



OpenMP programming Part I

Shaohao Chen

High performance computing @ Louisiana State University



Outline

Part I

- Introduction to OpenMP
- OpenMP language features

Part II

- Optimization for performance
- Trouble shooting and debug
- Use OpenMP with Intel Xeon Phi

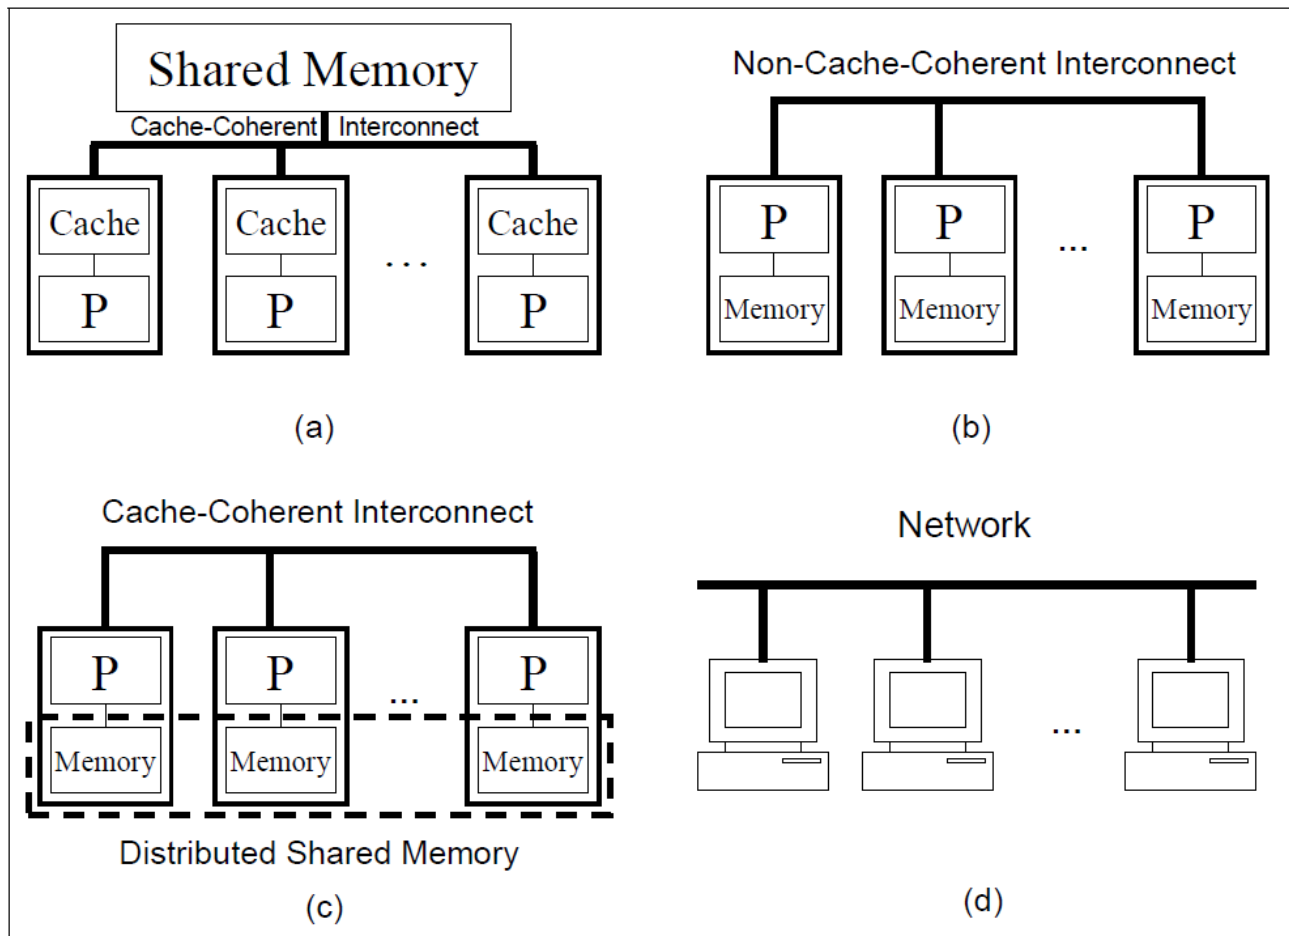


History

- ❑ The OpenMP Architecture Review Board (ARB) published its first API specifications, OpenMP 1.0 for Fortran, in October 1997. October the following year they released the C/C++ standard.
- ❑ 2000 saw version 2.0 of the Fortran specifications with version 2.0 of the C/C++ specifications being released in 2002. Version 2.5 is a combined C/C++/Fortran specification that was released in 2005.
- ❑ Version 3.0 was released in May 2008. Included in the new features in 3.0 is the concept of tasks and the task construct. Version 3.1 of the OpenMP specification was released July 9, 2011.
- ❑ Version 4.0 of the specification was released in July 2013. It adds or improves the following features: support for accelerators; atomics; error handling; thread affinity; tasking extensions; user defined reduction; SIMD support; Fortran 2003 support.



Distributed and shared memory



(a) Physically shared-memory system (e.g. one compute node with multi processors).

(b) Distributed memory system (e.g. multi compute nodes).

(c) Distributed memory system, however the distributed memories are accessible to all processors.

(d) A set of independent computers linked by a network.

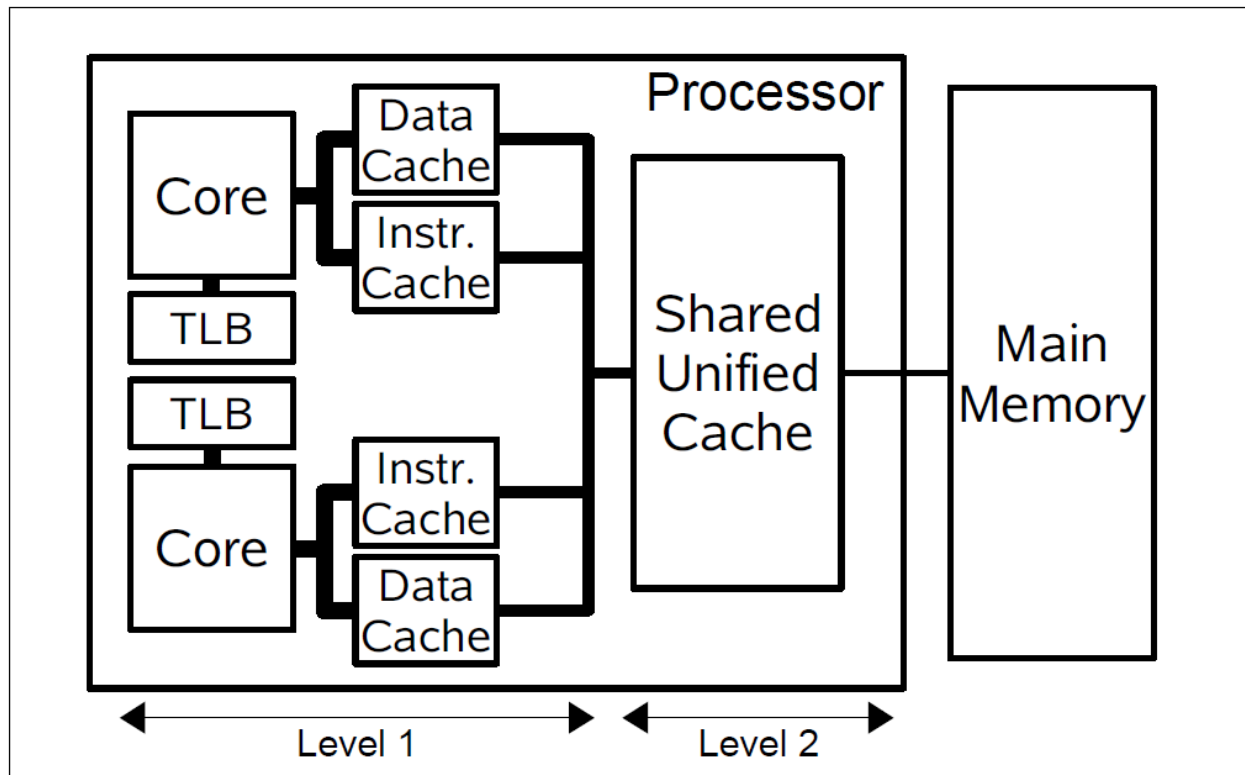
Parallel schemes:

❑ Shared-memory parallel (SMP) system:
OpenMP

❑ Distributed-memory system (e.g. SIMD):
MPI



Multicore processor with shared memory

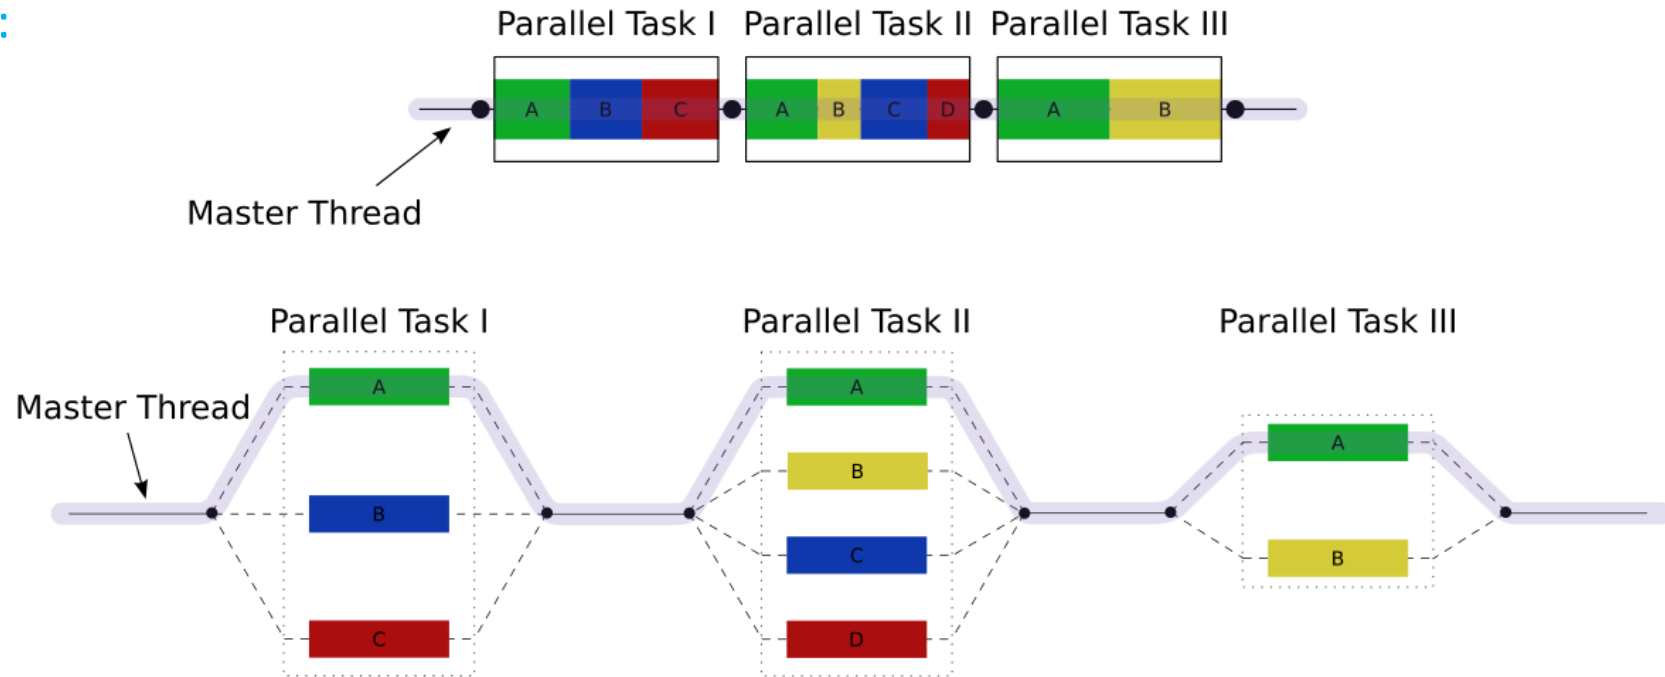


- The level-1 cache is private to the core, but the cache at the second level is shared. Both cores can use it to store and retrieve instructions, as well as data.
- Data is copied into cache from main memory.
- ❑ OpenMP can be applied to a multicore processor with shared memory.
- ❑ Works are spread to multi threads and each thread is assigned to one core.



Parallelism of OpenMP

- **Multithreading:** a master thread forks a specified number of slave threads and the system divides a task among them. The threads then run concurrently, with the runtime environment allocating threads to different processors (or cores).
- **Fork-join model:**





- Hello world in Fortran language

```
program hello
  use omp_lib
  implicit none
  integer i
  !$omp parallel private(i)
  i = omp_get_thread_num()
  if (mod(i,2).eq.1) then
    print *, 'Hello from thread', i, ', I am odd!'
  else
    print *, 'Hello from thread', i, ', I am even!'
  endif
  !$omp end parallel
end program hello
```

- #pragma omp** *directive-name* [*clause*[[,] *clause*]. . .]

- !\$omp** *directive-name* [*clause*[[*,*] *clause*]. . .]

- **directive-name**: a specific keyword, for example parallel, that defines and controls the action(s) taken.
- **clauses**: for example private, can be used to further specify the behavior.

```
> icc/icpc/ifort -openmp name.c/name.f90 -o name
> gcc/g++/gfortran -fopenmp name.c/name.f90 -o name
> pgcc/pgc++/pgf90 -mp name.c/name.f90 -o name
```

```
> export OMP_NUM_THREADS=20      # set number of threads
> ./name
> time ./name                     # run and measure the time.
```



OpenMP language features

- Parallel Construct
 - Work-Sharing Constructs
 - Loop Construct
 - Sections Construct
 - Single Construct
 - Workshare Construct (Fortran only)
 - Basic clauses
 - shared, private, lastprivate, firstprivate,
default, nowait, schedule
 - Synchronization constructs
 - Barrier Construct
 - Master Construct
 - Critical Construct
 - Atomic Construct
 - Advanced clauses:
 - reduction, if, num_thread
 - Nested parallelism
-
- **Construct** : An OpenMP executable directive and the associated statement, loop, or structured block, not including the code in any called routines.



Parallel construct

- Syntax in C/C++ programs

```
#pragma omp parallel [clause[[,] clause]. . . ]
```

```
..... code block .....
```

- Syntax in Fortran programs

```
!$omp parallel [clause[[,] clause]. . . ]
```

```
..... code block .....
```

```
!$omp end parallel
```

- Parallel construct is used to specify the computations that should be executed in parallel.
- A team of threads is created to execute the associated parallel region.
- The work of the region is replicated for every thread.
- At the end of a parallel region, there is an implied barrier that forces all threads to wait until the work inside the region has been completed.



- Clauses supported by the parallel construct

- **if**(*scalar-expression*) (C/C++)
- **if**(*scalar-logical-expression*) (Fortran)
- **num_threads**(*integer-expression*) (C/C++)
- **num_threads**(*scalar-integer-expression*) (Fortran)
- **private**(*list*)
- **firstprivate**(*list*)
- **shared**(*list*)
- **default**(**none** | **shared**) (C/C++)
- **default**(**none** | **shared** | **private**) (Fortran)
- **copyin**(*list*)
- **reduction**(*operator:list*) (C/C++)
- **reduction**(*{operator | intrinsic procedure name}:list*) (Fortran)



Work-sharing constructs

Functionality	Syntax in C/C++	Syntax in Fortran
Distribute iterations	#pragma omp for	!\$omp do
Distribute independent works	#pragma omp sections	!\$omp sections
Use only one thread	#pragma omp single	!\$omp single
Parallelize array syntax	N/A	!\$omp workshare

- Many applications can be parallelized by using just a parallel region and one or more of these constructs, possibly with clauses.



- The parallel and work-sharing (except single) constructs can be combined.
- Following is the syntax for combined parallel and work-sharing constructs,

Combine parallel construct with ...	Syntax in C/C++	Syntax in Fortran
Loop construct	#pragma omp parallel for	!\$omp parallel do
Sections construct	#pragma omp parallel sections	!\$omp parallel sections
Workshare construct	N/A	!\$omp parallel workshare



Loop construct

- The loop construct causes the iterations of the loop immediately following it to be executed in parallel.

- Syntax in C/C++ programs

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

- Syntax in Fortran programs

```
!$omp do [clause[,] clause]. . . ]  
..... do loop .....
```

[!\$omp end do]

- The terminating **!\$omp end do** directive in Fortran is optional but recommended.



- Distribute iteration in a parallel region

```
#pragma omp parallel for shared(n,a) private(i)
{
    for (i=0; i<n; i++)
        a[i]=i+1;
}    /*-- End of parallel region --*/
```

- **shared clause:** All threads can read from and write to the variable.
- **private clause:** Each thread has a local copy of the variable.
- The maximum iteration number n is shared, while the iteration number i is private.
- Each thread executes a **subset** of the total iteration space $i = 0, \dots, n - 1$
- The mapping between iterations and threads can be controlled by the schedule clause.



- Two work-sharing loops in one parallel region

```
#pragma omp parallel shared(n,a,b) private(i)
{
    #pragma omp for
    for (i=0; i<n; i++)  a[i] = i;
    // there is an implied barrier

    #pragma omp for
    for (i=0; i<n; i++)  b[i] = 2 * a[i];
}  /*-- End of parallel region --*/
```

- The distribution of iterations to threads could be different for the two loops.
- The **implied barrier** at the end of the first loop ensures that all the values of $a[i]$ are updated before they are used in the second loop.



Sections construct

- Syntax in C/C++ programs

```
#pragma omp sections [clause[[,] clause]. . . ]  
{  
  [#pragma omp section ]  
  ..... code block 1 .....  
  [#pragma omp section  
  ..... code block 2 ..... ]  
  . . .  
}
```

- Syntax in Fortran programs

```
!$omp sections [clause[[,] clause]. . . ]  
[!$omp section ]  
..... code block 1 .....  
[!$omp section  
..... code block 2 ..... ]  
.  
.  
.  
!$omp end sections
```

- The work in each section must be **independent**.
- Each section is distributed to one thread.

- Syntax in Fortran programs

!\$omp single [clause[[,] clause]. . .]

```
..... code block .....
```

!\$omp end single

- The code block following the single construct is executed by one thread only.
- The executing thread could be any thread (not necessary the master one).
- The other threads wait at a barrier until the executing thread has completed.



Workshare construct

- Workshare construct is only available for Fortran.

- Syntax in Fortran programs

```
!$omp workshare [clause[[,] clause]. . . ]
```

```
..... code block .....
```

```
!$omp end workshare
```

- Units of works within the block are executed in parallel in a manner that respects the semantics of Fortran array operations.
- For example, if the workshare directive is applied to an array assignment statement, the assignment of each element is a unit of work.



- Example of workshare construct

```
!$OMP PARALLEL SHARED(n,a,b,c)
```

```
!$OMP WORKSHARE
```

```
  b(1:n) = b(1:n) + 1
```

```
  c(1:n) = c(1:n) + 2
```

```
  a(1:n) = b(1:n) + c(1:n)
```

```
!$OMP END WORKSHARE
```

```
!$OMP END PARALLEL
```

- These array operations are parallelized.
- There is no control over the assignment of array updates to the threads.
- The OpenMP compiler must generate code such that the updates of b and c have completed before a is computed.

- ```
#pragma omp parallel for private(i) lastprivate(a)

for (i=0; i<n; i++) {
 a = i+1;
 printf("Thread %d has a value of a = %d for i = %d\n", omp_get_thread_num(),a,i);
} /*-- End of parallel for --*/

printf("After parallel for: i = %d , a = %d\n", i, a);
```

- Alternative code with shared clause

```
#pragma omp parallel for private(i) private(a) shared(a_shared)
for (i=0; i<n; i++) {
 a = i+1;
 printf("Thread %d has a value of a = %d for i = %d\n",
 omp_get_thread_num(),a,i);
 if (i == n-1) a_shared = a;
} /*-- End of parallel for --*/
```

- All behavior of the lastprivate clause can be reproduced by the shared clause, **but the lastprivate clause is more recommended**.
- A performance penalty is likely to be associated with the use of lastprivate, because the OpenMP library needs to keep track of which thread executes the last iteration.



## Firstprivate clause

- **private clause:** Preinitialized value of variables are not passed to the parallel region.
- **firstprivate clause:** Each thread has a preinitialized copy of the variable. This variable is still private, so threads can update it individually.
- Firstprivate clause is available for parallel, loop, sections and single constructs.

```
int i, vtest=10, n=20;
#pragma omp parallel for private(i) firstprivate(vtest) shared(n)
for(i=0; i<n; i++) {
 printf("thread %d: initial value = %d\n", omp_get_thread_num(), vtest);
 vtest=i;
}
printf("value after loop = %d\n", vtest);
```



- Syntax in C programs

- Syntax in Fortran programs

- **An example:** declares all variables to be shared, with the some exceptions.

- If default(`none`) is specified, the programmer is forced to specify a data-sharing attribute for each variable in the construct.



## Nowait clause

- If the nowait clause is added to a construct, **the implicit barrier at the end of the associated construct will be suppressed**. When a thread is finished with the work associated with the parallelized for loop, it continues and no longer waits for the other threads to finish.
- Note, however, that the barrier at the end of a parallel region cannot be suppressed.
- An example for C program
- An example for Fortran program

```
#pragma omp for nowait
for (i=0; i<n; i++)
{
.....
} // no barrier here
```

```
!$OMP DO
.....
!$OMP END DO NOWAIT ! no barrier here
```



## Schedule clause

- Specifies how iterations of the loop are assigned to the threads in the team.
- Supported on the loop construct only.
- The iteration space is divided into chunks. Chunk represents the granularity of workload distribution, a contiguous nonempty subset of the iteration space.

- Syntax

`schedule(kind [,chunk_size] )`

- The **static schedule** works best for regular workloads and is the default on many OpenMP compilers.
- The **dynamic and guided schedules** are useful for handling poorly balanced and unpredictable workloads.
- There is a performance penalty for using dynamic and guided schedules.



- Schedule kind

| kind    | description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| static  | The chunks are assigned to the threads statically in a round-robin manner, in the order of the thread number. If <i>chunk_size</i> is not specified, the chunk size is approximately equal to the total number of iteration divided by the number of threads.                                                                                                                                                                                                                                                                                                                  |
| dynamic | The chunks are assigned to threads as the threads request them. The last chunk may have fewer iterations than chunk size. If <i>chunk_size</i> is not specified, it defaults to 1.                                                                                                                                                                                                                                                                                                                                                                                             |
| guided  | The chunks are assigned to threads as the threads request them. For a <i>chunk_size</i> of 1, the size of each chunk is proportional to the number of unassigned iterations, divided by the number of threads, decreasing to 1. For a <i>chunk_size</i> of “k” ( $k > 1$ ), the size of each chunk is determined in the same way, with the restriction that the chunks do not contain fewer than k iterations (with a possible exception for the last chunk to be assigned, which may have fewer than k iterations). When no <i>chunk_size</i> is specified, it defaults to 1. |
| runtime | The schedule and (optional) chunk size are set through the OMP_SCHEDULE environment variable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



- Example of schedule clause:

The workload in the inner loop depends on the value of the outer loop iteration variable  $i$ . Therefore, **the workload is not balanced**, and the static schedule is probably not the best choice. Dynamic or guided schedules are required.

```
#pragma omp parallel for default(none) schedule(runtime) private(i,j) shared(n)
for (i=0; i<n; i++)
{
 printf("Iteration %d executed by thread %d\n", i, omp_get_thread_num());
 for (j=0; j<i; j++)
 system("sleep 1");
}
```

# Barrier construct

- A barrier is a point in the execution of a program where threads wait for each other: no thread in the team of threads it applies to may proceed beyond a barrier until all threads in the team have reached that point.

- Syntax in C/C++ programs

## #pragma omp barrier

- Syntax in Fortran programs

## !\$omp barrier

## Two important restrictions apply to the barrier construct:

- Each barrier must be encountered by all threads in a team, or by none at all.
- The sequence of work-sharing regions and barrier regions encountered must be the same for every thread in the team.



- Example of barrier construct:

A thread waits at the barrier until the last thread in the team arrives.

```
#pragma omp parallel private(TID)
{
 TID = omp_get_thread_num();
 if (TID < omp_get_num_threads()/2) system("sleep 3");
 bt1 = time(NULL);
 printf("Thread %d before barrier at %s \n", omp_get_thread_num(), ctime(&t1));
 #pragma omp barrier
 t2 = time(NULL);
 printf("Thread %d after barrier at %s \n", omp_get_thread_num(), ctime(&t2));
} /*-- End of parallel region --*/
```





# Master construct

- The master construct defines a block of code that is guaranteed to be executed by the master thread only.
- **It does not have an implied barrier on entry or exit.** In the cases where a barrier is not required, the master construct may be preferable compared to the single construct.

- Syntax in C/C++ programs

```
#pragma omp master
```

```
..... code block
```

- Syntax in Fortran programs

```
!$omp master
```

```
..... code block
```

```
!$omp end master
```

- The master construct is often used (in combination with barrier construct) to initialize data.



## Critical construct

- The critical construct provides a means to **ensure that multiple threads do not attempt to update the same shared data simultaneously.**
- When a thread encounters a critical construct, it waits until no other thread is executing a critical region with the same name.

- Syntax in C/C++ programs

```
#pragma omp critical [(name)]
..... code block
```

- Syntax in Fortran programs

```
!$omp critical [(name)]
..... code block
```

```
!$omp end critical [(name)]
```

- The code block is executed by all threads, but only one at a time executes the block.



- Example 1 of critical construct: Avoiding garbled output

A critical region helps to avoid intermingled output when multiple threads print from within a parallel region.

```
#pragma omp parallel private(TID)
{
 TID = omp_get_thread_num();
 #pragma omp critical (print_tid)
 {
 printf("Thread %d : Hello, ",TID);
 printf("world!\n");
 }
} /*-- End of parallel region --*/
```



## Race condition

- Race conditions arise when the result depends on the sequence or timing of processes or threads, for example, **when multithreads read or write the same shared data simultaneously**.
- Example:** two threads each want to increment the value of a shared integer variable by one.

### Correct sequence

| Thread 1       | Thread 2       |   | value |
|----------------|----------------|---|-------|
|                |                |   | 0     |
| read value     |                | ← | 0     |
| Increase value |                |   | 0     |
| write back     |                | → | 1     |
|                | read value     | ← | 1     |
|                | increase value |   | 1     |
|                | write back     | → | 2     |

### Incorrect sequence

| Thread 1       | Thread 2       |   | value |
|----------------|----------------|---|-------|
|                |                |   | 0     |
| read value     |                | ← | 0     |
|                | read value     | ← | 0     |
| increase value |                |   | 0     |
|                | increase value |   | 0     |
| write back     |                | → | 1     |
|                | write back     | → | 1     |



- Example of data racing: sums up elements of a vector

Multithreads can read and write the shared data sum simultaneously.

**A data race condition arises!**

If a thread reads sum before sum is updated by another thread, the final result of sum is wrong!

```
sum = 0;
#pragma omp parallel for shared(sum,a,n) private(i)
for (i=0; i<n; i++)
{
 sum = sum + a[i];
} /*-- End of parallel for --*/
printf("Value of sum after parallel region: %f\n",sum);
```



- A partially parallel scheme to avoid data racing

Step 1: Calculate local sums in parallel

Thread 1

$$\begin{array}{c} \boxed{a_0} \\ + \\ \boxed{a_1} \\ + \\ \vdots \\ + \\ \boxed{a_{m-1}} \\ \parallel \\ \boxed{LS_1} \end{array}$$

Thread 2

$$\begin{array}{c} \boxed{a_m} \\ + \\ \boxed{a_{m+1}} \\ + \\ \vdots \\ + \\ \boxed{a_{2m-1}} \\ \parallel \\ \boxed{LS_2} \end{array}$$

.....

.....

.....

.....

Thread m

$$\begin{array}{c} \boxed{a_{n-m-1}} \\ + \\ \boxed{a_{n-m}} \\ + \\ \vdots \\ + \\ \boxed{a_n} \\ \parallel \\ \boxed{LS_m} \end{array}$$

m: number of threads

n: array length

LS: local sum



## Step 2: Update total sum sequentially

| Thread 1       | Thread 2       | ..... | Thread m       |
|----------------|----------------|-------|----------------|
| Read initial S |                |       |                |
| $S = S + LS_1$ |                |       |                |
| Write S        |                |       |                |
|                | Read S         |       |                |
|                | $S = S + LS_2$ |       |                |
|                | Write S        |       |                |
|                |                | ..... |                |
|                |                |       | Read S         |
|                |                |       | $S = S + LS_m$ |
|                |                |       | Write S        |

m: number of threads

LS: local sum

S: total sum



- Example 2 of critical construct: sums up the elements of a vector

The critical region is needed to avoid a data race condition when updating variable sum.

```
sum = 0;
#pragma omp parallel shared(n,a,sum) private(sumLocal)
{
 sumLocal = 0;
 #pragma omp for
 for (i=0; i<n; i++) sumLocal += a[i];
 #pragma omp critical (update_sum)
 {
 sum += sumLocal;
 printf("TID=%d: sumLocal=%d sum = %d\n", omp_get_thread_num(), sumLocal, sum);
 }
} /*-- End of parallel region --*/
printf("Value of sum after parallel region: %d\n",sum);
```





- Example 3 of critical construct: To determine minimum value

The critical region is needed to avoid a data race condition when comparing the value of the private variable LScale with the shared variable Scale and when updating it and ssq. The execution order does not matter in the case.

```
#pragma omp parallel private(ix, LScale, lssq, Temp) shared(Scale, ssq)
{
 #pragma omp for
 for(ix = 1, ix<N, ix++) LScale = array[ix];
 #pragma omp critical
 {
 if(Scale < LScale){
 ssq = (Scale/LScale) *ssq + lssq;
 Scale = LScale; }
 else { ssq = ssq + (LScale / Scale) * lssq; }
 } /* End of critical region --*/
} /*-- End of parallel region --*/
```



# Atomic construct

- The atomic construct also enables multiple threads to update shared data without interference.
- It is applied only to the (single) assignment statement that immediately follows it.
- If a thread is atomically updating a value, then no other thread may do so simultaneously.

## C/C++ programs

- Syntax

```
#pragma omp atomic
..... a single statement
```

- Supported operators

```
+, *, -, /, &, ^, |, <<, >>.
```

## Fortran programs

```
!$omp atomic
..... a single statement
```

```
!$omp end atomic
```

```
+, *, -, /, .AND., .OR., .EQV., .NEQV..
```



- Example 1 of atomic construct: sums up the elements of a vector

The atomic construct ensures that no updates are lost when multiple threads update the variable sum. Atomic construct can be an alternative to the critical construct in this case.

```
sum = 0;
#pragma omp parallel shared(n,a,sum) private(sumLocal)
{
 sumLocal = 0;
 #pragma omp for
 for (i=0; i<n; i++) sumLocal += a[i];
 #pragma omp atomic
 sum += sumLocal;
} /*-- End of parallel region --*/
printf("Value of sum after parallel region: %d\n",sum);
```



- Example 2 of atomic construct: sums up the values of functions

The atomic construct does not prevent multiple threads from **executing the function *bigfunc* parallelly**.

It is only the update to the memory location of the variable sum that will occur atomically.

```
sum = 0;
#pragma omp parallel for shared(n,a,sum) private(i)
for (i=0; i<n; i++)
{
 #pragma omp atomic
 sum = sum + bigfunc();
} /*-- End of parallel for --*/
printf("Value of sum after parallel region: %d\n",sum);
```

- An OpenMP compiler will generate **a roughly equivalent machine code** for the two cases: using critical construct and using reduction clause, meaning that their performance is almost the same.
- The result sum will be shared and it is not necessary to specify it explicitly as “shared”.
- The order in which thread-specific values are combined is unspecified. Therefore, where floating-point data are concerned, **there may be numerical differences** between the results of a sequential and parallel run, or even of two parallel runs using the same number of threads.



- Operators and statements supported by the reduction clause

|                                    | C/C++                                                                                                                                                               | Fortran                                                                                                                                                                                         |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Typical statements                 | $x = x \text{ op } \text{expr}$<br>$x \text{ binop } = \text{expr}$<br>$x = \text{expr} \text{ op } x$ (except for subtraction)<br>$x++$<br>$++x$<br>$x--$<br>$--x$ | $x = x \text{ op } \text{expr}$<br>$x = \text{expr} \text{ op } x$ (except for subtraction)<br>$x = \text{intrinsic } (x, \text{expr\_list})$<br>$x = \text{intrinsic } (\text{expr\_list}, x)$ |
| <i>op</i> could be                 | +, *, -, &, ^,  , &&, or                                                                                                                                            | +, *, -, .and., .or., .eqv., or .neqv.                                                                                                                                                          |
| <i>binop</i> could be              | +, *, -, &, ^, or                                                                                                                                                   | N/A                                                                                                                                                                                             |
| <i>Intrinsic</i> function could be | N/A                                                                                                                                                                 | max, min, iand, ior, ieor                                                                                                                                                                       |





## Num\_threads clause

- The `num_threads` clause is supported on the `parallel` construct only and can be used to specify how many threads should be in the team executing the parallel region

```
omp_set_num_threads(4);
#pragma omp parallel if (n > 5) num_threads(n) default(none) shared(n)
{
 #pragma omp single
 {
 printf("Number of threads in parallel region: %d\n", omp_get_num_threads());
 }

 printf("Print statement executed by thread %d\n", omp_get_thread_num());
} /*-- End of parallel region --*/
```





## Nested parallelism

- If a thread in a team executing a parallel region encounters another parallel construct, it creates a new team and becomes the master of that new team.

```
#pragma omp parallel private(TID)
{
 TID = omp_get_thread_num();
 #pragma omp parallel num_threads(2) firstprivate(TID)
 {
 printf("Outer thread number: %d. Inner thread number: %d.\n", TID, omp_get_thread_num());
 } /*-- End of inner parallel region --*/
} /*-- End of outer parallel region --*/
```

- The function `omp_get_thread_num()` returns the thread number of the current parallel region.
- The thread number of the first level can be passed on to the second level by `firstprivate` clause.



# OpenMP functions

- Enable the usage of OpenMP functions:

C/C++ program: include omp.h .

Fortran program: include omp\_lib.h or use omp\_lib module.

- List of OpenMP functions:

`omp_set_num_threads(integer)` : set the number of threads

`omp_get_num_threads()`: returns the number of threads

`omp_get_thread_num()`: returns the number of the calling thread.

**omp\_set\_dynamic(integer|logical):** dynamically adjust the number of threads

`omp_get_num_procs()`: returns the total number of available processors when it is called.

`omp_in_parallel()`: returns true if it is called within an active parallel region. Otherwise, it returns false.

# OpenMP runtime variables

**OMP\_NUM\_THREADS** : the number of threads (=integer)

**OMP\_SCHEDULE** : the schedule type (=kind,chunk . Kind could be static, dynamic or guided)

**OMP\_DYNAMIC**: dynamically adjust the number of threads (=true | =false).

**KMP\_AFFINITY** : for intel compiler, to bind OpenMP threads to physical processing units.  
(=compact | =scatter | =balanced).

Example usage: `export KMP_AFFINITY=compact,granularity=fine,verbose.`

