

Parallel Programming

CPSC 6109 - Algorithms Analysis and Design

Dr. Hyrum D. Carroll

March 27, 2024

Parallel Algorithms Module Objectives

The goal of this module is that by the end of it you will be able to:

1. Appreciate that modern computers have multiple processing units
2. Describe the role of a concurrency platform
3. Calculate speedup given performance results
4. Identify a race condition in code and explain why it is a race condition
5. Describe how and why an algorithm is parallized (for example, matrix multiple and mergesort)

Parallel Systems

We have access to several parallel architectures:

- ▶ Multicore machines
- ▶ GPUs (graphics processing units) with multiple processors
- ▶ Clusters
- ▶ Cloud computing

Parallel Systems

We have access to several parallel architectures:

- ▶ Multicore machines
- ▶ GPUs (graphics processing units) with multiple processors
- ▶ Clusters
- ▶ Cloud computing

Which ones have you used?

Parallelism in Your Computers


How many cores does your computer have?

Parallelism in Your Computers

How many cores does your computer have?


How many cores does your computer (your phone) have?

Parallelism in Your Computers



Source: Advanced Micro Devices

Parallelism in Your Computers



Source: OpenSPARC, Oracle.com

Memory Models

Two memory models:

1. Shared memory
2. Distributed memory

Memory Models

Two memory models:

1. Shared memory
 - ▶ All cores can access all of the memory
2. Distributed memory

Memory Models

Two memory models:

1. Shared memory

- ▶ All cores can access all of the memory

2. Distributed memory

- ▶ Cores have private memory, need to send messages to other cores

Memory Models

Two memory models:

1. Shared memory
 - ▶ All cores can access all of the memory
2. Distributed memory
 - ▶ Cores have private memory, need to send messages to other cores

Both memory models exist in practice

Shared Memory

Two main features of a the shared memory model:

1. Multiple threads
2. Concurrency platform handles scheduling (including load-balancing)

Parallel Programming in a Nutshell

Load balancing vs Communication

This is the eternal problem in parallel computing. The basic approaches to this problem include:

- ▶ Data partitioning - moving different parts of the data set across several nodes
- ▶ Task partitioning - give separate tasks to different nodes

Definition of Terms

Parallel processing terms:

- ▶ node - a box usually containing processors, local memory, disks and network connection
- ▶ cluster - a group of nodes networked together
- ▶ speedup: $S_p = \frac{T_1}{T_p}$
- ▶ efficiency: $\frac{S_p}{p} = \frac{T_1}{pT_p}$

(T_i is the execution time for i processors, p is the number of processors)

Speedup

- ▶ Adding more processors does not always improve the speed a code runs.
- ▶ Usually, better speedup can be found by increasing the problem size, at least to a point.
- ▶ The non-parallel part of a code generally scales linearly with the problem size. The parallel part usually scales with the problem size to some power.
- ▶ Generally increasing the problem size without increasing the node number helps performance.

Scalability

Good parallel algorithms run faster when more nodes are available. In the best case, doubling the number of nodes decreases the execution time by a factor of two.

One way to consider scaling of a code is Amdahl's law:

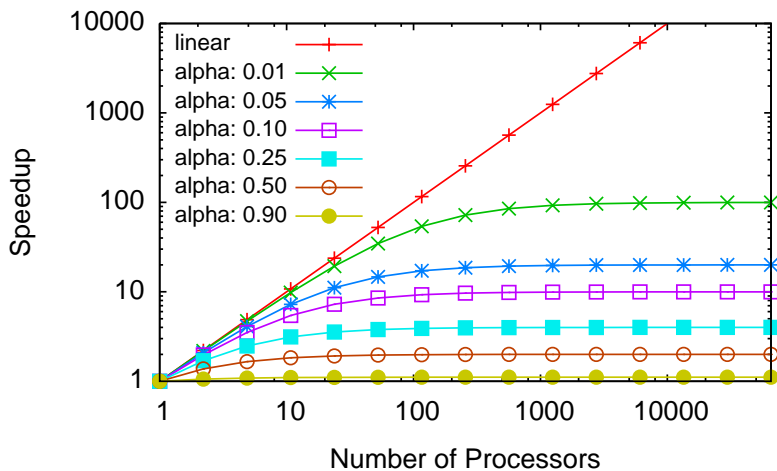
$$S_p = \frac{1}{\alpha + \frac{1-\alpha}{p}}$$

where α is the portion of the code which cannot be parallelized and p is the number of processors.

This is a simplification, but-

Speedup is limited by the slowest portion of the code.

Amdahl's Law



Communication

- ▶ Communication between nodes takes a great deal of time.
- ▶ Typically you can do thousands of computations in the time it takes to pass the simplest message.
- ▶ The time it takes for a message to be passed is limited by bandwidth b and latency l . To pass a message of size s , you need

$$\frac{s}{b} + l$$

(Assuming b , l , and s are in consistent units.)

Introduction to OpenMP

OpenMP is a concurrency platform.

It is as a set of simple program additions to make codes run efficiently on shared memory computers. The formal API for OpenMP is only about 50 pages long, and contains compiler directives and library functions.

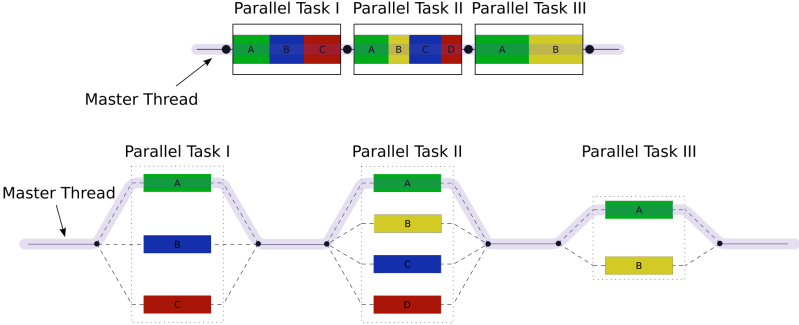
<http://www.llnl.gov/computing/tutorials/openMP/>

OpenMP Threads

OpenMP uses threads for parallel programming

- ▶ Forks and joins are used for most of the internal programming
- ▶ Speedup is achieved by the operating system splitting the threads across multiple CPUs.
- ▶ New threads are created explicitly by the program directives dynamically.

Forks and Joins



Source: Wikipedia user A1

Goals of OpenMP - from LLNL

- ▶ Standardization
- ▶ Lean and Mean - only 3-4 directives
- ▶ Ease of use
- ▶ Portability - F77, F90, F95, C, C++

Goals of OpenMP - from LLNL

- ▶ Standardization
- ▶ Lean and Mean - only 3-4 directives
- ▶ Ease of use
- ▶ Portability - F77, F90, F95, C, C++

Note: You can not use OpenMP and Java together

OpenMP Programming Model - from LLNL

- ▶ Shared Memory, thread based
- ▶ Explicit Parallelism
- ▶ Fork-Join Model
- ▶ Compiler Directives
- ▶ Nested Parallelism Support - in most implementations
- ▶ Dynamic Threads
- ▶ Not tied to I/O

Explicit Parallelism

- ▶ You must tell the computer what sections of code to parallelize using compiler directives.
- ▶ The compiler directives vary between languages, but are ignored when OpenMP flags are not set with the compiler.
- ▶ Codes written with OpenMP can run easily on serial machines.

Environment and Library Routines

- ▶ Some environmental variables are needed to make the code execute using the correct number of threads
- ▶ Some library routines allow the programmer to set and access system variables

Not Message Passing

This is NOT a set of message passing routines. Instead, you give directives to the compiler of what parts of the code can be executed in parallel.

In some ways, OpenMP is a set of directives to tell the compiler how to more efficiently handle loops.

General Syntax

- ▶ Fortran:

```
!$OMP <directive>  
do useful stuff  
!$OMP end <directive>
```

- ▶ C/C++:

```
#pragma omp <directive-name> clause  
{  
do useful stuff in a structured block  
}
```

A Trivial Example

Basic Code

```
1 program trivial
2   print*, 'Hello World!'
3 end program
```

```
% gfortran trivial.f90
% ./a.out
Hello World!
```

OMP Additions

```
1 program trivial
2
3 !$OMP PARALLEL
4   print*, 'Hello World!'
5 !$OMP END PARALLEL
6
7 end program trivial
```

```
% gfortran trivialOpenMP.f90
% ./a.out
Hello World!
```

A Trivial Example

Basic Code

```
1 program trivial
2   print*, 'Hello World!'
3 end program
```

```
% gfortran trivial.f90
% ./a.out
Hello World!
```

OMP Additions

```
1 program trivial
2
3 !$OMP PARALLEL
4   print*, 'Hello World!'
5 !$OMP END PARALLEL
6
7 end program trivial
```

```
% gfortran trivialOpenMP.f90
% ./a.out
Hello World!
```

What went wrong?

Execution of the Trivial Example

```
% gfortran trivialOpenMP.f90 -fopenmp
```

```
% ./a.out
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```

```
% export OMP_NUM_THREADS=3
```

```
% ./a.out
```

```
Hello World!
```

```
Hello World!
```

```
Hello World!
```


Thread ID

```
1 program trivial1
2   implicit none
3   integer :: OMP_GET_THREAD_NUM, OMP_GET_MAX_THREADS
4   integer :: tid , nthreads
5 !$OMP PARALLEL PRIVATE(nthreads , tid)
6   tid = OMP_GET_THREAD_NUM()
7   nthreads = OMP_GET_MAX_THREADS()
8   print*,'Hello World! from ', tid , nthreads
9 !$OMP END PARALLEL
10 end program
```

Note the PRIVATE key word, indicating that all threads have their own copy of the variable.

Thread ID (2)

```
% gfortran -fopenmp trivial1.f90
% ./a.out
Hello World! from      0      1
Hello World! from      2      1
Hello World! from      3      1
Hello World! from      4      1
Hello World! from      1      1
Hello World! from      7      1
Hello World! from      5      1
Hello World! from      6      1
```

Thread ID

```
1 program trivial2
2   implicit none
3   integer :: OMP_GET_THREAD_NUM, OMP_GET_MAX_THREADS
4   integer :: tid , nthreads
5   nthreads = OMP_GET_MAX_THREADS()
6   !$OMP PARALLEL PRIVATE(tid)
7   tid = OMP_GET_THREAD_NUM()
8   print*, 'Hello World! from ', tid , nthreads
9   !$OMP END PARALLEL
10 end program
```

Note that `nthreads` is outside of the OMP directives

Thread ID (2)

```
% gfortran -fopenmp trivial2.f90
```

```
% ./a.out
```

```
Hello World! from      5      8  
Hello World! from      0      8  
Hello World! from      1      8  
Hello World! from      2      8  
Hello World! from      7      8  
Hello World! from      3      8  
Hello World! from      4      8  
Hello World! from      6      8
```

Parallelizing Loops

To parallelize a loop, you need to help the compiler figure out the most efficient way to use threads. There are simple defaults, but giving it more details can help efficiency.

The basic directives are:

```
!$OMP PARALLEL
```

```
!$OMP DO
```

```
some parallel loop
```

```
!$OMP END DO
```

```
!$OMP END PARALLEL
```


A Simple OMP Example

omptest1

```
1 program omptest1
2   integer , parameter :: n = 10000
3   integer , parameter :: dble = selected_real_kind
4     (15,307)
5   real (kind=dble), dimension(n) :: a
6   integer :: i, j
7 !$OMP PARALLEL
8 !$OMP DO
9   do j = 1, 100000
10    do i = 1, n
11      a(i) = log(real(i)) + j
12    enddo
13  enddo
14 !$OMP END DO
15 !$OMP END PARALLEL
16   print*,a(1)
17 end program omptest1
```


Results

omptest1




Results

omptest1




Results

omptest1



Results

omptest1



Combining Directives


You do not have to have a separate directive on each line. For example,

```
!$OMP PARALLEL  
!$OMP DO  
!$OMP PRIVATE(NTHREADS, TID)
```

Becomes

```
!$OMP PARALLEL DO PRIVATE(NTHREADS, TID)
```

Numerical Integration



Numerical Integration

Integrating

$$\pi = \int_{-1/2}^{1/2} \frac{4}{1+x^2} dx \quad (1)$$

We can approximate this integral using Simpson's algorithms

- ▶ Input the number of partitions to be used
- ▶ Divide the domain into n partitions
- ▶ Evaluate the function at each partition
- ▶ Multiply the function evaluation times the width of the function to find a differential area
- ▶ Add the differential areas together
- ▶ Output the result

Parallel Integration

In parallel, the problem is nearly the same.

- ▶ Have processing element (PE) zero, get the number of partitions, n
- ▶ Determine the number of PEs: m
- ▶ Divide the domain into $\frac{n}{m}$ partitions on each PE
- ▶ Evaluate the function at each partition
- ▶ Multiply the function evaluation times the width of the function to find a differential area
- ▶ Add the differential areas together across all the PEs
- ▶ On PE zero, output the result

Simple Code to Calculate PI (serial version)

```
1 program reduce
2   integer :: i, num_steps
3   double precision :: x, pi, step, sum
4   sum = 0.0d0; nsteps = 10000
5   step = 1.0d0 / dble(nsteps)
6   do i = 1, nsteps
7       x = (dble(i) + 0.5d0) * step
8       sum = sum + 4.0d0 / (1.0d0 + x*x)
9   enddo
10  pi = step * sum
11  print *, "Estimate of Pi with ", nsteps, " steps is "
12        , pi
13 end program reduce
```

```
$ ./reduce
```

```
Estimate of Pi with 10000 steps is 3.1413926444243838
```

Simple Code to Calculate PI (parallel version)

```
1 program reduceOMP
2   integer :: i, num_steps
3   double precision :: x, pi, step, sum
4   sum =0.0d0; num_steps = 10000
5   step = 1.0d0 / dble(num_steps)
6   !$OMP PARALLEL DO
7     do i = 1, num_steps
8       x = (dble(i) + 0.5d0) * step
9       sum = sum + 4.0d0 / (1.0d0 + x*x)
10    enddo
11 !$OMP END PARALLEL DO
12   pi = step * sum
13   print *, "Estimate of Pi with ", num_steps, " steps
14         is ", pi
end program reduceOMP
```

```
$ gfortran -fopenmp reduceOMP.f90 -o reduceOMP
```

```
$ ./reduceOMP
```

```
Estimate of Pi with 10000 steps is 7.5588335781770253
```


Simple Code to Calculate PI (parallel version)

```
1 program reduceOMP
2   integer :: i, num_steps
3   double precision :: x, pi, step, sum
4   sum = 0.0d0; num_steps = 10000
5   step = 1.0d0 / dble(num_steps)
6   !$OMP PARALLEL DO
7     do i = 1, num_steps
8       x = (dble(i) + 0.5d0) * step
9       sum = sum + 4.0d0 / (1.0d0 + x*x)
10    enddo
11 !$OMP END PARALLEL DO
12   pi = step * sum
13   print *, "Estimate of Pi with ", num_steps, " steps
14         is ", pi
15 end program reduceOMP
```

```
$ gfortran -fopenmp reduceOMP.f90 -o reduceOMP
```

```
$ ./reduceOMP
```

```
Estimate of Pi with 10000 steps is 7.5588335781770253
```

What happened?


Race Conditions

A race condition exists when two processing units are accessing the same resource and one or both of them are writing to it.

Example from Introduction to Algorithms:

RACE-EXAMPLE()

```
1  $x = 0$ 
2 parallel for  $i = 1$  to 2
3      $x = x + 1$ 
4 print  $x$ 
```



(a)

step	x	r_1	r_2
1	0	-	-
2	0	0	-
3	0	1	-
4	0	1	0
5	0	1	1
6	1	1	1
7	1	1	1

(b)

Race Condition Examples

Identify the race condition (if any) from example from Introduction to Algorithms:

MAT-VEC-WRONG(A, x)

```
1   $n = A.rows$ 
2  let  $y$  be a new vector of length  $n$ 
3  parallel for  $i = 1$  to  $n$ 
4       $y_i = 0$ 
5  parallel for  $i = 1$  to  $n$ 
6      parallel for  $j = 1$  to  $n$ 
7           $y_i = y_i + a_{ij}x_j$ 
8  return  $y$ 
```

P-SQUARE-MATRIX-MULTIPLY(A, B)

```
1   $n = A.rows$ 
2  let  $C$  be a new  $n \times n$  matrix
3  parallel for  $i = 1$  to  $n$ 
4      parallel for  $j = 1$  to  $n$ 
5           $c_{ij} = 0$ 
6          for  $k = 1$  to  $n$ 
7               $c_{ij} = c_{ij} + a_{ik} \cdot b_{kj}$ 
8  return  $C$ 
```

Race Condition Examples

Identify the race condition (if any) from example from Introduction to Algorithms:

MAT-VEC-WRONG(A, x)

```
1  $n = A.rows$ 
2 let  $y$  be a new vector of length  $n$ 
3 parallel for  $i = 1$  to  $n$ 
4      $y_i = 0$ 
5 parallel for  $i = 1$  to  $n$ 
6     parallel for  $j = 1$  to  $n$ 
7          $y_i = y_i + a_{ij}x_j$ 
8 return  $y$ 
```

P-SQUARE-MATRIX-MULTIPLY(A, B)

```
1  $n = A.rows$ 
2 let  $C$  be a new  $n \times n$  matrix
3 parallel for  $i = 1$  to  $n$ 
4     parallel for  $j = 1$  to  $n$ 
5          $c_{ij} = 0$ 
6         for  $k = 1$  to  $n$ 
7              $c_{ij} = c_{ij} + a_{ik} \cdot b_{kj}$ 
8 return  $C$ 
```

In MAT-VEC-WRONG, multiple threads are writing to y_i

Race Condition Examples

Identify the race condition (if any) from example from Introduction to Algorithms:

MAT-VEC-WRONG(A, x)

```
1   $n = A.rows$ 
2  let  $y$  be a new vector of length  $n$ 
3  parallel for  $i = 1$  to  $n$ 
4       $y_i = 0$ 
5  parallel for  $i = 1$  to  $n$ 
6      parallel for  $j = 1$  to  $n$ 
7           $y_i = y_i + a_{ij}x_j$ 
8  return  $y$ 
```

P-SQUARE-MATRIX-MULTIPLY(A, B)

```
1   $n = A.rows$ 
2  let  $C$  be a new  $n \times n$  matrix
3  parallel for  $i = 1$  to  $n$ 
4      parallel for  $j = 1$  to  $n$ 
5           $c_{ij} = 0$ 
6          for  $k = 1$  to  $n$ 
7               $c_{ij} = c_{ij} + a_{ik} \cdot b_{kj}$ 
8  return  $C$ 
```

In MAT-VEC-WRONG, multiple threads are writing to y_i . In P-SQUARE-MATRIX-MULTIPLY, only one thread is writing to c_{ij} .

Reductions

Because the loops are executing separately, you may wish to combine the results from different threads to a final answer. You need to use reduction to make this work.

```
$!OMP PARALLEL PRIVATE(X) REDUCTION(+:SUM)
```

OpenMP Modifications

```
1 program reduceOMP2
2   integer :: i, num_steps
3   double precision :: x, pi, step, sum
4   sum = 0.0d0 ; nsteps = 100000000
5   step = 1.0d0 / dble(nsteps)
6   !$OMP PARALLEL DO PRIVATE(X) REDUCTION(+:SUM)
7   do i = 1, nsteps
8     x = (dble(i) + 0.5d0) * step
9     sum = sum + 4.0d0 / (1.0d0 + x*x)
10  enddo
11  !$OMP END PARALLEL DO
12  pi = step * sum
13  print *, "Estimate of Pi with ", nsteps, " steps is "
14        , pi
15 end program reduceOMP2
```


```
$ gfortran -fopenmp reduceOMP2.f90 -o reduceOMP2
```

```
$ ./reduceOMP2
```

```
Estimate of Pi with 10000 steps is 3.1413926444243732
```


Results

reduceOMP2



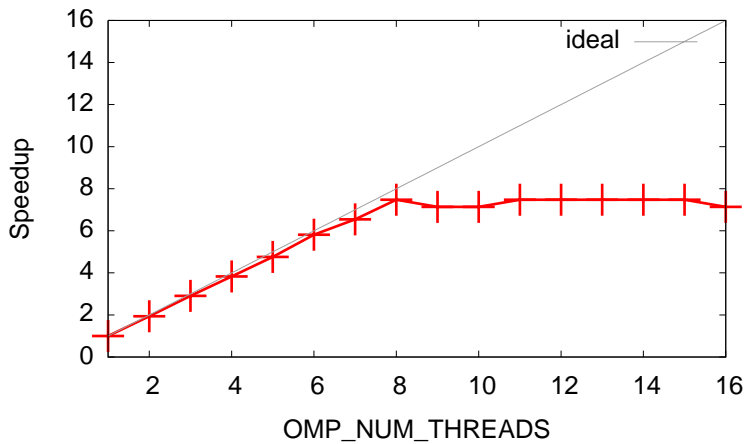
Results

reduceOMP2




Results

reduceOMP2



Results

reduceOMP2



Loop Splitting

One of the key ideas to remember is that loops often contain several operations that can be split. Taking an example from the Patterns in Parallel Programming book, imagine we have a loop with two functions:

- ▶ *BIG_COMPUTATION* - a big computation that executes independently on each element in the loop
- ▶ *COMBINE* - an element that cannot be parallelized and must execute in order

Loop Splitting

```
do i = 1, nsteps
  x = BIG_COMPUTATION(i)
  call COMBINE(x,answer)
enddo
```

can be split into

```
do i = 1, nsteps
  x(i) = BIG_COMPUTATION(i)
enddo
do i = 1, nsteps
  call COMBINE(x(i),answer)
enddo
```

Using OpenMP in Loop Splitting

```
!$OMP PARALLEL DO PRIVATE(I)
do i = 1, nsteps
  x(i) = BIG_COMPUTATION(i)
enddo
!$OMP END PARALLEL DO

do i = 1, nsteps
  call COMBINE(x(i),answer)
enddo
```

Controlling Loops

There are many options for controlling the execution of threads.

```
!$OMP DO SCHEDULE(TYPE, integer)
```

- ▶ `schedule(static[,chunk])` - groups of size `chunk` statically assigned in a round-robin fashion
- ▶ `schedule(dynamic[,chunk])` - threads dynamically grab work as it is completed
- ▶ `schedule(guided[,chunk])` - chunk size is reduced automatically during iteration toward a minimum level of `chunk`
- ▶ `schedule(runtime)` - checks the `OMP_SCHEDULE` environmental variable

Controlling Loops

```
integer, parameter :: chunk = 10

!$OMP PARALLEL PRIVATE(i,j,z,c,it) DEFAULT(SHARED)
!$OMP DO SCHEDULE(DYNAMIC,CHUNK)
  do i = 1, n
    do j = 1, n
...

```


Controlling Loops

```
setenv OMP_SCHEDULE static
```

```
11.477u 0.012s 0:08.24 139.3%
```

```
setenv OMP_SCHEDULE dynamic
```

```
11.239u 0.006s 0:05.67 198.0%
```

```
setenv OMP_SCHEDULE guided
```

```
11.453u 0.005s 0:06.52 175.6%
```

```
setenv OMP_SCHEDULE static,20
```

```
11.439u 0.028s 0:05.89 194.3%
```

```
no omp
```

```
11.280u 0.004s 0:11.28 100.0%
```

Examples

Multithread Matrix Multiplication

To multiply matrix A by matrix B to get matrix C we can divide the matrixes into submatrices:

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}, B = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix}, C = \begin{pmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{pmatrix}$$
$$\begin{pmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \quad (2)$$

$$\begin{pmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{pmatrix} = \begin{pmatrix} A_{11}B_{11} & A_{11}B_{12} \\ A_{21}B_{11} & A_{21}B_{12} \end{pmatrix} \begin{pmatrix} A_{12}B_{21} & A_{12}B_{22} \\ A_{22}B_{21} & A_{22}B_{22} \end{pmatrix} \quad (3)$$

Multithread Matrix Multiplication Pseudocode

P-MATRIX-MULTIPLY-RECURSIVE(C, A, B)

```
1   $n = A.rows$ 
2  if  $n == 1$ 
3       $c_{11} = a_{11}b_{11}$ 
4  else let  $T$  be a new  $n \times n$  matrix
5      partition  $A, B, C$ , and  $T$  into  $n/2 \times n/2$  submatrices
            $A_{11}, A_{12}, A_{21}, A_{22}; B_{11}, B_{12}, B_{21}, B_{22}; C_{11}, C_{12}, C_{21}, C_{22};$ 
           and  $T_{11}, T_{12}, T_{21}, T_{22}$ ; respectively
6      spawn P-MATRIX-MULTIPLY-RECURSIVE( $C_{11}, A_{11}, B_{11}$ )
7      spawn P-MATRIX-MULTIPLY-RECURSIVE( $C_{12}, A_{11}, B_{12}$ )
8      spawn P-MATRIX-MULTIPLY-RECURSIVE( $C_{21}, A_{21}, B_{11}$ )
9      spawn P-MATRIX-MULTIPLY-RECURSIVE( $C_{22}, A_{21}, B_{12}$ )
10     spawn P-MATRIX-MULTIPLY-RECURSIVE( $T_{11}, A_{12}, B_{21}$ )
11     spawn P-MATRIX-MULTIPLY-RECURSIVE( $T_{12}, A_{12}, B_{22}$ )
12     spawn P-MATRIX-MULTIPLY-RECURSIVE( $T_{21}, A_{22}, B_{21}$ )
13     P-MATRIX-MULTIPLY-RECURSIVE( $T_{22}, A_{22}, B_{22}$ )
14     sync
15     parallel for  $i = 1$  to  $n$ 
16         parallel for  $j = 1$  to  $n$ 
17              $c_{ij} = c_{ij} + t_{ij}$ 
```

Source: Introduction to Algorithms, 3rd Edition

Multithread Mergesort Pseudocode

P-MERGE-SORT(A, p, r, B, s)

```
1   $n = r - p + 1$ 
2  if  $n == 1$ 
3       $B[s] = A[p]$ 
4  else let  $T[1..n]$  be a new array
5       $q = \lfloor (p + r)/2 \rfloor$ 
6       $q' = q - p + 1$ 
7      spawn P-MERGE-SORT( $A, p, q, T, 1$ )
8      P-MERGE-SORT( $A, q + 1, r, T, q' + 1$ )
9      sync
10     P-MERGE( $T, 1, q', q' + 1, n, B, s$ )
```


Source: Introduction to Algorithms, 3rd Edition

Multithread Merge Pseudocode

```
P-MERGE( $T, p_1, r_1, p_2, r_2, A, p_3$ )
1   $n_1 = r_1 - p_1 + 1$ 
2   $n_2 = r_2 - p_2 + 1$ 
3  if  $n_1 < n_2$                                 // ensure that  $n_1 \geq n_2$ 
4      exchange  $p_1$  with  $p_2$ 
5      exchange  $r_1$  with  $r_2$ 
6      exchange  $n_1$  with  $n_2$ 
7  if  $n_1 == 0$                                 // both empty?
8      return
9  else  $q_1 = \lfloor (p_1 + r_1)/2 \rfloor$ 
10      $q_2 = \text{BINARY-SEARCH}(T[q_1], T, p_2, r_2)$ 
11      $q_3 = p_3 + (q_1 - p_1) + (q_2 - p_2)$ 
12      $A[q_3] = T[q_1]$ 
13     spawn P-MERGE( $T, p_1, q_1 - 1, p_2, q_2 - 1, A, p_3$ )
14     P-MERGE( $T, q_1 + 1, r_1, q_2, r_2, A, q_3 + 1$ )
15     sync
```

Source: Introduction to Algorithms, 3rd Edition

Multithread Merge Diagram



Source: Introduction to Algorithms, 3rd Edition