

WWW.TACC.UTEXAS.EDU



XSEDE

Extreme Science and Engineering Discovery Environment

# Introduction to OpenMP Programming

RIKEN R-CCS

IHPCSS July 7-12, 2019 **PRESENTED BY:** 

Kent Milfeld / Lars Koesterke

milfeld@tacc.utexas.edu

lars@tacc.utexas.edu



Lecture and Lab slides available at: tinyurl.com/tacc-2019-ihpcss tinyurl.com/tacc-2019-ihpcss

#### Outline

#### **OpenMP Directive Review – The Concepts**

- Execution Model: Parallel Fork-Join
- Directive Syntax
- Memory Model
- Using OpenMP Directives
- Worksharing
- Directive Scope
- Clauses
- Thread Scheduling
- Data Model / Data Environment
- Synchronization
- API functions / Environment Variables

THE UNIVERSITY OF TEXAS AT AUSTIN

# OpenMP?

- OpenMP stands for **Open Multi-Processing**
- Three primary components:
  - Directives to the Compiler within code (pragmas in C/C++, comments in Fortran)
  - Runtime Library Routines
  - Environment Variables
- <u>http://www.openmp.org/</u>
  - OpenMP 5.0 Complete Specifications (Nov. 2018)
  - OpenMP 4.5 Complete Specifications (Nov. 2015)
  - OpenMP Examples 4.5 (Nov. 2015)
  - OpenMP Summary Card, Reference Guide (C/C++/Fortran)

3

THE UNIVERSITY OF TEXAS AT AUSTIN

#### **OpenMP** Constructs



## **Execution Model**

- Program begins as a single process: master thread executes serially.
- At a parallel directive (construct) a team of threads is forked.

tinyurl.com/tacc-2019-ihpcss

• At end of parallel region, (team of) threads synchronize and terminate (join) but master continues.



**TEXAS ADVANCED COMPUTING CENTER** 

## **OpenMP** Directive Scope

• OpenMP directives are are formed with comments/pragmas in the source:

F90 :!\$omp ...C/C++ :#pragma omp ...



THE UNIVERSITY OF TEXAS AT AUSTIN

6

#### **OpenMP** Syntax

- Compiler directive syntax:
   #pragma omp construct [clause [[,]clause]...]
   C
   !\$omp construct [clause [[,]clause]...]
   F90
- Example

Fortran print*,"serial"	printf("serial\n");	C/C++
<pre>!\$omp parallel num_threads(4)</pre>	<pre>#pragma omp parallel num_threads(4) {</pre>	
!\$omp end parallel	}	
print*,"serial"	printf("serial\n");	

THE UNIVERSITY OF TEXAS AT AUSTIN

tinyurl.com/tacc-2019-ihpcss

#### **TEXAS ADVANCED COMPUTING CENTER**

# Memory Model and Synchronization

- Every thread has access to (shared) memory. When multiple threads are created all threads share the same address space.
- Simultaneous updates to shared memory can create a *race condition*-- Causing results change with different thread scheduling.
- Use mutual exclusion type directives to avoid race conditions but don't use too many because this will serialize performance.
- Barriers are used to synchronize threads.
- A flush exists at barriers, to make local changes to data (writes) seen by all threads.

8

THE UNIVERSITY OF TEXAS AT AUSTIN

## C/C++

# **Parallel Regions**

```
1 #pragma omp parallel
2 {
3     //code block
4     work(...);
5 }
```

- Line 1 Team of threads formed at beginning of parallel region
- Lines 3-4 Each thread executes code block

No branching into or out of a parallel region

Line 5 All threads synchronize at end of parallel region (implied barrier)

How to make individualized work: use the thread number to divide work among threads.



THE UNIVERSITY OF TEXAS AT AUSTIN



**F90** 

# **Parallel Regions**

- 1 !\$OMP PARALLEL
- 2 ! code block
- 3 call work(...)
- 4 !\$OMP END PARALLEL
- Line 1 Team of threads formed at beginning of parallel region.
- Lines 2-3 Each thread executes code block and subroutine calls. No branching into or out of a parallel region.
- Line 4 All threads synchronize at end of parallel region (implied barrier).

How to make individualized work: use the thread number to divide work among threads.

10



THE UNIVERSITY OF TEXAS AT AUSTIN

TEXAS ADVANCED COMPUTING CENTER

#### C/C++ Parallel Region & Number of Threads

• For example, to create a 4-thread Parallel region.

```
gcc -fopenmp code.c
export OMP_NUM_THREADS=4
a.out
```



- Each thread executes the code within the parallel region
- Thread numbers range from 0 to Nthreads-1
- Each thread calls foo(ID,A) with a different **ID** {=0,1,2,or3}

THE UNIVERSITY OF TEXAS AT AUSTIN

#### **F90**

# Parallel Region & Number of Threads

• For example, to create a 4-thread Parallel region.



- Each thread executes the code within the parallel region
- Thread numbers range from 0 to Nthreads-1
- Each thread calls foo(ID, A) with a different **ID** {= 0, 1, 2 or 3}

12

THE UNIVERSITY OF TEXAS AT AUSTIN

# Parallel Region & Worksharing

Use OpenMP directives to specify Worksharing inside a Parallel Region



"!\$omp" and "#pragma omp" not shown here.

Work-sharing Directive assigns threads to unit(s) of work. There is an implied barrier at the end of a worksharing construct!





THE UNIVERSITY OF TEXAS AT AUSTIN

**TEXAS ADVANCED COMPUTING CENTER** 

### C/C++

# Worksharing: Loop

```
1 #pragma omp parallel
2 {
3     #pragma omp for
4     for (i=0; i<N; i++)
5     {
6         a[i] = b[i] + c[i];
7     }
8 }</pre>
```

Line 1 Team of threads formed (parallel region).

Line 3-7 Loop iterations are split among threads. implied <u>barrier at }</u>

Each loop iteration must be independent of other iterations.

```
Or
  #pragma omp parallel
  #pragma omp for
  for (i=0; i<N; i++)
        a[i] = b[i] + c[i];</pre>
```

```
Or
    #pragma omp parallel for
    for (i=0; i<N; i++)
        a[i] = b[i] + c[i];</pre>
```

THE UNIVERSITY OF TEXAS AT AUSTIN

14

F9()

# Worksharing: Loop



```
Or
 !$OMP PARALLEL DO
 do i=1, N
    a(i) = b(i) + c(i)
 enddo
```

- Line 1 Team of threads formed (parallel region).
- Line 3-4 Loop iterations are split among threads by DO construct.
- Line 5 !\$OMP END DO is optional after the enddo.
- Implied barrier at enddo. Line 5

Each loop iteration must be independent of other iterations.



THE UNIVERSITY OF TEXAS AT AUSTIN

# Worksharing: Sections

```
1 #pragma omp parallel sections
2 {
3    #pragma omp section
4    {
5        work_1();
6    }
7    #pragma omp section
8       { work_2(); }
9 }
```

Line 1 Team of threads formed (parallel region).

- Line 3-8 Only **One** thread works on a section.
- Line 9 End of parallel sections with an implied barrier.

Scales only to the number of sections.



#### **F90**

# Worksharing: Sections

```
1 !$omp parallel sections
2 !$omp section
4 call work_1()
6 !$omp section
7 !$omp section
8 call work_2()
9
10 !$omp end parallel sections
```

Line 1 Team of threads formed (parallel region).Line 3-9 Only **One** thread works on a section.Line 10 End of parallel sections with an implied barri

Line 10 End of parallel sections with an implied barrier.

Scales only to the number of sections.



17

# Question



18

THE UNIVERSITY OF TEXAS AT AUSTIN

## **OpenMP** Parallel + Worksharing

: Work blocks are executed by all threads. Replicated **Work-Sharing** : Work is divided among threads.



tinyurl.com/tacc-2019-ihpcss TEXAS ADVANCED COMPUTING CENTER

19

## **OpenMP** Parallel + Worksharing (sections)

**Replicated** : Work blocks are executed by all threads. **Work-Sharing** : Work is divided among threads.



#### **OpenMP** Parallel + Worksharing



21

## **OpenMP** Directive Scope

- construct the lexical extent of executable directive
- all code encountered in a construct (construct + routines) region



22

# **OpenMP Clauses**

#### Clauses control the behavior of an OpenMP directive:

Control Schedule for *for/do* worksharing Data Scoping Initialization Parallelize a region or not Number of threads to use

tinyurl.com/tacc-2019-ihpcss

Clause schedule() private(), shared(), default() firstprivate(), copyin() if() num\_threads()

TACC

THE UNIVERSITY OF TEXAS AT AUSTIN

23

**TEXAS ADVANCED COMPUTING CENTER** 

#### Schedule Clause for loop worksharing

E.G. **!\$omp do schedule(static)** Fort. do i=1,N A(i)=B(i)+C(i) enddo

Туре	Description		
schedule (static)	Each CPU receives one set of contiguous iterations		
schedule (static, C)	Iterations are divided round-robin fashion in chunks of size C iterations		
schedule (dynamic, C)	Iterations handed out in chunks of size C as CPUs become available		
schedule (guided, C)	Each of the iterations are handed out in pieces of exponentially decreasing size, with C minimum number of iterations to dispatch each time		
schedule (runtime)	Schedule and chunk size taken from the OMP_SCHEDULE environment variable		
schedule (auto)	Let the runtime decide		
	Before Execution: export OMP_SCHEDULE="static,100"		
	E.G. <b>#pragma omp for schedule(runtime)</b> C for(i=0;i <n;i++) 24<="" a[i]="B[i]+C[i];" th=""></n;i++)>		
FACC	tinyurl.com/tacc-2019-ihpcss TEXAS ADVANCED COMPUTING CENTER		

Example - schedule(static, 16), threads = 4



THE UNIVERSITY OF TEXAS AT AUSTIN

25

#### **Comparison of Scheduling Options**

type	compute overhead	distribution	chunk size	static or dynami c
static	lowest	partitioned	N/P	static
static, c	low	round-robin	С	static
dynamic [, c]	medium	simple dynamic	С	dynami c
guided [, c]	Less than medium	guided	decreasing unassigned/P	dynami c
runtime	varies	runtime	varies	varies

[, C] = optional default = 1

26

THE UNIVERSITY OF TEXAS AT AUSTIN

## Data Model – for parallel region



- Threads Execute on Cores/HW-threads
- In a parallel region, team threads are assigned (tied) to implicit tasks to do work. Think of tasks and threads as being synonymous.
- Tasks by "default" share memory declared in scope before a parallel region.
- Data: shared or private
  - Shared data: accessible by all tasks
  - Private data: only accessible by the owner task

THE UNIVERSITY OF TEXAS AT AUSTIN

27

# **OpenMP Data Environment**

• The following clauses control the data-sharing attributes of variables:

shared, private, reduction, firstprivate, lastprivate

Default variable scope (in parallel region):

- 1. Variables declared in main/program (C/F90) are shared by default
- 2. Global variables are shared by default
- 3. Automatic variables within function | subroutine, called from within a parallel region, are private (reside on a stack private to each thread)
- 4. Loop index of worksharing loops are private.
- 5. Default scoping rule can be changed with **default** clause

THE UNIVERSITY OF TEXAS AT AUSTIN

28

# Private & Shared Data

**shared** - Variable is shared (seen) by all threads

**private** - Each thread has a private instance (copy) of the variable Defaults: The for-loop index is private, all other variables are shared

```
#pragma omp parallel for shared(a,b,c,n)
                                                   private(i)
                                                           OK to be explicit;
   for (i=0; i<n; i++) {</pre>
                                                           but not necessary.
       a[i] = b[i] + c[i];
   }
```

All threads have access to the same storage areas for a, b, c, and n, but each loop has its own private copy of the loop index, i



THE UNIVERSITY OF TEXAS AT AUSTIN

29

# Private & Shared Data

- shared Variable is shared (seen) by all threads
- private Each thread has a private instance (copy) of the variable

Defaults: The for-loop index is private, all other variables are shared

```
OK to be explicit;
!$omp parallel do shared(a,b,c,n) private(i) but not necessary.
do i = 1,n
        a(i) = b(i) + c(i)
        enddo
```

All threads have access to the same storage areas for a, b, c, and n, but each loop has its own private copy of the loop index, i



# C/C++

# Private Data Example

• In the following loop, each thread needs its own private copy of temp

If temp were shared, the result would be unpredictable since each thread would be writing/reading to/from the same memory location

```
#pragma omp parallel for shared(a,b,c,n) private(temp,i)
for (i=0; i<n; i++) {
   temp = a[i] / b[i];
   c[i] = temp + cos(temp);
}</pre>
```

- A lastprivate(temp) clause will copy the last loop(stack) value of temp to the (global) temp storage when the parallel *for* is complete.
- A firstprivate(temp) would copy the global temp value to each stack's temp.

31

THE UNIVERSITY OF TEXAS AT AUSTIN



# Private Data Example

• In the following loop, each thread needs its own private copy of temp

If temp were shared, the result would be unpredictable since each thread would be writing/reading to/from the same memory location

!\$omp parallel for shared(a,b,c,n) private(temp,i)
do i = 1,n
 temp = a(i) / b(i)
 c(i) = temp + cos(temp)
endo

- A lastprivate(temp) clause will copy the last loop(stack) value of temp to the (global) temp storage when the parallel *do* is complete.
- A firstprivate(temp) would copy the global temp value to each stack's temp.

32

THE UNIVERSITY OF TEXAS AT AUSTIN

# C/C++

#### Reduction

- Reduction: Operation that combines multiple elements to form a single result
- A variable that accumulates the result is called a reduction variable
- In parallel loops reduction operators and variables must be declared

```
float asum=0.0, aprod=1.0;
#pragma omp parallel for reduction(+:asum) reduction(*:aprod)
for (i=0; i<n; i++) {
    asum = asum + a[i];
    aprod = aprod * a[i];
}</pre>
```

Each thread has a private **asum** and **aprod**, initialized to the operator's identity

• After the loop execution, the master thread collects the private values of each thread and finishes the (global) reduction

tinyurl.com/tacc-2019-ihpcss TEXAS ADVANCED COMPUTING CENTER

THE UNIVERSITY OF TEXAS AT AUSTIN

#### **F90**

#### Reduction

- Reduction: Operation that combines multiple elements to form a single result
- A variable that accumulates the result is called a reduction variable
- In parallel loops reduction operators and variables must be declared

```
real asum=0.0, aprod=1.0
!$omp parallel do reduction(+:asum) reduction(*:aprod)
    do i = 1,n
        asum = asum + a(i)
        aprod = aprod * a(i)
    enddo
    print*, asum, aprod
```

Each thread has a private asum and aprod, initialized to the operator's identity

• After the loop execution, the master thread collects the private values of each thread and finishes the (global) reduction

THE UNIVERSITY OF TEXAS AT AUSTIN

# Synchronization

- Synchronization is used to impose order constraints and to protect access to shared data
- High-Level Synchronization
  - critical
  - atomic
  - barrier
  - ordered (not explored here)
- Low-Level Synchronization
  - locks (no explored here)

THE UNIVERSITY OF TEXAS AT AUSTIN

35

# C/C++

#### Synchronization: Critical/Atomic Directives

- When each thread must execute a section of code serially the region must be marked with critical/end critical directives
- Use the **#pragma omp atomic** directive for simple cases: can use hardware support



#### Synchronization: Critical/Atomic Directives

- When each thread must execute a section of code serially the region must be marked with critical/end critical directives
- Use the **!\$ omp atomic** directive for simple cases: can use hardware support



## Single Construct

• Single: Any single thread executes the construct. Since Single is worksharing there is an implied barrier.



THE UNIVERSITY OF TEXAS AT AUSTIN

tinyurl.com/tacc-2019-ihpcss TEXAS ADVANCED COMPUTING CENTER

38

#### Master Construct

• Master: Only the master executes the construct. There is no implied barrier for this construct.



THE UNIVERSITY OF TEXAS AT AUSTIN

# C/C++

# Synchronization: Barrier

• Barrier: Each thread waits until all threads arrive and a <u>flush occurs</u>

```
#pragma omp parallel shared (A, B, C) private(id)
   id=omp get thread num();
   A[id] = big calc1(id);
   #pragma omp barrier
   #pragma omp for
   for (i=0; i<N; i++) {</pre>
          C[i]=big calc3(i,A);
                Implicit barrier
   #pragma omp for nowait
   for(i=0;i<N;i++) {</pre>
           B[i]=big calc2(C, i);
                        No implicit barrier due to nowait
   } <
   A[id] = big calc4(id);

    Implicit barrier
```

40

THE UNIVERSITY OF TEXAS AT AUSTIN

# Synchronization: Barrier

F90

• Barrier: Each thread waits until all threads arrive and flush occurs

```
!$omp parallel shared (A, B, C) private(id)
  id=omp get thread num()
  A(id) = big calc1(id)
   !$omp barrier
   !$omp do
    do i = 1, N; C(i) = big_calc3(i, A); enddo
   !$omp end do < Implicit barrier
   !$omp do
    do i = 1, N; B(i) = big calc2(C, i); enddo
   !$omp end do nowait <------ No implicit barrier due to nowait
  A(id) = big calc4(id);
!$omp end parallel < Implicit barrier
```

41

THE UNIVERSITY OF TEXAS AT AUSTIN

# C/C++

#### NOWAIT

- When a work-sharing region is exited, a barrier is implied - all threads must reach the barrier before any can proceed.
- By using the NOWAIT clause at the end of each loop inside the parallel region, an unnecessary synchronization of threads can be avoided.

```
#pragma omp parallel
  #pragma omp for nowait
          for (i=0; i<n; i++)</pre>
            {work(i);}
  #pragma omp for schedule(guided,k)
          for (i=0; i<m; i++)</pre>
            {x[i]=y[i]+z[i];}
```

THE UNIVERSITY OF TEXAS AT AUSTIN

42



#### **F90**

#### NOWAIT

- When a work-sharing region is exited, a barrier is implied - all threads must reach the barrier before any can proceed.
- By using the NOWAIT clause at the end of each loop inside the parallel region, an unnecessary synchronization of threads can be avoided.

```
!SOMP PARALLEL
   !$OMP DO
      do i=1,n
         work(i)
      enddo
   !$OMP END DO NOWAIT
   !$OMP DO schedule(guided,k)
      do i=1,m
         x(i) = y(i) + z(i)
      enddo
   !$OMP END DO
$OMP END PARALLEL
```

tinyurl.com/tacc-2019-ihpcss TEXAS ADVANCED COMPUTING CENTER

THE UNIVERSITY OF TEXAS AT AUSTIN

43

## **Runtime Library Routines**

function	description
omp_get_num_threads()	Number of threads in team, N
omp_get_thread_num()	Thread ID {0 -> N-1}
omp_get_num_procs()	Number of machine CPUs
omp_in_parallel()	True if in parallel region & multiple thread executing
omp_set_num_threads(#)	Set the number of threads for subsequent parallel regions



THE UNIVERSITY OF TEXAS AT AUSTIN

44

# **Environment Variables**

variable	description
OMP_NUM_THREADS=integer	Set to default no. of threads to use
OMP_SCHEDULE="schedule-type[, chunk_size]"	Sets "runtime" in loop schedule clause: "omp for/do schedule(runtime)"
OMP_DISPLAY_ENV=anyvalue	Prints runtime environment at beginning of code execution.

٦

[...] = optional 45

THE UNIVERSITY OF TEXAS AT AUSTIN



#### **OpenMP Wallclock Timers**

real\*8 :: omp\_get\_wtime, omp\_get\_wtick() (Fortran)
double omp\_get\_wtime(), omp\_get\_wtick(); (C)

double t0, t1, dt, res; .... t0 = omp\_get\_wtime(); <work> t1 = omp\_get\_wtime(); dt = t1 - t0; res = 1.0/omp\_get\_wtick(); printf("Elapsed time = %lf\n",dt); printf("clock resolution = %lf\n",res);

THE UNIVERSITY OF TEXAS AT AUSTIN

46

#### Progression

Concept	What to learn	Level	Optimize	
Setup	How to compile OMP_NUM_THREADS	Basic	1 thread per 'core'	
Parallel region	Forking/joining threads	Easy	Minimize number of fork/join	
Work-sharing/replicated work	What do the threads do? 'omp do/for'	Work-sharing: easy Replicated: medium	Optimize scheduling Remove implicit barriers	
Avoiding race conditions		Will take effort!		
- Private variables	Why/how to shelter data	medium		
- reduction	Condensing a result from pieces	medium		
- Critical/atomic	All threads, but one thread at a time	harder	Do not serialize everything	
- Single/master	One thread, and only one thread	harder	See 'min. fork/join'	
Advanced				
- Hybrid	MPI + OpenMP	medium	Interplay MPI/OpenMP	
- Thread/memory pinning	Affinity	medium	Utilize all cores	
- SIMD	Vectorization with OpenMP	hard	Utilize vector lanes	
- Tasking	Irregular problems	hard		

THE UNIVERSITY OF TEXAS AT AUSTIN

# OpenMP 3.0 and above

- First update to the spec since 2005
- Tasking: move beyond loops with generalized tasks and support complex and dynamic control flows (dependences), priority and other features.
- Loop collapse: combine nested loops automatically to expose more concurrency
- Nested parallelism support: better definition of and control over nested parallel regions, and new API routines to determine nesting structure
- OpenMP Affinity
- User-defined reductions
- Specify SIMD instructions for loop
- Offloading to devices (GPUs)
- Provide features implementation must provide (requires) and conditional directives (metadirective) 48



#### References

- http://www.openmp.org/
- **Parallel Programming in OpenMP**, by Chandra, Dagum, Kohr,  $\bullet$ Maydan, McDonald, Menon
- **Using OpenMP**, by Chapman, Jost, Van der Pas (OpenMP2.5)
- Using OpenMP The Next Step Affinity, Accelerators, Tasking and SIMD, Rudd van der Pas, Eric Stotzer, and Christian Terboven
- http://www.nic.uoregon.edu/iwomp2005/iwomp2005 tutorial openmp rvdp.pdf •

THE UNIVERSITY OF TEXAS AT AUS



Kent Milfeld Lars Koesterke milfeld|lars@tacc.utexas.edu

> For more information: www.tacc.utexas.edu



TACC

THE UNIVERSITY OF TEXAS AT AUSTIN

50

Miscellaneous and backups



THE UNIVERSITY OF TEXAS AT AUSTIN

51

**TEXAS ADVANCED COMPUTING CENTER** 

# Default variable scoping (Fortran example)

```
Program Main
Integer, Parameter :: nmax=100
Integer :: n, j
                                    SAVE array sum
Real*8 :: x(n,n)
                                    Integer :: i, m
Common /vars/ y(nmax)
                                    Real*8 :: a(m,m)
• • •
n=nmax; y=0.0
                                    do i=1,m
!$OMP Parallel do
  do j=1,n
                                    end do
   call Adder(x,n,j)
   end do
. . .
End Program Main
```

Subroutine Adder(a,m,col) Common /vars/ y(nmax)

```
y(col) = y(col) + a(i, col)
array sum=array sum+y(col)
```

End Subroutine Adder

52





# Default data scoping in Fortran (cont.)

Variable	Scope	Is use safe?	Reason for scope
n	shared	yes	declared outside parallel construct
j	private	yes	parallel loop index variable
x	shared	yes	declared outside parallel construct
У	shared	yes	common block
i	private	yes	parallel loop index variable
m	shared	yes	actual variable <i>n</i> is shared
а	shared	yes	actual variable x is shared
col	private	yes	actual variable <i>j</i> is private
array_sum	shared	no	declared with SAVE attribute



THE UNIVERSITY OF TEXAS AT AUSTIN

# C/C++

# Loop Collapse

- Allow collapsing of perfectly nested loops
- Will form a single loop and then parallelize it:

```
#pragma omp parallel do collapse(2)
for(i=0;i<n;i++){
   for(j=0;j<n;j++){
        .....
   }
}</pre>
```

54

THE UNIVERSITY OF TEXAS AT AUSTIN

#### **F90**

## Loop Collapse

- Allow collapsing of perfectly nested loops
- Will form a single loop and then parallelize it:

```
!$omp parallel do collapse(2)
do i=1,n
    do j=1,n
        ....
    end do
end do
```



55

### Loop Nesting



While OpenMP 3.0 supports nested parallelism, many implementations may ignore the nesting by serializing the inner parallel regions



56

# Parallel Regions & Modes

There are two OpenMP "modes"

- *static* mode
  - Fixed number of threads -- set in the OMP\_NUM\_THREADS env.

Or the threads may be set by a function call (or clause) inside the code:

- omp\_set\_num\_threads runtime function
  num\_threads(#) clause
- *dynamic* mode:
  - Number of threads can change under OS control from one parallel region to another using:

*Note*: the user can only define the maximum number of threads, compiler can use a smaller number

