

Parallel Programming: Background Information and Tips



Oregon State
University

Mike Bailey
mjb@cs.oregonstate.edu



Three Reasons to Study Parallel Programming

1. Increase performance: do more work in the same amount of time
2. Increase performance: take less time to do the same amount of work
3. Make some programming tasks more convenient to implement

Example:
Decrease the time to compute an existing simulation program

Example:
Increase the resolution, and thus the accuracy, of a simulation program

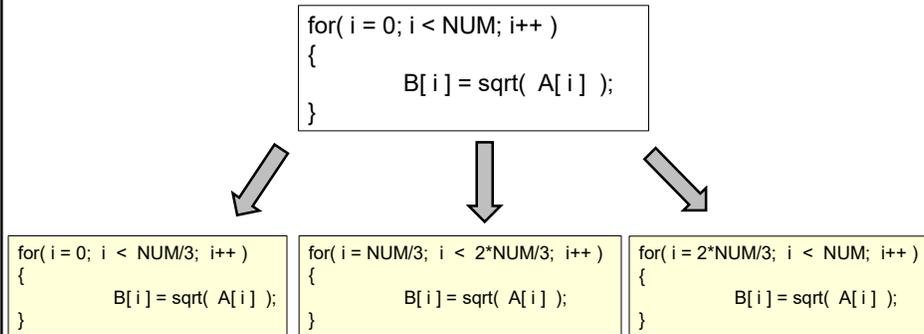
Example:
Create a web browser where the tasks of monitoring the user interface, downloading text, and downloading multiple images are happening simultaneously



Two Types of Parallelism: 1. Data Level Parallelism (DLP)

3

Threads are executing the same instructions on different data



mjb - March 14, 2024

3

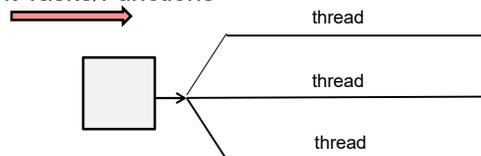
Two Types of Parallelism: 2. Thread (or Task or Functional) Level Parallelism (TLP)

4

Threads are executing *different* instructions

Example: processing a variety of incoming transaction requests

Different Tasks/Functions



In TLP you can have more threads than cores

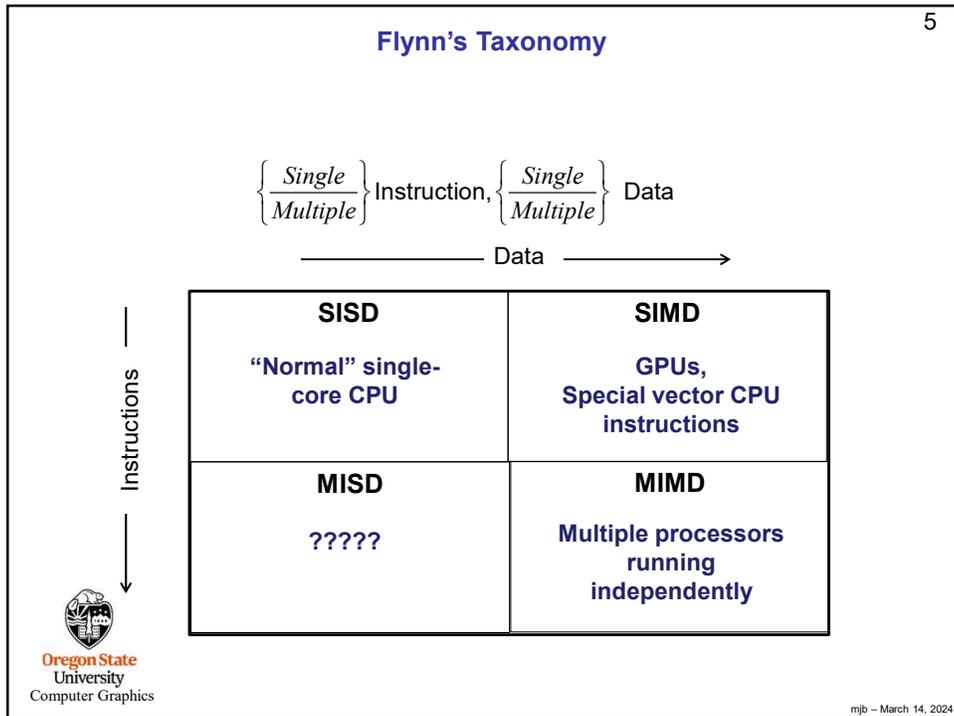


Thread execution switches when a thread blocks or uses up its time slice

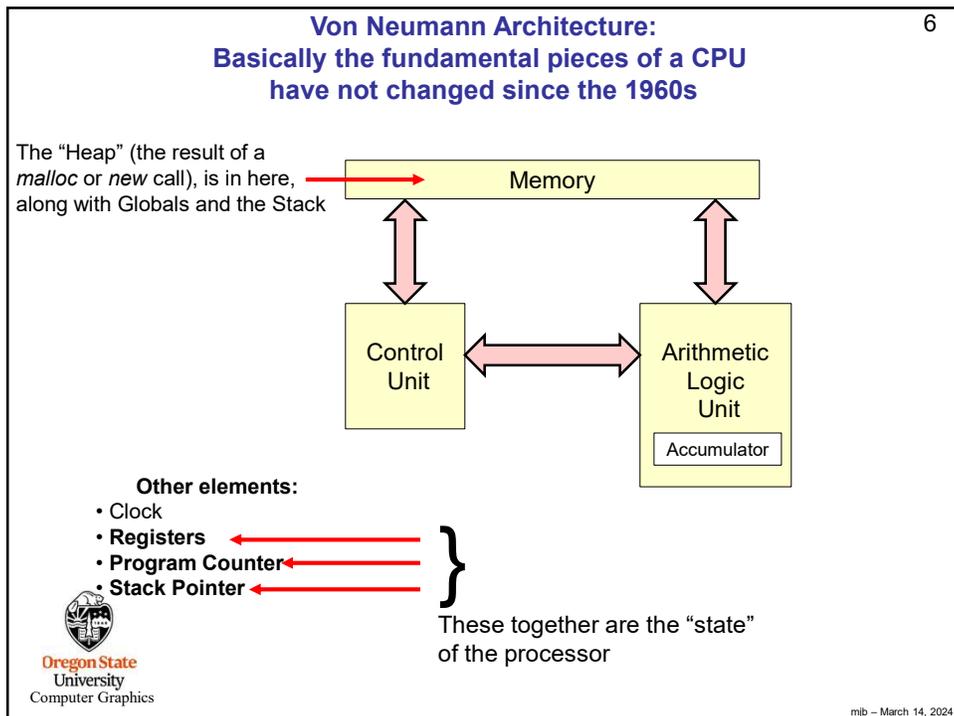


mjb - March 14, 2024

4



5

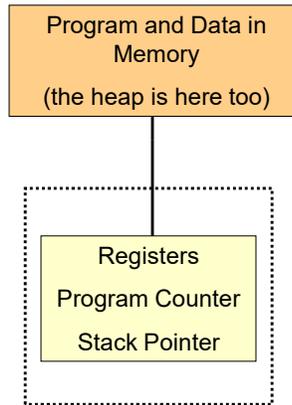


6

What Exactly is a Process?

7

Processes execute a program in memory. The process keeps a state (program counter, registers, and stack).



Other elements:

- Clock
- **Registers**
- **Program Counter**
- **Stack Pointer**



Oregon State
University
Computer Graphics

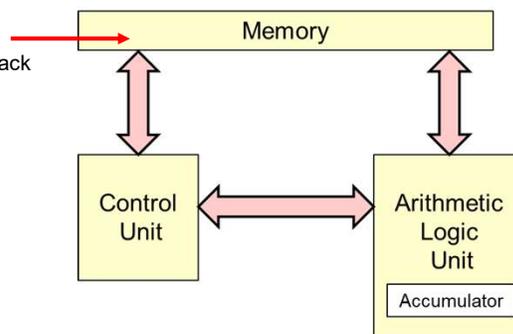
mjb - March 14, 2024

7

Von Neumann Architecture: Basically the fundamental pieces of a CPU have not changed since the 1960s

8

The "Heap" (the result of a *malloc* or *new* call), is in here, along with Globals and the Stack



Other elements:

- Clock
- **Registers** ←
- **Program Counter** ←
- **Stack Pointer** ←



Oregon State
University
Computer Graphics

What if we include more than one set of these?

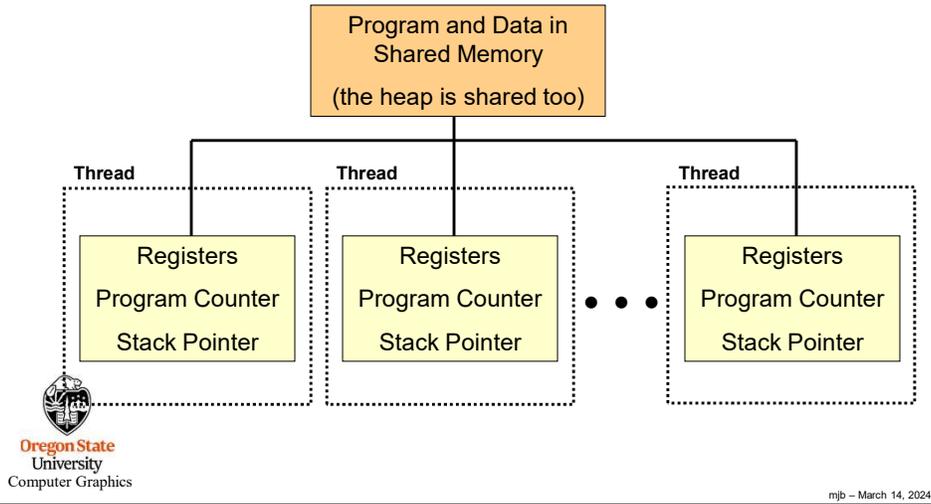
mjb - March 14, 2024

8

What Exactly is a Thread?

9

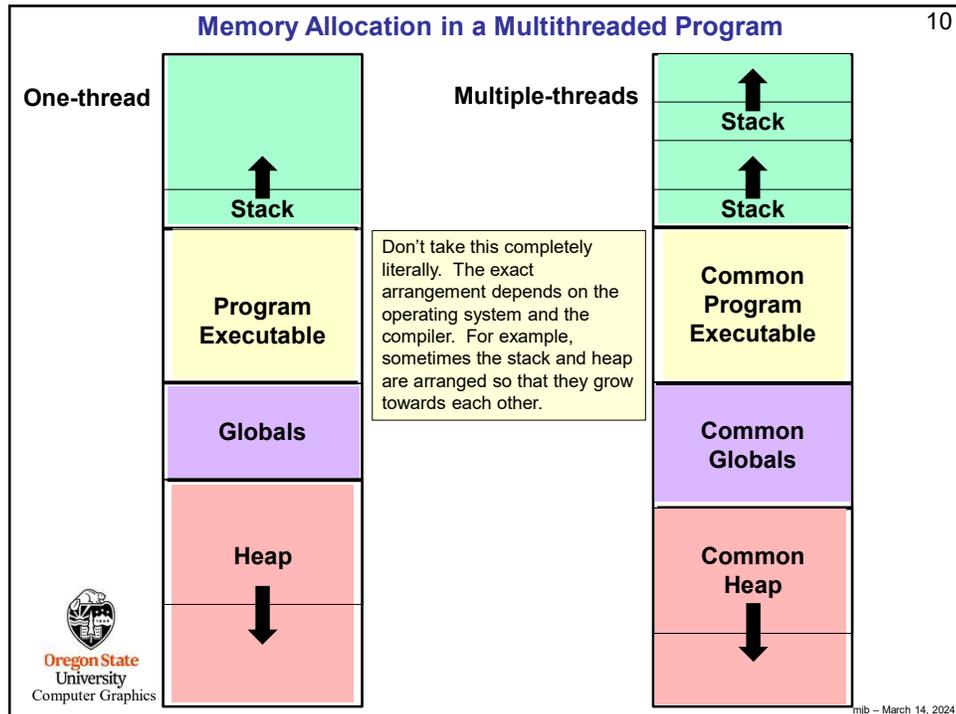
Threads are separate independent processes, all executing a common program and sharing memory. Each thread has its own state (program counter, registers, and stack pointer).



9

Memory Allocation in a Multithreaded Program

10



10

What Exactly is a Thread?

11

A “thread” is an independent path through the program code. Each thread has its own **Program Counter, Registers, and Stack Pointer**. But, since each thread is executing some part of the same program, each thread has access to the same global data in memory. Each thread is scheduled and swapped just like any other process.

Threads can share time on a single processor. You don't have to have multiple processors (although you can – the *multicore* topic is coming soon!).

This is useful, for example, in a web browser when you want several things to happen autonomously:

- User interface
- Communication with an external web server
- Web page display
- Image loading
- Animation



Oregon State
University
Computer Graphics



mjb – March 14, 2024

11

When is it Good to use Multithreading?

12

- When certain operations can become blocked, waiting for something else to happen
- When certain operations can be CPU-intensive
- When certain operations must respond to asynchronous I/O, including the user interface (UI)
- To manage independent behaviors in interactive simulations
- When you want to accelerate a single program on multicore CPU chips

Threads can make it easier to have many things going on in your program at one time and can absorb the dead-time of other threads.



Oregon State
University
Computer Graphics

mjb – March 14, 2024

12

Some Definitions

13

Atomic An operation that takes place to completion with no chance of being interrupted by another thread

Barrier A point in the program where *all* threads must reach before *any* of them are allowed to proceed

Coarse-grained parallelism Breaking a task up into a small number of large tasks

Deterministic The same set of inputs always gives the same outputs

Dynamic scheduling Dividing the total number of tasks T up so that each of N available threads has *less than* T/N sub-tasks to do, and then doling out the remaining tasks to threads as they become available

Fine-grained parallelism Breaking a task up into lots of small tasks



mjb - March 14, 2024

13

Some More Definitions

14

Private variable After a fork operation, a variable which has a private copy within each thread

Reduction Combining the results from multiple threads into a single sum or product, continuing to use multithreading. Typically, this is performed so that it takes $O(\log_2 N)$ time instead of $O(N)$ time:

Shared variable After a fork operation, a variable which is shared among threads, i.e., has a single value

Speed-up(N) T_1 / T_N
Speed-up Efficiency $\text{Speed-up}(N) / N$

Static Scheduling Dividing the total number of tasks T up so that each of N available threads has exactly T/N sub-tasks to do



mjb - March 14, 2024

14

Parallel Programming Tips



15

Tip #1 -- Don't Keep Internal State

```

int
GetLastPositiveNumber( int x )
{
    static int savedX; ← Internal state
    if( x >= 0 )
        savedX = x;
    return savedX;
}
  
```

If you do keep internal state between calls, there is a chance that a second thread will hop in and change it, then the first thread will use that state thinking it has not been changed.

Ironically, some of the standard C functions that we use all the time (e.g., *strtok*) keep internal state:

```
char * strtok ( char * str, const char * delims );
```

16

Tip #1 -- Don't Keep Internal State

17

Thread #1

```
char * tok1 = strtok( Line1, DELIMS ); 1
while( tok1 != NULL )
{
    ...
    tok1 = strtok( NULL, DELIMS ); 3
};
```

Execution Order

Thread #2

```
char * tok2 = strtok( Line2, DELIMS ); 2
while( tok2 != NULL )
{
    ...
    tok2 = strtok( NULL, DELIMS );
};
```

1. Thread #1 sets the internal character array pointer to somewhere in Line1[].
2. Thread #2 resets the same internal character array pointer to somewhere in Line2[].
3. Thread #1 uses that internal character array pointer, but it is not pointing into Line1[] where Thread #1 thinks it left it.



Oregon State
University
Computer Graphics

mjb - March 14, 2024

17

Tip #1 -- Keep External State Instead

18

Moral: if you will be multithreading, don't use internal static variables to retain state inside of functions.

In this case, using `strtok_r` is preferred:

```
char * strtok_r( char *str, const char *delims, char **sret );
```

`strtok_r` returns its internal state to you so that you can store it locally and then can pass it back when you are ready. (The 'r' stands for "re-entrant".)



Oregon State
University
Computer Graphics

mjb - March 14, 2024

18

Tip #1 -- Keep External State Instead

19

Thread #1

```
char *retValue1;
char * tok1 = strtok_r( Line1, DELIMS, &retValue1 );
while( tok1 != NULL )
{
    ...
    tok1 = strtok( NULL, DELIMS, &retValue1 );
};
```

Thread #2

```
char *retValue2;
char * tok2 = strtok_r( Line2, DELIMS, &retValue2 );
while( tok2 != NULL )
{
    ...
    tok2 = strtok( NULL, DELIMS, &retValue2 );
};
```

Now, execution order no longer matters!



mjb - March 14, 2024

19

Tip #1 – Note that Keeping *Global* State is Just as Dangerous

20

Internal state:

```
int
GetLastPositiveNumber( int x )
{
    static int savedX;
    if( x >= 0 )
        savedX = x;
    return savedX;
}
```

Global state:

```
int savedX;
int
GetLastPositiveNumber( int x )
{
    if( x >= 0 )
        savedX = x;
    return savedX;
}
```



mjb - March 14, 2024

20

Tip #2 – Avoid Deadlock

21

Deadlock is when two threads are each waiting for the other to do something

Worst of all, the way these problems occur is not usually deterministic!



mjb – March 14, 2024

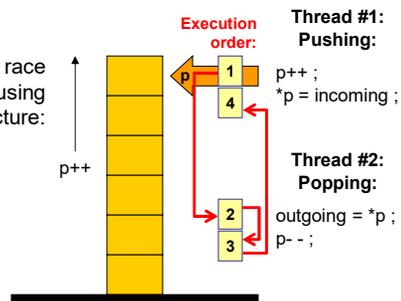
21

Tip #3 – Avoid Race Conditions

22

- A Race Condition is where it matters which thread gets to a particular piece of code first.
- This often comes about when one thread is modifying a variable while the other thread is in the midst of using it

A good example of a potential race condition situation is maintaining and using the pointer in a stack data structure:



Worst of all, the way these problems occur is not usually deterministic!

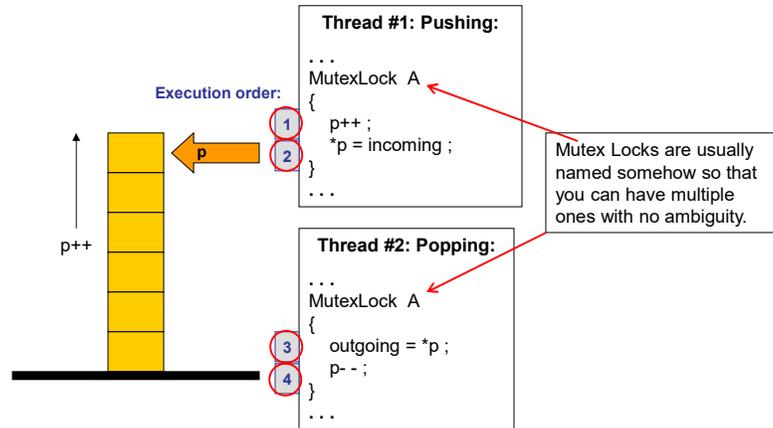


mjb – March 14, 2024

22

BTW, Race Conditions can often be fixed through the use of Mutual Exclusion Locks (Mutexes)

23



We will talk about these a little later.
But note that, while solving a race condition, we can accidentally create a deadlock condition if the thread that owns the lock is waiting for the other thread to do something

Computer Graphics

mjb - March 14, 2024

23

Tip #4 -- Sending a Message to the Optimizer: The *volatile* Keyword

24

The *volatile* keyword is used to let the compiler know that another thread might be changing a variable "in the background", so don't make any assumptions about what can be optimized away.

```
int val = 0;
...
while( val != 0 );
```

A good compiler optimizer will *eliminate* this code because it "knows" that, for all time, $val == 0$

```
volatile int val = 0;
...
while( val != 0 );
```

The *volatile* keyword tells the compiler optimizer that it cannot count on val being $== 0$ here

Oregon State
University
Computer Graphics

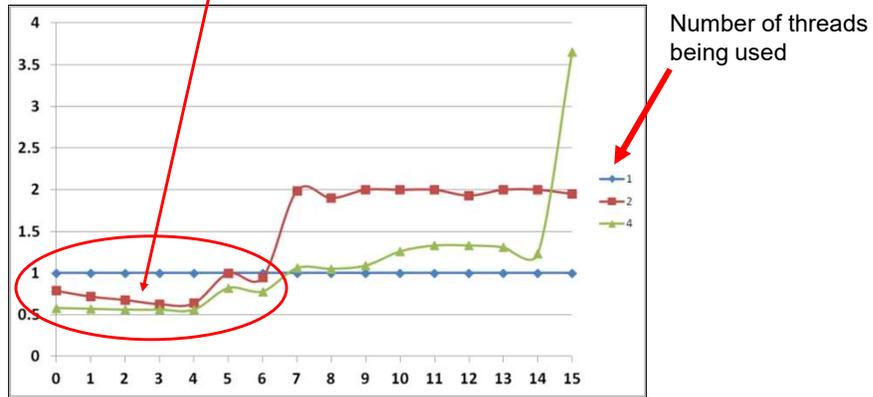
mjb - March 14, 2024

24

Tip #5 – Beware of False Sharing Caching Issues

25

Note that using more threads initially gives a *drop* in performance!



We will get to this in the Caching notes!

