



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/)



Oregon State University  
Computer Graphics

openmp.cpp

mjb - March 14, 2023


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 uses the fork-join model
- OpenMP is both directive- and library-based
- OpenMP threads share a single executable, global memory, and heap (malloc, new)
- Each OpenMP thread has its own stack (function arguments, function return address, local variables)
- Using OpenMP 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 the executable

**#pragma omp directive [clause]**



Oregon State University  
Computer Graphics

mjb - March 14, 2023

3

Who is in the OpenMP Consortium?





Oregon State University  
Computer Graphics


mjb - March 14, 2023

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
- 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



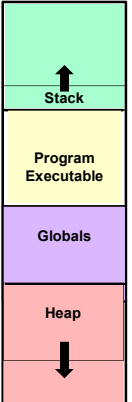
Oregon State University  
Computer Graphics

mjb - March 14, 2023


5

Memory Allocation in a Multithreaded Program


One-thread



Multiple-threads



Don't take this completely literally. The exact arrangement depends on the operating system and the compiler. For example, sometimes the stack and heap are arranged so that they grow towards each other.



Oregon State University  
Computer Graphics

mjb - March 14, 2023

6

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 2019 and get a compile message that looks like this:  
**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



Oregon State University  
Computer Graphics

mjb - March 14, 2023

## Seeing if OpenMP is Supported on Your System

```
#ifndef _OPENMP
fprintf( stderr, "OpenMP version %d is supported here!\n", _OPENMP );
#else
fprintf( stderr, "OpenMP is not supported here – sorry!\n" );
exit( 0 );
#endif
```

This gives you a year and month of the OpenMP you are using

To get an OpenMP version number:

OpenMP 5.0 – November 2018  
 OpenMP 4.5 – November 2015  
 OpenMP 4.0 – July 2013  
 OpenMP 3.1 – July 2011



- By default, flip is using g++ 4.8.5, which uses OpenMP version 3.1
- Flip's g++ 9.2.0 uses OpenMP version 4.5
- Looks like Visual Studio 2019's is even older (?)

mp – March 14, 2023

## 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( );
```

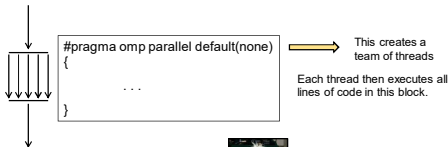
Asking which thread number this one is:

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



mp – March 14, 2023

## Creating an OpenMP Team of Threads



Think of it this way:



mp – March 14, 2023

## Creating an OpenMP Team of Threads

```
#include <stdio.h>
#include <omp.h>
int
main( )
{
    omp_set_num_threads( 8 );
    #pragma omp parallel default(none)
    {
        printf( "Hello, World, from thread #%d !\n" , omp_get_thread_num( ) );
    }
    return 0;
}
```

Hint: run it several times in a row. What do you see? Why?



mp – March 14, 2023

## Uh-oh...

### First Run

```
Hello, World, from thread #6 !
Hello, World, from thread #1 !
Hello, World, from thread #7 !
Hello, World, from thread #4 !
Hello, World, from thread #5 !
Hello, World, from thread #3 !
Hello, World, from thread #2 !
Hello, World, from thread #0 !
```

### Second Run

```
Hello, World, from thread #0 !
Hello, World, from thread #7 !
Hello, World, from thread #4 !
Hello, World, from thread #6 !
Hello, World, from thread #1 !
Hello, World, from thread #3 !
Hello, World, from thread #5 !
Hello, World, from thread #2 !
```

### Third Run

```
Hello, World, from thread #2 !
Hello, World, from thread #5 !
Hello, World, from thread #0 !
Hello, World, from thread #7 !
Hello, World, from thread #1 !
Hello, World, from thread #3 !
Hello, World, from thread #4 !
Hello, World, from thread #6 !
```

### Fourth Run

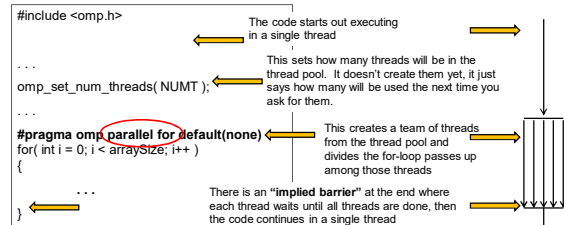
```
Hello, World, from thread #1 !
Hello, World, from thread #5 !
Hello, World, from thread #2 !
Hello, World, from thread #4 !
Hello, World, from thread #7 !
Hello, World, from thread #6 !
Hello, World, from thread #0 !
```



There is no guarantee of thread execution order!

mp – March 14, 2023

## Creating OpenMP threads in Loops



This tells the compiler to parallelize the for-loop into multiple threads. Each thread automatically gets its own personal copy of the variable *i* because it is defined within the for-loop body.

The default(none) directive forces you to explicitly declare all variables declared outside the parallel region to be either private or shared while they are in the parallel region. Variables declared within the for-loop are automatically private.

mp – March 14, 2023

### OpenMP for-Loop Rules

**#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?

Oregon State University  
Computer Graphics  
mp - March 14, 2023

### OpenMP For-Loop Rules

```

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

```

Oregon State University  
Computer Graphics  
mp - March 14, 2023

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

`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)**

Oregon State University  
Computer Graphics  
mp - March 14, 2023

### For-loop "Fission"

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.;
}

```

Oregon State University  
Computer Graphics  
mp - March 14, 2023

### For-loop "Collapsing"

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!

How many for-loops to collapse into one loop

```

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

```

Oregon State University  
Computer Graphics  
mp - March 14, 2023

### Single Program Multiple Data (SPMD) in OpenMP

```

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

```

```

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

```

Oregon State University  
Computer Graphics  
mp - March 14, 2023

## 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




mp - March 14, 2023

## 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	<b>chunksize = 1</b> 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	<b>chunksize = 2</b> 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	<b>chunksize = 4</b> Each thread is assigned four iterations, then the assignments start over



mp - March 14, 2023

## Arithmetic Operations Among Threads – A Problem

21

```
#pragma omp parallel for private(myPartialSum), 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!**



mp - March 14, 2023


## Here's a trapezoid integration example.

22

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, 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.**

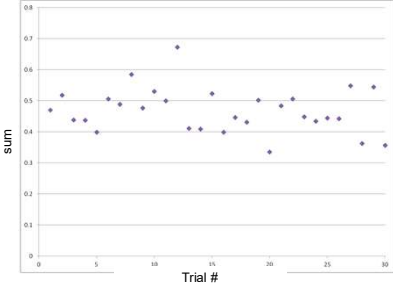


mp - March 14, 2023


## Here's a trapezoid integration example.

23

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, they are not even consistently bad!



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



mp - March 14, 2023

## Synchronization

24

**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 * );
```

( 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)



mp - March 14, 2023

### Synchronization Example

25

```


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 );

```

 Oregon State University  
Computer Graphics

mp - March 14, 2023

### Synchronization Example

26

```


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

...

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

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

```

 Oregon State University  
Computer Graphics


mp - March 14, 2023

### Single-thread-execution Synchronization

27

**#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

mp - March 14, 2023

### Creating Sections of OpenMP Code

28

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)

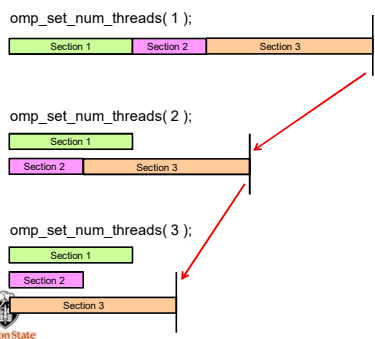
 Oregon State University  
Computer Graphics


mp - March 14, 2023

### What do OpenMP Sections do for You?

29

They decrease your overall execution time.



 Oregon State University  
Computer Graphics

mp - March 14, 2023

### A Functional Decomposition Sections Example

30

```


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

mp - March 14, 2023

### A Potential OpenMP/Visual Studio Compiler Problem

31

If you are using Visual Studio 2019 and get a compile message that looks like this:

```
1>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



mp - March 14, 2023

### Another Potential OpenMP/Visual Studio Compiler Problem

32

If you print to standard error (stderr), like I do, then you think that you need to include *stderr* in the shared list because, well, you use 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:\CS575\SQ22\robertw5-01\Project1\Project1.cpp(113,98): error C2059: syntax error: '('
```



mp - March 14, 2023