



1


Parallel Programming using OpenMP



Oregon State University
Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/)



openmp.pptx

mjb - April 9, 2024

1


2

OpenMP Multithreaded Programming

- OpenMP stands for "Open Multi-Processing"
- OpenMP is a multi-vendor (see next page) standard to perform shared-memory multithreading
- OpenMP is both compiler-directive- and library-based
- OpenMP threads share a single executable, a single global memory, and a single heap (malloc, new)
- Each OpenMP thread has its own stack (function arguments, function return address, local variables)
- Using OpenMP usually requires no dramatic code changes
- OpenMP probably gives you the biggest multithread benefit per amount of work you have to put in to using it

Much of your use of OpenMP will be accomplished by issuing C/C++ "pragmas" to tell the compiler how to build the threads into your executable, like this:

#pragma omp directive [clause]



mjb - April 9, 2024

2

3

Who is in the OpenMP Consortium?





mjb - April 9, 2024


3

4

What OpenMP Isn't:

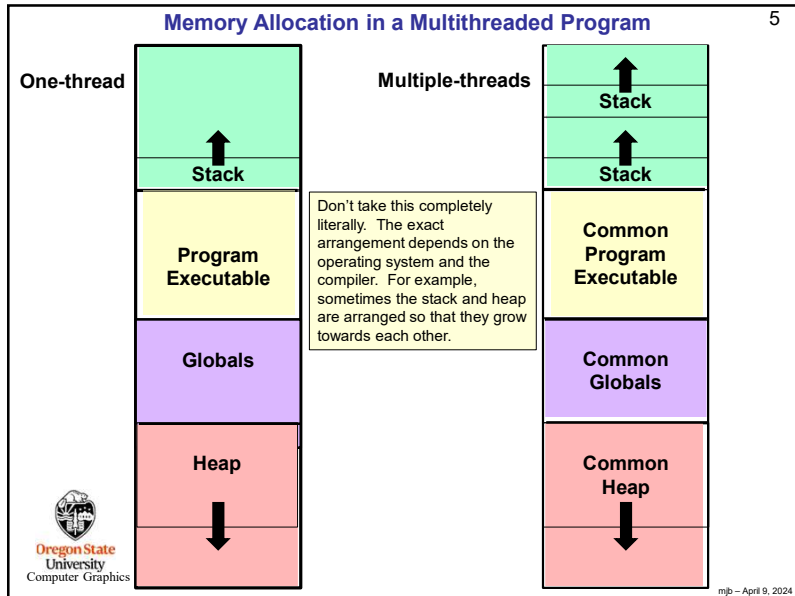
- OpenMP doesn't check for data dependencies, data conflicts, deadlocks, or race conditions. You are responsible for avoiding those yourself
- OpenMP doesn't check for non-conforming code sequences (we'll talk about what this means)
- OpenMP doesn't guarantee **identical** behavior across vendors or hardware, or even between multiple runs on the same vendor's hardware
- OpenMP doesn't guarantee the **order** in which threads execute, just that they do execute
- OpenMP is not overhead-free
- OpenMP does not prevent you from writing code that triggers cache performance problems (such as in false-sharing), in fact, *it makes it really easy*

We will get to "false sharing" in the cache notes



mjb - April 9, 2024

4



5

Using OpenMP on Linux

`g++ -o proj proj.cpp -lm -fopenmp`

Using OpenMP in Microsoft Visual Studio

1. Go to the Project menu → Project Properties
2. Change the setting Configuration Properties → C/C++ → Language → OpenMP Support to **"Yes (/openmp)"**

If you are using Visual Studio and get a compile message that looks like this:
`>c1xx: error C2338: two-phase name lookup is not supported for C++/CLI, C++/CX, or OpenMP; use /Zc:twoPhase-`
 then do this:

1. Go to "Project Properties" → "C/C++" → "Command Line"
2. Add `/Zc:twoPhase-` in "Additional Options" in the bottom section
3. Press OK

mjb - April 9, 2024

6

Seeing if OpenMP is Supported on Your System

```

#ifdef _OPENMP
    fprintf( stderr, "OpenMP release %d is supported here\n", _OPENMP );
#else
    fprintf( stderr, "OpenMP is not supported here - sorry!\n" );
    return 1;
#endif
    
```

Printing `_OPENMP` gives you a year and month of the OpenMP release that you are using

To get the OpenMP version number from the year and month, check here:

OpenMP 5.0 – November 2018
 OpenMP 4.5 – November 2015
 OpenMP 4.0 – July 2013
 OpenMP 3.1 – July 2011
 OpenMP 2.0 – March 2002
 OpenMP 1.0 – October 1998

- By default, flip uses g++ 11.4, which uses OpenMP version 4.5
- Visual Studio 2022 uses OpenMP 2.0

mjb - April 9, 2024

7

Numbers of OpenMP threads

How to specify how many OpenMP threads you want to have available:

```
omp_set_num_threads( num );
```

Asking how many cores this program has access to:

```
num = omp_get_num_procs( );
```

← Actually returns the number of hyperthreads, not the number of *physical* cores

Setting the number of available threads to the exact number of cores available:

```
omp_set_num_threads( omp_get_num_procs( ) );
```

Asking how many OpenMP threads this program is using right now:

```
num = omp_get_num_threads( );
```

Asking which thread number this one is:

```
me = omp_get_thread_num( );
```

mjb - April 9, 2024

8

OpenMP for-Loop Rules

13

```
#pragma omp parallel for default(none), shared(...), private(...)  
for( int index = start ; index terminate condition; index changed )
```

- The *index* must be an *int* or a *pointer*
- The *start* and *terminate* conditions must have compatible types
- Neither the *start* nor the *terminate* conditions can be changed during the execution of the loop
- The *index* can only be modified by the *changed* expression (i.e., not modified inside the loop itself)
- You cannot use a *break* or a *goto* to get out of the loop
- There can be no inter-loop data dependencies such as:
a[i] = a[i-1] + 1.;

```
a[101] = a[100] + 1.; // what if this is the last line of thread #0's work?  
a[102] = a[101] + 1.; // what if this is the first line of thread #1's work?
```



mjb - April 9, 2024

13

OpenMP For-Loop Rules

14

```
for( index = start ; index < end ; index <= end ; index > end ; index >= end )  
    index++  
    ++index  
    index--  
    --index  
    index += incr  
    index = index + incr  
    index = incr + index  
    index -= decr  
    index = index - decr
```



mjb - April 9, 2024

14

What to do about Variables Declared Before the for-loop Starts?

15

```
float x = 0.;;  
#pragma omp parallel for ...  
for( int i = 0; i < N; i++ )  
{  
    x = (float) i;  
    float y = x*x;  
    << more code... >  
}
```

i and *y* are automatically *private* because they are defined within the loop.

Good practice demands that *x* be explicitly declared to be shared or private!

private(x)

Means that each thread will get its own version of the variable

shared(x)

Means that all threads will share a common version of the variable

default(none)

I recommend that you include this in your OpenMP for-loop directive. This will force you to explicitly flag all of your externally-declared variables as *shared* or *private*. Don't make a mistake by leaving it up to the default!

```
Example:  
#pragma omp parallel for default(none), private(x)
```



mjb - April 9, 2024

15

For-loop "Fission"

16

Because of the loop dependency, this whole thing is not parallelizable:

```
x[ 0 ] = 0.;;  
y[ 0 ] *= 2.;;  
for( int i = 1; i < N; i++ )  
{  
    x[ i ] = x[ i-1 ] + 1.;;  
    y[ i ] *= 2.;;  
}
```

But it *can* be broken into one loop that is not parallelizable, plus one that is:

```
x[ 0 ] = 0.;;  
for( int i = 1; i < N; i++ )  
{  
    x[ i ] = x[ i-1 ] + 1.;;  
}  
  
#pragma omp parallel for shared(y)  
for( int i = 0; i < N; i++ )  
{  
    y[ i ] *= 2.;;  
}
```



mjb - April 9, 2024

16

For-loop "Collapsing"

17

Uh-oh, which for-loop do you put the #pragma on?

```
for( int i = 1; i < N; i++ )
{
    for( int j = 0; j < M; j++ )
    {
        ...
    }
}
```

Ah-ha – trick question. You put it on both!

```
#pragma omp parallel for collapse(2)
for( int i = 1; i < N; i++ )
{
    for( int j = 0; j < M; j++ )
    {
        ...
    }
}
```


How many for-loops to collapse into one loop

17

Single Program Multiple Data (SPMD) in OpenMP

18

```
#define NUM 1000000
float A[NUM], B[NUM], C[NUM];
...
int total = omp_get_num_threads( );
#pragma omp parallel default(none),shared(total)
{
    int me = omp_get_thread_num( );
    DoWork( me, total );
}
```



```
void DoWork( int m, int t )
{
    int first = NUM * m / t;
    int last = NUM * (m+1)/t - 1;
    for( int i = first; i <= last; i++ )
    {
        C[ i ] = A[ i ] * B[ i ];
    }
}
```

18

OpenMP Allocation of Work to Threads

19

Static Threads

- All work is allocated and assigned at runtime

Dynamic Threads

- The pool is statically assigned some of the work at runtime, but not all of it
- When a thread from the pool becomes idle, it gets a new assignment
- "Round-robin assignments"

OpenMP Scheduling

```
schedule(static [, chunksize])
schedule(dynamic [, chunksize])
Defaults to static
chunksize defaults to 1
```

19

OpenMP Allocation of Work to Threads

20

```
#pragma omp parallel for default(none),schedule(static,chunksize)
for( int index = 0 ; index < 12 ; index++ )
```

Static, 1

```
0 0,3,6,9
1 1,4,7,10
2 2,5,8,11
```

chunksize = 1

Each thread is assigned one iteration, then the assignments start over

Static, 2

```
0 0,1,6,7
1 2,3,8,9
2 4,5,10,11
```

chunksize = 2

Each thread is assigned two iterations, then the assignments start over

Static, 4

```
0 0,1,2,3
1 4,5,6,7
2 8,9,10,11
```

chunksize = 4

Each thread is assigned four iterations, then the assignments start over

Think of dealing for-loop passes to threads the same way as dogs deal cards ©



20

Arithmetic Operations Among Threads – A Problem

21

```
float sum = 0.;
#pragma omp parallel for default(none), shared(sum)
for( int i = 0; i < N; i++ )
{
    float myPartialSum = ...

    sum = sum + myPartialSum;
}

```

- There is no guarantee when each thread will execute this line
- There is not even a guarantee that each thread will finish this line before some other thread interrupts it. (Remember that each line of code usually generates multiple lines of assembly.)
- This is non-deterministic !

Assembly code:
 Load sum
 Add myPartialSum
 Store sum

What if the scheduler decides to switch threads right here?

Conclusion: Don't do it this way!

mjb - April 9, 2024

21

Here's a trapezoid integration example.

The partial sums are added up, as shown on the previous page.
 The integration was done 30 times.
 The answer is supposed to be exactly 2.
 None of the 30 answers is even close.
 And not only are the answers *bad*, but they are not even consistently *bad*!

0.469635	0.398893
0.517984	0.446419
0.438868	0.431204
0.437553	0.501783
0.398761	0.334996
0.506564	0.484124
0.489211	0.506362
0.584810	0.448226
0.476670	0.434737
0.530668	0.444919
0.500062	0.442432
0.672593	0.548837
0.411158	0.363092
0.408718	0.544778
0.523448	0.356299

Don't do it this way! We'll talk about how to do it correctly in the Trapezoid Integration noteset.

mjb - April 9, 2024

22

Synchronization

23

Mutual Exclusion Locks (Mutexes)

```
omp_init_lock( omp_lock_t* );
omp_set_lock( omp_lock_t* );
omp_unset_lock( omp_lock_t* );
omp_test_lock( omp_lock_t* );
```

Blocks if the lock is not available
Then sets it and returns when it is available

If the lock is not available, returns 0
If the lock is available, sets it and returns 10

(omp_lock_t is really an array of 4 unsigned chars)

Critical sections

```
#pragma omp critical
```

Restricts execution to one thread at a time

```
#pragma omp single
```

Restricts execution to a single thread ever

Barriers

```
#pragma omp barrier
```

Forces each thread to wait here until all threads arrive

(Note: there is an implied barrier after parallel for loops and OpenMP sections, unless the *nowait* clause is used)

mjb - April 9, 2024

23

Synchronization Example

24

```
omp_lock_t Sync;
...
omp_init_lock( &Sync );
...

```

Thread #0:

```
omp_set_lock( &Sync );
<< code that needs the mutual exclusion >>
omp_unset_lock( &Sync );
```

Thread #1:

```
omp_set_lock( &Sync );
<< code that needs the mutual exclusion >>
omp_unset_lock( &Sync );
```

mjb - April 9, 2024

24

Synchronization Example

25

```
omp_lock_t   Sync;
...
omp_init_lock( &Sync );

...

Thread #0:   Thread #1:
while( omp_test_lock( &Sync ) == 0 )   while( omp_test_lock( &Sync ) == 0 )
{                                           {
    DoSomeUsefulWork_0();                 DoSomeUsefulWork_1();
}
```



Oregon State
University
Computer Graphics

mjb - April 9, 2024

25

Single-thread-execution Synchronization

26

#pragma omp single

Restricts execution to a single thread ever. This is used when an operation only makes sense for one thread to do. Reading data from a file is a good example.



Oregon State
University
Computer Graphics

mjb - April 9, 2024

26

Creating Sections of OpenMP Code

27

Sections are independent blocks of code, able to be assigned to separate threads if they are available.

#pragma omp parallel sections

```
{
    #pragma omp section
    {
        Task 1
    }
    #pragma omp section
    {
        Task 2
    }
}
```

(Note: there is an **implied** barrier after parallel for loops and OpenMP sections, unless the *nowait* clause is used)



Oregon State
University
Computer Graphics

mjb - April 9, 2024

27

A Functional Decomposition Sections Example

28

```
omp_set_num_threads( 3 );

#pragma omp parallel sections
{
    #pragma omp section
    {
        Watcher( );
    }

    #pragma omp section
    {
        Animals( );
    }

    #pragma omp section
    {
        Plants( );
    }
} // implied barrier -- all functions must return to get past here
```



Oregon State
University
Computer Graphics

mjb - April 9, 2024

28

A Potential OpenMP/Visual Studio Compiler Problem

If you print to standard error (stderr) from inside a for-loop, like I do, then you think that you need to include *stderr* in the shared list because, well, the loops share it:

```
#pragma omp parallel for default(none) shared(a,b,stderr)
```

This turns out to be true for *g++/gcc only*.

If you are using Visual Studio, then *do not* include *stderr* in the list.

If you do, you will get this error:

```
1>Y:\ICS575\SQ22\robertw5-01\Project1\Project1.cpp(113,98): error C2059: syntax error: '('
```

This is because:

- In *g++/gcc*, *stderr* is a *variable*
- In *Visual Studio*, *stderr* is a *defined macro*

