



Parallel Computing

CSCI 201

Principles of Software Development

Jeffrey Miller, Ph.D.
jeffrey.miller@usc.edu



Outline

- Parallel Computing
- Program

Parallel Computing





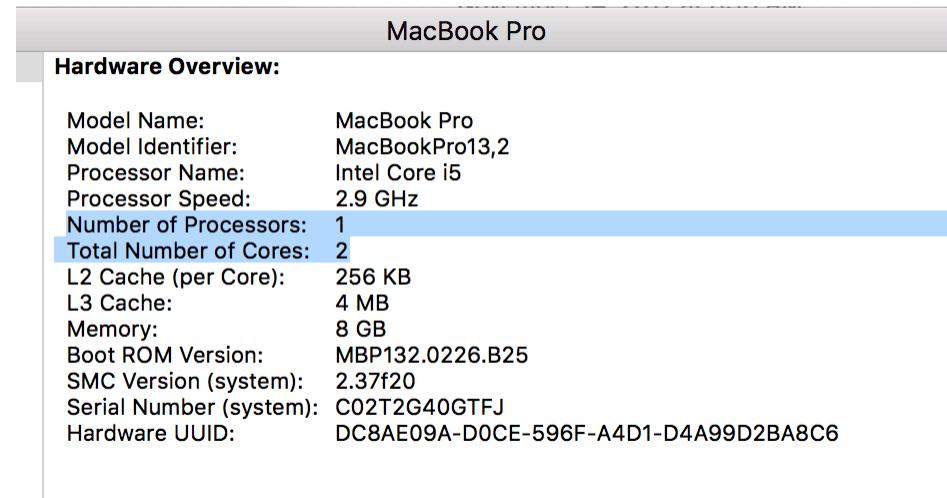
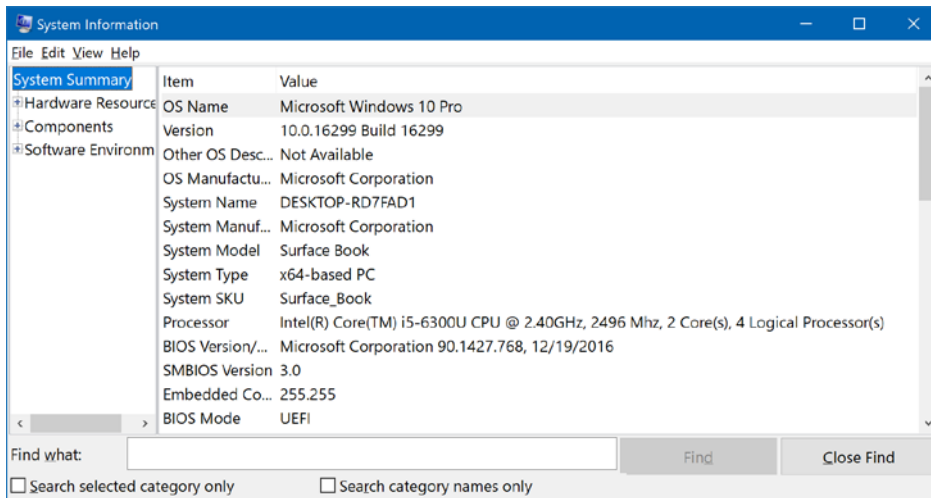
- **Parallel computing** studies software systems where components located on connected components communicate through message passing
 - › Individual threads have only a partial knowledge of the problem
 - › Parallel computing is a term used for programs that operate within a shared memory space with multiple processors or cores



Redundant Hardware Determination




- To determine what CPU or CPUs you have in your computer
 - › On Windows, Run ( + R) and type `msinfo32`
 - › On Mac, go to  ->About this Mac -> System Report



Fork/Join Framework



- The fork/join framework is an implementation of the `ExecutorService` interface that allows taking advantage of multiple cores or multiple processors
- It is designed for work that can be broken into smaller pieces recursively, allowing all available processing power
- We can create threads and add them to a thread pool
- The `ForkJoinPool` class implements the core algorithm for parallel computing in Java and can execute `ForkJoinTask` processes



Fork/Join Framework Structure



- Pseudocode for the fork/join framework is:

```
if (condition)
    do the work directly
else
    split work into multiple pieces
    invoke the pieces and wait for results
```
- The above code should be inside a child of `ForkJoinTask`
 - › `RecursiveTask` can return a result
 - › `RecursiveAction` does not return a result
- Create the object that represents the work to be done and pass it to the `invoke(...)` or `execute(...)` method of a `ForkJoinPool`



ForkJoinTask Class



- The `ForkJoinTask` class has a method `compute()` that should be overridden
- The `compute()` method will contain the main computation performed by the task
- The `fork()` method allows asynchronous execution (starts the thread and calls `compute()`)
- The `join()` method will not proceed until the task's `compute()` method has completed as a result of the call to `fork()`
- The `invoke()` method on a `ForkJoinPool` will execute the `compute()` method asynchronously and return the value (through a `fork()` then a `join()` call)
- The `execute()` method on a `ForkJoinPool` will execute the `compute()` method asynchronously but will not return any value (through a `fork()` call with no `join()`)

Multi-Threading vs Parallel Comparison



	Multi-Threading	Parallel
Class/Interface from which to inherit	<code>Thread/Runnable</code>	<code>RecursiveAction</code> or <code>RecursiveTask</code>
Method to override	<code>void run()</code>	<code>protected abstract void compute()</code> <i>(RecursiveAction)</i> <hr/> <code>protected abstract V compute()</code> <i>(RecursiveTask)</i>
Method called to start thread	<code>public void start()</code>	<code>public final ForkJoinTask<V> fork()</code>
Managing class	<code>ExecutorService</code>	<code>ForkJoinPool</code>
Method to call on managing class to start thread	<code>void execute(Runnable)</code>	<code>public void execute(ForkJoinTask<?>)</code> <i>(RecursiveAction)</i> <hr/> <code>public T invoke(ForkJoinTask<T>)</code> <i>(RecursiveTask)</i>
Method to stop current thread from continuing until another thread completes	<code>public final void join()</code>	<code>public final V join()</code>

Running Time



- Any time you fork a task, there is overhead in moving that to a new CPU, executing it, and getting the response
- Just because you are parallelizing code *does not mean* that you will have an improvement in execution speed
- If you fork more threads than you have CPUs, the threads will execute in a concurrent manner (time-slicing similar to multi-threading) in each CPU



Not Parallel Sum



```
1 public class SumNoParallel {
2     public static void main(String [] args) {
3         long maxNumber = 1_000_000_000;
4         long before = System.currentTimeMillis();
5         Sum sum = new Sum(maxNumber);
6         long num = sum.compute();
7         long after = System.currentTimeMillis();
8         System.out.println("time without parallelism = " + (after - before));
9         System.out.println("SUM(0.." + maxNumber + ") = " + num);
10    }
11 }
12 class Sum {
13     private long maxNum;
14     Sum(long maxNum) {
15         this.maxNum = maxNum;
16     }
17     protected Long compute() {
18         long sum = 0;
19         for (int i=0; i <= maxNum; i++) {
20             sum += i;
21         }
22         return sum;
23     }
24 }
```

```
time without parallelism = 902
SUM(0..1000000000) = 500000000500000000
```

```
time without parallelism = 928
SUM(0..1000000000) = 500000000500000000
```

```
time without parallelism = 812
SUM(0..1000000000) = 500000000500000000
```

Parallel Sum



```
1 import java.util.concurrent.ForkJoinPool;
2 import java.util.concurrent.RecursiveTask;
3
4 public class SumParallel {
5     public static void main(String [] args) {
6         long minNumber = 0;
7         long maxNumber = 1_000_000_000;
8         long numElements = 1000;
9         int numThreads = (int)((maxNumber - minNumber) / numElements);
10        long before = System.currentTimeMillis();
11
12        ForkJoinPool pool = new ForkJoinPool();
13        SumP sum[] = new SumP[numThreads];
14        long start = minNumber;
15        long end = numElements;
16        for (int i=0; i < numThreads; i++) {
17            sum[i] = new SumP(start, end);
18            start = end + 1;
19            end = start + numElements - 1;
20            pool.execute(sum[i]); // no return value, so we will join later
21        }
22        long num = 0;
23        for (int i=0; i < numThreads; i++) {
24            num += sum[i].join();
25        }
26        long after = System.currentTimeMillis();
27        System.out.println("time with parallelism = " + (after-before));
28        System.out.print("SUM(" + minNumber + ".." + maxNumber + ") = ");
29        System.out.println(num);
30 }
```

```
31 class SumP extends RecursiveTask<Long> {
32     private long minNum;
33     private long maxNum;
34     public static final long serialVersionUID = 1;
35     public SumP(long minNum, long maxNum) {
36         this.minNum = minNum;
37         this.maxNum = maxNum;
38     }
39     protected Long compute() {
40         long sum = 0;
41         for (long i=minNum; i <= maxNum; i++) {
42             sum += i;
43         }
44         return sum;
45     }
46 }
```

```
time with parallelism = 457
SUM(0..1000000000) = 500000000500000000
```

```
time with parallelism = 436
SUM(0..1000000000) = 500000000500000000
```

```
time with parallelism = 433
SUM(0..1000000000) = 500000000500000000
```

Not Parallel Mergesort



```
1 public class NonParallelMergeSort {
2     public static void main(String[] args) {
3         int SIZE = 2_000_000;
4         int[] list = new int[SIZE];
5         for (int i = 0; i < SIZE; i++) {
6             list[i] = (int)(Math.random() * Integer.MAX_VALUE);
7         }
8         long timing = 0;
9         long sum = 0;
10        // run it 8 times to see if there are variations
11        for (int i=0; i < 8; i++) {
12            timing = nonParallelMergeSort((int[])list.clone());
13            System.out.println(timing + " ms");
14            sum += timing;
15        }
16        System.out.println("average = " + (sum / 8) + " ms");
17    }
18
19    public static long nonParallelMergeSort(int[] list) {
20        long before = System.currentTimeMillis();
21        new SortTask(list).compute();
22        long after = System.currentTimeMillis();
23        long time = after - before;
24        return time;
25    }
26    private static class SortTask {
27        private int[] list;
28        SortTask(int[] list) {
29            this.list = list;
30        }
31    }
32 }
```

```
31 protected void compute() {
32     if (list.length < 2) return; // base case
33     // split into halves
34     int[] firstHalf = new int[list.length / 2];
35     System.arraycopy(list, 0, firstHalf, 0, list.length / 2);
36     int secondLength = list.length - list.length / 2;
37     int[] secondHalf = new int[secondLength];
38     System.arraycopy(list, list.length / 2, secondHalf, 0, secondLength);
39
40     // recursively sort the two halves
41     new SortTask(firstHalf).compute();
42     new SortTask(secondHalf).compute();
43     // merge halves together
44     merge(firstHalf, secondHalf, list);
45 }
46 }
47
48 public static void merge(int[] list1, int[] list2, int[] merged) {
49     int i1 = 0, i2 = 0, i3 = 0; // index in list1, list2, out
50     while (i1 < list1.length && i2 < list2.length) {
51         merged[i3++] = (list1[i1] < list2[i2]) ? list1[i1++] : list2[i2++];
52     }
53     // any trailing ends
54     while (i1 < list1.length) {
55         merged[i3++] = list1[i1++];
56     }
57     while (i2 < list2.length) {
58         merged[i3++] = list2[i2++];
59     }
60 }
61 }
```

```
458 ms
416 ms
412 ms
451 ms
401 ms
378 ms
417 ms
397 ms
average = 416 ms
```

Parallel Merge Sort Example



```
1 import java.util.concurrent.ForkJoinPool;
2 import java.util.concurrent.RecursiveAction;
3 public class ParallelMergeSort {
4     public static void main(String[] args) {
5         int SIZE = 2_000_000;
6         int[] list = new int[SIZE];
7
8         for (int i = 0; i < SIZE; i++) {
9             list[i] = (int)(Math.random() * Integer.MAX_VALUE);
10        }
11        int numProcessors = Runtime.getRuntime().availableProcessors();
12        System.out.println("num processors: " + numProcessors);
13
14        // timing[i] : time to sort on i processors
15        long[] timing = new long[numProcessors*2+1];
16
17        for (int i=1; i <= numProcessors * 2; i++) {
18            timing[i] = parallelMergeSort((int[])list.clone(), i);
19            System.out.println(i + " processors=" + timing[i] + " ms");
20        }
21    }
22
23    public static long parallelMergeSort(int[] list, int proc) {
24        long before = System.currentTimeMillis();
25        ForkJoinPool pool = new ForkJoinPool(proc);
26        pool.invoke(new SortTask(list));
27        pool.shutdown();
28        while (!pool.isTerminated()) {
29            Thread.yield();
30        }
31        long after = System.currentTimeMillis();
32        long time = after - before;
33        return time;
34    }
35
36    private static class SortTask extends RecursiveAction {
37        public static final long serialVersionUID = 1;
38        private int[] list;
39        SortTask(int[] list) {
40            this.list = list;
41        }
42
43        protected void compute() {
44            if (list.length < 2) return; // base case
45            // split into halves
46            int[] firstHalf = new int[list.length / 2];
47            System.arraycopy(list, 0, firstHalf, 0, list.length / 2);
48            int secondLength = list.length - list.length / 2;
49            int[] secondHalf = new int[secondLength];
50            System.arraycopy(list, list.length / 2, secondHalf, 0, secondLength);
51
52            // recursively sort the two halves
53            SortTask st1 = new SortTask(firstHalf);
54            SortTask st2 = new SortTask(secondHalf);
55            st1.fork();
56            st2.fork();
57            st1.join();
58            st2.join();
59        }
60
61        public static void merge(int[] list1, int[] list2, int[] merged) {
62            int i1 = 0, i2 = 0, i3 = 0; // index in list1, list2, out
63            while (i1 < list1.length && i2 < list2.length) {
64                merged[i3++] = (list1[i1] < list2[i2]) ? list1[i1++] : list2[i2++];
65            }
66            // any trailing ends
67            while (i1 < list1.length) {
68                merged[i3++] = list1[i1++];
69            }
70            while (i2 < list2.length) {
71                merged[i3++] = list2[i2++];
72            }
73        }
74    }

```

```
num processors: 4
1 processors=902 ms
2 processors=704 ms
3 processors=269 ms
4 processors=257 ms
5 processors=214 ms
6 processors=240 ms
7 processors=249 ms
8 processors=236 ms
```

```
num processors: 4
1 processors=912 ms
2 processors=732 ms
3 processors=273 ms
4 processors=255 ms
5 processors=228 ms
6 processors=224 ms
7 processors=261 ms
8 processors=235 ms
```



Outline

- Parallel Computing
- Program

Program



- Modify the `ParallelMergeSort` code to split the array into the same number of sub-arrays as processors/cores on your computer. Does that provide a lower execution time than splitting the array into two sub-arrays? Why or why not?