

Parallel Programming

Gnkggo, Informatik B. Sc. 2. Semester

Contents

1	Creating a Thread	2
2	Joining threads	2
3	Thread states	3
4	Data Races	3
5	Important to Remember	3
5.1	Speedup	3
5.2	Amdahl	4
5.3	Gustafson	4
5.4	Workspan	4
5.5	Divide and Conquer with ExecutorService	4
6	Using Recursive Task in Java	5
7	Throughput	5
8	Latency	5
9	MPI	6
9.1	Synchronous, asynchronous, blocking, non-blocking	6
10	Transactional Memory	6
11	Important Code	6
11.1	Barrier	6
11.2	Semaphore	7
11.3	PetersonLock	8
11.4	Filterlock	8
11.5	BakeryLock	8
11.6	TAS	9
11.7	CAS	9
11.8	TASLock	10
11.9	TATASLock	10

1 Creating a Thread

```
1 public class Useless extends Thread {
2     int i;
3     Useless(int i) {
4         this.i = i;
5     }
6     public void run() {
7         System.out.println("Thread says hi" + i);
8         System.out.println("Thread says bye" + i);
9     }
10 }
11
12 public class M {
13     public static void main(String[] args) {
14         for (int i = 0; i < 20; i++) {
15             Thread t = new Useless(i + 1);
16             t.start(); // Important: you cannot use t.run() --> t.run doesn't create a new
17                       // Thread
18         }
19     }
20 }
```

Important: You can only create a Thread when you use `start()`. The `run` method has no return value and no argument.

2 Joining threads

```
1 public class Useless extends Thread {
2     int i;
3     Useless(int i) {
4         this.i = i;
5     }
6     public void run() {
7         System.out.println("Thread says hi" + i);
8         System.out.println("Thread says bye" + i);
9     }
10 }
11
12 public class M {
13     public static void main(String[] args) {
14         Thread[] threads = new Thread[20];
15         for (int i = 0; i < 20; i++) {
16             Thread t = new Useless(i + 1);
17             t.start();
18             threads[i] = t;
19         }
20         for (int i = 0; i < 20; i++) {
21             try { // need catch block around join
22                 threads[i].join();
23             } catch (InterruptedException e) {
24                 // Some catch block
25             }
26         }
27         System.out.println("All done.");
28     }
29 }
```

With "join," you wait until every thread is done.

3 Thread states

If we want to be able to talk about the effects of different thread operations, we need some notion of thread states. In short, a Java thread typically goes through the following states:

- **Non-Existing:** Before the thread is created, this is where it is. We don't know too much about this place, as it's not actually on our plane of reality, but it's somewhere out there.
- **New:** Once the Thread object is created, the thread enters the new state.
- **Runnable:** Once we call `start()` on the new thread object, it becomes eligible for execution, and the system can start scheduling the thread as it wishes.
- **Blocked:** When the thread attempts to acquire a lock, it goes into a blocked state until it has obtained the lock, upon which it returns to a runnable state. In addition, calling the `join()` method will also transfer a thread into a blocked state.
- **Waiting:** The thread can call `wait()` to go into a waiting state. It'll return to a runnable state once another thread calls `notify()` or `notifyAll()` and the thread is removed from the waiting queue.
- **Terminated:** At any point during execution, we can use `interrupt()` to signal the thread to stop its execution. It will then transfer to a terminated state. Note that when the thread is in a runnable state, it needs to check whether its interrupted flag is set itself; it won't transfer to the terminated state automatically. Of course, exiting the run method is equivalent to entering a terminated state. Once the garbage collector realizes that the thread has been terminated and is no longer reachable, it will garbage collect the thread and return it to a non-existing state, completing the cycle.

4 Data Races

A data race is a specific kind of race condition that is better described as a simultaneous access error, although nobody uses that term. There are two kinds of data races:

- When one thread might read an object field at the same moment that another thread writes the same field.
- When one thread might write an object field at the same moment that another thread also writes the same field.

```
1 class C {
2     private int x = 0;
3     private int y = 0;
4
5     void f() {
6         x = 1; //line A
7         y = 1; //line B
8     }
9     void g() {
10        int a = y; //line C
11        int b = x; //line D
12        assert (b >= a);
13    }
14 }
```

The code has data races, but it doesn't occur, which can be proven by contradiction.

5 Important to Remember

5.1 Speedup

$$S_p := \frac{T_1}{T_P}$$

Where T_1 is the sequential time (Time with one processor) and T_P with the time with P processors.

Reasons why the program is nevertheless slower:

Additional overheads caused by inter-thread dependencies, creating threads, communicating between them, and memory-hierarchy issues can greatly limit the speedup we gain from adding more processors.

5.2 Amdahl

- Fixed workload and upper bound on the speedup achievable when increasing the number of processors at our disposal.
- Let f denote the non-parallelizable, serial fraction of the total work done in a program, and P the number of processors at our disposal. Then, the following inequality holds:

$$S_P \leq \frac{1}{f + \frac{1-f}{P}}$$

- If P is infinity, then:

$$S_\infty \leq \frac{1}{f}$$

How to derive it:

$$T = T_s + T_p$$

Where T is the total time, T_s is the sequential time, and T_p is the parallel time. $T_p = \frac{T_s}{p}$ If you have more than one processor, you can rewrite the function to:

$$T = T_s + \frac{T_p}{p}$$

$$S_p = \frac{W_{\text{seq}} + W_{\text{par}}}{W_{\text{seq}} + \frac{W_{\text{par}}}{p}}$$

We know that $W_{\text{seq}} + W_{\text{par}} = 1$. So we can rewrite the function to:

$$\frac{1}{W_{\text{seq}} + \frac{W_{\text{par}}}{p}} = \frac{1}{f + \frac{1-f}{p}}$$

5.3 Gustafson

- We increase the problem size as we improve the resources at our disposal. We consider the time interval to be fixed and look at the problem size.
- Let f denote the non-parallelizable, serial fraction of the total work done in the program, and P the number of processors at our disposal. Then, we get:

$$S_P = f + P(1 - f) = P - f(P - 1)$$

5.4 Workspan

$$T_x \leq T_\infty + \frac{T_1 - T_\infty}{x}$$

In a graph:

- Work: all jobs summed up
- Span: Longest critical path

5.5 Divide and Conquer with ExecutorService

```

“Java class MaxTask implements Callable<int> {
    int l; int h; int[] arr; ExecutorService ex;
    public MaxTask(ExecutorService ex, int lo, int hi, int[] arr) // ... public Integer call() throws Exception //
    Check base case int size = h - l; if (size == 1) return arr[l]; // Split work int mid = size / 2; MaxTask m1 = new
    MaxTask(ex, l, l + mid, arr); MaxTask m2 = new MaxTask(ex, l + mid, h, arr); // Start subtasks Future<Integer> f1
    = ex.submit(m1); Future<Integer> f2 = ex.submit(m2); // Combine results try return Math.max(f1.get(), f2.get());
    catch (Exception e) return 0;
    public static void main(String[] args) int[] arr = new int[] {15, 7, 9, 8, 4, 22, 42, 13}; ExecutorService ex =
    Executors.newFixedThreadPool(8); // Attention, has to be a minimum number of threads that are needed -j otherwise
    you have an endless loop MaxTask top = new MaxTask(ex, 0, arr.length, arr); Future<Integer> max = ex.submit(top);
    try System.out.println(max.get()); catch (Exception e) // something ex.shutdown();
}

```

6 Using Recursive Task in Java

To avoid "knowing" how many threads you need, you can use Recursive Task:

```

1 class MaxForkJoin extends RecursiveTask<Integer> {
2     int l;
3     int h;
4     int[] arr;
5
6     public MaxForkJoin(int lo, int hi, int[] arr) {
7         this...
8     }
9
10    public Integer compute() {
11        //Check base case
12        int size = h - l;
13        if (size == 1) {
14            return arr[l];
15        } //split work
16        int mid = size / 2;
17        MaxForkJoin m1 = new MaxForkJoin(l, l + mid, arr);
18        MaxForkJoin m2 = new MaxForkJoin(l + mid, h, arr);
19        //Run subtasks
20        m1.fork();
21        int max2 = m2.compute();
22        int max1 = m1.join();
23        //Combine results
24        return Math.max(m1.get(), m2.get());
25    }
26 }
27
28 public static void main(String[] args) {
29     int[] arr = new int[] {15, 7, 9, 8, 4, 22, 42, 13};
30     MaxForkJoin top = new MaxTask(0, arr.length, arr);
31     ForkJoinPool fjp = new ForkJoinPool();
32     int res = fjp.invoke(top);
33     System.out.println(res);
34 }

```

Note the following similarities:

- Instead of extending Thread, we extend RecursiveTask<T> (with return value) or RecursiveAction (without return value).
- Instead of overriding run, we override compute.
- Instead of calling start, we call fork.
- Instead of a topmost call to run, we create a ForkJoinPool and call invoke.

Also, note that in the case of RecursiveTask<T>, join now returns a result.

7 Throughput

$$\text{Throughput} \approx \frac{1}{\max(\text{computationtime}(\text{stages}))} \quad (1)$$

8 Latency

Time to perform a single computation, including wait time resulting from resource dependencies. A pipeline is **balanced** if the latency remains constant over time.

9 MPI

9.1 Synchronous, asynchronous, blocking, non-blocking

- Synchronous + blocking: try to call somebody until he answers.
- Synchronous + non-blocking: try to call, if the other person does not pick up, I do something else.
- Asynchronous + blocking: wait until your crush texts you back.
- Asynchronous + non-blocking: send an E-Mail and continue working until you get a response.

In the actor model, messages are sent in an **asynchronous, non-blocking fashion**. The sender places the message into the buffer of the receiver and continues execution. In contrast, when the sender sends **synchronous messages**, it blocks until the message has been received.

MPI collects processes into groups, where each group can have multiple **colors**. A group paired with its color uniquely identifies a communicator. Initially, all processes are collected in the same group and communicator MPI_COMM_WORLD. Within each communicator, a process is assigned a unique identifier, called the **rank**.

```
1 public void Send(  
2     Object buf, //Ptr to data to be sent  
3     int offset,  
4     int count, //number of items to be sent  
5     Datatype datatype, //datatype of items  
6     int dest, //destination process id  
7     int tag //data id tag  
8 );  
9  
10 public void Recv(  
11     Object buf,  
12     int offset,  
13     int count, // Number of items to be received  
14     Datatype datatype, // Datatype of items  
15     int dest, // Source process id  
16     int tag //Data id tag  
17 );
```

10 Transactional Memory

Definition 3.7.1 (Transactional Memory): Transactional Memory is a programming model whereby loads and stores on a particular thread can be grouped into transactions. The read set and write set of a transaction are the set of addresses read from and written to, respectively, during the transaction. A data conflict occurs in a transaction if another processor reads or writes a value from the transaction's write set or writes to an address in the transaction's read set. Data conflicts cause the transaction to abort, and all instructions executed since the start of the transaction (and all changes to the write set) to be discarded.

Transactions run in *isolation*: while a transaction is running, effects from other transactions are not observed. A good analogy is the one of a snapshot: transactional memory works as if a transaction takes a snapshot of the global state when it begins and then operates on that snapshot.

11 Important Code

11.1 Barrier

```
1 public class Barrier {  
2     private Semaphore mutex;  
3     private Semaphore barrier1;  
4     private Semaphore barrier2;  
5 }
```

```
6 private volatile int count = 0;
7 private final int n;
8
9 Barrier(int n) {
10     mutex = new Semaphore(1);
11     barrier1 = new Semaphore(0);
12     barrier2 = new Semaphore(1);
13
14     this.count = 0;
15     this.n = n;
16 }
17
18 void await() throws InterruptedException {
19     mutex.acquire();
20     ++count;
21     if (count == n) {
22         barrier2.acquire();
23         barrier1.release();
24     }
25     mutex.release();
26
27     barrier1.acquire();
28     barrier1.release();
29
30     mutex.acquire();
31     --count;
32     if (count == 0) {
33         barrier1.acquire();
34         barrier2.release();
35     }
36     mutex.release();
37
38     barrier2.acquire();
39     barrier2.release();
40 }
41 }
```

11.2 Semaphore

```
1 public class Semaphore {
2     private volatile int count;
3     private Object monitor = new Object();
4
5     public Semaphore(int count) {
6         this.count = count;
7     }
8
9     public void acquire() throws InterruptedException {
10         synchronized(monitor) {
11             while (count <= 0)
12                 monitor.wait();
13             --count;
14         }
15     }
16
17     public void release() {
```



```
18     synchronized(monitor) {
19         ++count;
20         monitor.notify();
21     }
22 }
23 }
```

11.3 PetersonLock

```
1 class PetersonLock {
2     volatile boolean flag[] = new boolean[2];
3     volatile int victim;
4
5     public void Acquire(int id) {
6         flag[id] = true;
7         victim = id;
8         while (flag[1-id] && victim == id);
9     }
10
11     public void Release(int id) {
12         flag[id] = false;
13     }
14 }
```

11.4 Filterlock

```
1 int[] level(#threads), int[] victim(#threads)
2
3 lock(me) {
4     for (int i=1; i<n; ++i) {
5         level[me] = i;
6         victim[i] = me;
7         while (exists(k != me): level[k] >= i && victim[i] == m){};
8     }
9 }
10
11 unlock(me) {
12     level[me] = 0;
13 }
```

Filterlock is not fair.

11.5 BakeryLock

```
1 class BakeryLock {
2     AtomicIntegerArray flag;
3     AtomicIntegerArray label;
4     final int n;
5
6     BakeryLock(int n) {
7         this.n = n;
8         flag = new AtomicIntegerArray(n);
9         label = new AtomicIntegerArray(n);
10    }
11
12    int MaxLabel() {
```

```
13     int max = label.get(0);
14     for (int i = 1; i < n; ++i)
15         max = Math.max(max, label.get(i));
16     return max;
17 }
18
19 boolean Conflict(int me) {
20     for (int i = 0; i < n; ++i)
21         if (i != me && flag.get(i) != 0) {
22             int diff = label.get(i) - label.get(me);
23             if (diff < 0 || diff == 0 && i < me)
24                 return true;
25         }
26     return false;
27 }
28
29 public void Acquire(int me) {
30     flag.set(me, 1);
31     label.set(me, MaxLabel() + 1);
32     while (Conflict(me));
33 }
34
35 public void Release(int me) {
36     flag.set(me, 0);
37 }
38 }
```

11.6 TAS

```
1 boolean TAS(memref s) {
2     if (mem[s] == 0) {
3         mem[s] = 1;
4         return true;
5     } else {
6         return false;
7     }
8 }
9
10 Init(lock) {
11     lock = 0;
12     Acquire(lock) {
13         while (!TAS(lock)); //wait
14     }
15     Release(lock) {
16         lock = 0;
17     }
18 }
```

11.7 CAS

```
1 int CAS(memref a, int old, int newValue) {
2     oldVal = mem[a];
3     if (old == oldVal) {
4         mem[a] = newValue
5     }
6 }
```

```
6     return oldVal;
7 }
8
9 Init (lock) {
10     lock = 0;
11     Acquire (lock) {
12         while (CAS(lock, 0, 1) != 0); //wait
13     }
14     Release(lock) {
15         CAS(lock, 1, 0); //ignore result
16     }
17 }
```

11.8 TASLock

```
1 public class TASLock implements Lock {
2     AtomicBoolean state = new AtomicBoolean(false);
3
4     public void lock() {
5         while (state.getAndSet(true)){}
6     }
7
8     public void unlock() {
9         state.set(false);
10    }
11 }
```

11.9 TATASLock

```
1 public class TATASLock implements Lock {
2     AtomicBoolean state = new AtomicBoolean(false);
3
4     public void lock() {
5         do {
6             while(state.get()) {}
7         } while (!state.compareAndSet(false, true))
8     }
9 }
```