# COMP 605: Introduction to Parallel Computing Topic: OpenMP Parallel For Directives

#### Mary Thomas

Department of Computer Science Computational Science Research Center (CSRC) San Diego State University (SDSU)

> Presented: 04/13/17 Last Update: 04/10/17

COMP 605: Topic Presented: 04/13/1	Last Update: 04/10/17	2/39 Mary Thomas	
------------------------------------	-----------------------	------------------	--







COMP 605: Topic Presented: 04/13/17 Last Update: 04/10/17

OpenMP: Parallel for Directive

Parallel For Operator



Good News/Bad News: threads divide up work themselves (Good), but you cannot control this (maybe Bad). Not useful when you need/want to avoid default data composition.

4/39 Mary Thomas

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	5/39
OpenMP:	Parallel <i>for</i> D	Directive		
Parallel	For Operator			





- pointer type (e.g., it can' t be a float).
- The expressions start, end, and incr must have a compatible type. For example, if index is a pointer, then incr must have integer type.



Copyright © 2010, Elsevier Inc. All rights Reserved









#### Pacheco Data Dependencies example: omp\_fibo.c

```
#include <stdio.h>
#include <stdlib.h>
#include <omp.h>
int main(int argc, char* argv[]) {
   int thread count. n. i:
  long long* fibo;
  if (argc != 3) Usage(argv[0]);
  thread_count = strtol(argv[1], NULL, 10);
  n = strtol(argv[2], NULL, 10);
  fibo = malloc(n*sizeof(long long));
  fibo[0] = fibo[1] = 1;
# pragma omp parallel for num threads(thread count)
  for (i = 2; i < n; i++)
     fibo[i] = fibo[i-1] + fibo[i-2]:
  printf("The first n Fibonacci numbers:\n");
  for (i = 0; i < n; i++)
     printf("%d\t%lld\n", i, fibo[i]);
  free(fibo);
  return 0;
} /* main */
```

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	10/39	Mary Thomas
OpenMP:	Parallel fo	or Directive			
Parallel	For Operation	ator			

Pacheco Data Dependencies example: local laptop had thread problems, but this did not happen on tuckoo (up to 10,000 threads, no problems)

[gidget:i	ntro-par-pgming-pacheco/ipp-source/ch5] mthomas% gcc -fopenmp -o omp_fibo omp_fibo.
[gidget:i	ntro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp_fibo 10 10
The first	n Fibonacci numbers:
0	1
1	1
2	2
3	1
4	3
5	1
6	0
7	0
8	0
9	0
[gidget:i	ntro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp_fibo 10 10
The first	n Fibonacci numbers:
0	1
1	1
2	2
3	1
4	3
5	4
6	7
7	0
8	0
9	0
Lgidget:i	ntro-par-pgming-pacheco/ipp-source/ch5] mthomas%



# What went wrong?

- OpenMP compilers dont check for dependences among iterations in a loop thats being parallelized with a parallel for directive.
- A loop in which the results of one or more iterations depend on other iterations cannot, in general, be correctly parallelized by OpenMP.
- Dependencies within the current loop is OK:
   #pragma omp parallel for numthreads (thread count)

$$x[i] = a + i * h$$
  
$$y[i] = exp(x[i])$$





M<

Copyright © 2010, Elsevier Inc. All rights Reserved

43







```
4\sum_{k=0}^{\infty}\frac{(-1)^k}{2k+1}
```

Examine the series: for pi: Replace  $(-1)^k$ 

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	15/39	Mary Thomas
OpenMP:	Parallel <i>for</i>	Directive			
Parallel	For Operato				

#### Pacheco Data Dependencies example: omp\_pi.c

```
int main(int argc, char* argv[]) {
  long long n, i;
   int thread count:
   double factor:
   double sum = 0.0:
   if (argc != 3) Usage(argv[0]):
   thread count = strtol(argv[1], NULL, 10);
   n = strtoll(argv[2], NULL, 10);
   if (thread_count < 1 || n < 1) Usage(argv[0]);
# pragma omp parallel for num threads(thread count) \
      reduction(+; sum) private(factor)
  for (i = 0; i < n; i++) {
      factor = (i \% 2 == 0) ? 1.0 ; -1.0;
      sum += factor/(2*i+1);
   ifdef DEBUG
#
      printf("Thread %d > i = %lld, my_sum = %f\n", my_rank, i, my_sum);
#
      endif
  3
   s_{11m} = 4.0 * s_{11m}:
  printf("With n = %11d terms and %d threads, \n", n, thread_count);
   printf(" Our estimate of pi = %.14f\n", sum);
  printf("
                             pi = %.14f\n", 4.0*atan(1.0));
   return 0;
} /* main */
```



#### Running omp\_pi.c: could reproduce erroneous condition on OS Mountain Lion - not every time

```
With n = 100 terms and 10 threads.
   Our estimate of pi = 3,13159290355855
                  pi = 3,14159265358979
[gidget:intro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp_pi 1 1000
With n = 1000 terms and 1 threads.
   Our estimate of pi = 3,14059265383979
                   pi = 3.14159265358979
[gidget:intro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp pi-noprivatfactor 2 1000
With n = 1000 terms and 2 threads.
   Our estimate of pi = 3,14803451430491
                   pi = 3.14159265358979
dget:intro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp_pi 10 1000000
With n = 1000000 terms and 10 threads.
  Our estimate of pi = 3.14159165358972
                  pi = 3,14159265358979
[gidget:intro-par-pgming-pacheco/ipp-source/ch5] mthomas% ./omp_pi 10 10000000
With n = 10000000 terms and 10 threads.
   Our estimate of pi = 3.14159255358977
                   pi = 3.14159265358979
```





COMP 605: Topic Presented: 04/13/17 Last Update: 04/10/17 18/39 Mary Thomas OpenMP: Parallel for Directive Parallel For Operator





## **Bubble Sort:** $\vartheta(n^2)$



Loop carried dependence in outer loop - different threads sorting different data could result in wrong ordering.





#### See Pacheco, Chpt 3.7.

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	21/39	Mary Thomas
OpenMP:	Parallel <i>for</i> Di	rective			
Loops:	Sorting				

# **Odd-Even Transposition Sort:** $\vartheta(n)$

Seria	l Od	d-	Ev	en	T	ra	ns	pc	sition Sort
			Su	bsci	ipt ir	Ar	ray		
	Phase	0		1		2	1.5	3	
	0	9	$\leftrightarrow$	7		8	$\leftrightarrow$	6	
		7		9		6		8	
	1	7		9	$\leftrightarrow$	6		8	
		7		6		9		8	
	2	7	$\leftrightarrow$	6		9	$\leftrightarrow$	8	
		6		7		8		9	
	3	6		7	$\leftrightarrow$	8		9	
		6		7		8		9	



Copyright © 2010, Elsevier Inc. All rights Reserved

COMP 605: Topic Presented: 04/13/17 Last Update: 04/10/17

OpenMP: Parallel for Directive

Loops: Sorting



Mary Thomas

22/39

Issues: implicit barrier after each phase completes; overhead associated with fork/join.





odate: 04/10/17	24/39	Mary Thomas
	date: 04/10/17	date: 04/10/17 24/39

thread_count         1         2         3         4           Two parallel for directives         0.770         0.453         0.358         0.305           Two for directives         0.732         0.376         0.294         0.239	Odd-even sort with two parallel (Times a	for dire	ctives a conds.)	nd two	for direct	ives.
Two parallel for directives         0.770         0.453         0.305           Two for directives         0.732         0.376         0.294         0.239	thread_count	1	2	3	4	
Two for directives         0.732         0.376         0.294         0.239	Two parallel for directives	0.770	0.453	0.358	0.305	
	Two for directives	0.732	0.376	0.294	0.239	
					10	



**OpenMP** assigns/distributes work (block partitioning) **E.g.**, for *n* iterations, then each thread gets  $\frac{n}{thread\_count}$ Alternative approach: cyclic partitioning.



 COMP 605:
 Topic
 Presented:
 04/13/17
 Last Update:
 04/10/17
 26/39
 Mary Thomas

 OpenMP:
 Parallel for Directive
 Loops:
 Scheduling
 Vertice
 Vertice

```
double f(int i) {
   int j, start = i*(i+1)/2, finish = start + i;
   double return_val = 0.0;
   for (j = start; j <= finish; j++) {</pre>
       return_val += sin(j);
   return return_val;
  /* f */
                     Our definition of function f.
                 Copyright © 2010, Elsevier Inc. All rights Reserved
```

57











#### Scheduling: OpenMP assignment of threads.

```
The Schedule Clause
 Default schedule:
      sum = 0.0;
      pragma omp parallel for num_threads(thread_count) \
#
         reduction(+:sum)
      for (i = 0; i <= n; i++)
         sum += f(i):
 Cyclic schedule:
     sum = 0.0:
     pragma omp parallel for num_threads(thread_count) \
#
        reduction(+:sum) schedule(static .1)
     for (i = 0; i \le n; i++)
        sum += f(i);
```

Copyright © 2010, Elsevier Inc. All rights Reserved

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	30/39	Mary Thomas
OpenMP:	Parallel <i>for</i> D	Directive			
Loops:	Scheduling				

#### Scheduling test code(Pacheco, 5.7, pg 238)

```
double funtst(int): /* Thread function */
/*----*/
int main(int argc, char* argv[]) {
  int thread count = strtol(argv[1], NULL, 10);
  int i. n:
  double sum:
  sum = 0.0:
  n=100000:
# pragma omp parallel for num threads(thread count) \
    reduction(+:sum) schedule(static,1)
  for( i=0; i <=n; i++) {
     sum += funtst(i);
  3
  printf("N=%d, sum=%f\n",n,sum);
  return 0:
} /* main */
/*-----
* Function: Hello
* Purpose: Thread function that prints message
*/
double funtst(int i) {
   int j, start=i*(i+1)/2, finish=start+i;
   for(j=start;j<=finish;j++) {</pre>
      ret_val += sin(j);
   3
   return ret_val;
} /* funtst */
```

COMP 605:	Topic	Presented: 04/13/17	Last Update: 04/10/17	31/39	Mary Thomas
OpenMP:	Parallel for	Directive			
Loops:	Scheduling				

### results

==== WITHOUT SCHEDULING ========	
[gidget] mthomas% date ; ./omp_myt 1	; date
Tue Nov 6 15:46:51 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:50:18 PST 2012	
[gidget] mthomas% date ; ./omp_myt 2	date
Tue Nov 6 15:51:08 PST 2012	
N=100000, sum=-1.027596	
Tue Nov 6 15:52:53 PST 2012	
[gidget] mthomas% date ; ./omp_myt 8	; date
Tue Nov 6 15:46:12 PST 2012	
N=100000, sum=-1.931278	
Tue Nov 6 15:46:40 PST 2012	
[gidget] mthomas% date ; ./omp_myt 16	; date
Tue Nov 6 15:42:40 PST 2012	
N=100000, sum=7.887650	
Tue Nov 6 15:43:12 PST 2012	
[gidget] mthomas% date ; ./omp_myt 32	; date
Tue Nov 6 15:43:56 PST 2012	
N=100000, sum=22.965623	
Tue Nov 6 15:44:27 PST 2012	

==== WITH SCHEDULING ==========	-
[gidget]mthomas% date ; ./omp_myt_sch 1 ; o	date
Tue Nov 6 15:33:20 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:37:05 PST 2012	
[gidget] mthomas% date ; ./omp_myt_sch 2 ;	date
Tue Nov 6 15:42:24 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:44:51 PST 2012	
[gidget] mthomas% date ; ./omp_myt_sch 8 ;	date
Tue Nov 6 15:32:43 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:33:12 PST 2012	
[gidget] mthomas% date ; ./omp_myt_sch 16	; date
Tue Nov 6 15:32:05 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:32:34 PST 2012	
[gidget] mthomas% date; ./omp_myt_sch 32; o	date
Tue Nov 6 15:31:11 PST 2012	
N=100000, sum=-0.284234	
Tue Nov 6 15:31:40 PST 2012	



• The chunksize is a positive integer.

Copyright © 2010, Elsevier Inc. All rights Reserved



#### chunks assigned in round-robin fashion.

















- Each thread also executes a chunk, and when a thread finishes a chunk, it requests another one.
- However, in a guided schedule, as chunks are completed the size of the new chunks decreases.
- If no chunksize is specified, the size of the chunks decreases down to 1.
- If chunksize is specified, it decreases down to chunksize, with the exception that the very last chunk can be smaller than chunksize.



#### COMP 605: Topic Presented: 04/13/17 Last Update: 04/10/17

OpenMP: Parallel for Directive

Loops: Scheduling

Thread	Chunk	Size of Chunk	Remaining Iterations
0	1 - 5000	5000	4999
1	5001 - 7500	2500	2499
1	7501 - 8750	1250	1249
1	8751 - 9375	625	624
0	9376 - 9687	312	312
1	9688 - 9843	156	156
0	9844 - 9921	78	78
1	9922 - 9960	39	39
1	9961 - 9980	20	19
1	9981 - 9990	10	9
1	9991 - 9995	5	4
0	9996 - 9997	2	2
1	9998 - 9998	1	1
0	9999 - 9999	1	0

Assignment of trapezoidal rule iterations 1–9999 using a guided schedule with two threads.



Copyright © 2010, Elsevier Inc. All rights Reserved

67



