```

Fortran

APPENDIX A

Document Revision History

A.1 Changes from 5.0.0 to 5.0.1

- Added version tags (**omp**_*x.y*) in example labels and the corresponding source codes for all examples that feature OpenMP 3.0 and later.
- Included additional examples for the 5.0 features:
  - Extension to the **defaultmap** clause (Section 4.2 on page 132)
  - Transferring noncontiguous data with the **target update** directive in Fortran (Section 4.6 on page 150)
  - **conditional** modifier for the **lastprivate** clause (Section 7.8 on page 312)
  - **task** modifier for the **reduction** clause (Section 7.9.2 on page 322)
  - Reduction on combined target constructs (Section 7.9.3 on page 327)
  - Task reduction with target constructs (Section 7.9.4 on page 332)
  - **scan** directive for returning the *prefix sum* of a reduction (Section 7.10 on page 355)
- Included additional examples for the 4.x features:
  - Dependence for undeferred tasks (Section 3.3.9 on page 107)
  - **ref**, **val**, **uval** modifiers for **linear** clause (Section 5.4 on page 229)
- Clarified the description of pointer mapping and pointer attachment in Section 4.3 on page 137.
- Clarified the description of memory model examples in Section 8.1 on page 369.

A.2 Changes from 4.5.0 to 5.0.0

- Added the following examples for the 5.0 features:
  - Extended **teams** construct for host execution (Section 1.3 on page 8)
  - **loop** and **teams loop** constructs specify loop iterations that can execute concurrently (Section 1.15 on page 38)
  - Task data affinity is indicated by **affinity** clause of **task** construct (Section 2.2 on page 52)
  - Display thread affinity with **OMP_DISPLAY_AFFINITY** environment variable or **omp_display_affinity()** API routine (Section 2.3 on page 53)
  - **taskwait** with dependences (Section 3.3.6 on page 95)
  - **mutexinoutset** task dependences (Section 3.3.7 on page 102)
  - Multidependence Iterators (in **depend** clauses) (Section 3.3.8 on page 105)
  - Combined constructs: **parallel master taskloop** and **parallel master taskloop simd** (Section 3.7 on page 118)
  - Reverse Offload through *ancestor* modifier of **device** clause. (Section 4.1.6 on page 129)
  - Pointer Mapping - behavior of mapped pointers (Section 4.3 on page 137)
  - Structure Mapping - behavior of mapped structures (Section 4.4 on page 143)
  - Array Shaping with the *shape-operator* (Section 4.6 on page 150)
  - The **declare mapper** construct (Section 4.7 on page 153)
  - Acquire and Release Semantics Synchronization: Memory ordering clauses **acquire**, **release**, and **acq_rel** were added to flush and atomic constructs (Section 6.7 on page 258)
  - **depobj** construct provides dependence objects for subsequent use in **depend** clauses (Section 6.9 on page 270)
  - **reduction** clause for **task** construct (Section 7.9.2 on page 322)
  - **reduction** clause for **taskloop** construct (Section 7.9.5 on page 337)
  - **reduction** clause for **taskloop simd** construct (Section 7.9.5 on page 337)
  - Memory Allocators for making OpenMP memory requests with traits (Section 8.2 on page 376)
  - **requires** directive specifies required features of implementation (Section 9.5 on page 395)
  - **declare variant** directive - for function variants (Section 9.6 on page 397)
  - **metadirective** directive - for directive variants (Section 9.7 on page 403)

  - **OMP_TARGET_OFFLOAD** Environment Variable - controls offload behavior (Section 9.10 on page 420)
- Included the following additional examples for the 4.x features:
  - more taskloop examples (Section 3.6 on page 114)
  - user-defined reduction (UDR) (Section 7.9.6 on page 344)

A.3 Changes from 4.0.2 to 4.5.0

- Reorganized into chapters of major topics
- Included file extensions in example labels to indicate source type
- Applied the explicit **map(tofrom)** for scalar variables in a number of examples to comply with the change of the default behavior for scalar variables from **map(tofrom)** to **firstprivate** in the 4.5 specification
- Added the following new examples:
  - **linear** clause in loop constructs (Section 1.9 on page 25)
  - **priority** clause for **task** construct (Section 3.2 on page 88)
  - **taskloop** construct (Section 3.6 on page 114)
  - *directive-name* modifier in multiple **if** clauses on a combined construct (Section 4.1.5 on page 126)
  - unstructured data mapping (Section 4.9 on page 173)
  - **link** clause for **declare target** directive (Section 4.11.5 on page 188)
  - asynchronous target execution with **nowait** clause (Section 4.13 on page 200)
  - device memory routines and device pointers (Section 4.14.4 on page 213)
  - doacross loop nest (Section 6.10 on page 274)
  - locks with hints (Section 6.11 on page 280)
  - C/C++ array reduction (Section 7.9.1 on page 314)
  - C++ reference types in data sharing clauses (Section 7.13 on page 364)

A.4 Changes from 4.0.1 to 4.0.2

- Names of examples were changed from numbers to mnemonics
- Added SIMD examples (Section 5.1 on page 216)
- Applied miscellaneous fixes in several source codes
- Added the revision history

A.5 Changes from 4.0 to 4.0.1

Added the following new examples:

- the **`proc_bind`** clause (Section 2.1 on page 46)
- the **`taskgroup`** construct (Section 3.4 on page 109)

A.6 Changes from 3.1 to 4.0

- Beginning with OpenMP 4.0, examples were placed in a separate document from the specification document.
- Version 4.0 added the following new examples:
  - task dependences (Section 3.3 on page 90)
  - **`target`** construct (Section 4.1 on page 121)
  - array sections in device constructs (Section 4.5 on page 146)
  - **`target data`** construct (Section 4.8 on page 160)
  - **`target update`** construct (Section 4.10 on page 176)
  - **`declare target`** construct (Section 4.11 on page 180)
  - **`teams`** constructs (Section 4.12 on page 191)
  - asynchronous execution of a **`target`** region using tasks (Section 4.13.1 on page 200)
  - device runtime routines (Section 4.14 on page 209)
  - Fortran ASSOCIATE construct (Section 7.14 on page 365)
  - cancellation constructs (Section 9.4 on page 390)