

# Parallele Programmierung FS25

Exercise Session 5

Jonas Wetzel

# Plan für heute

- Organisation
- Nachbesprechung Exercise 4
- Theory Recap
- Intro Exercise 5
- Exam Questions
- Kahoot

# Organisation

- Mein Name ist Jonas Wetzel
- Meine Website (Materialien und Inhalt der Übungen):  
[n.ethz.ch/~jwetzel](http://n.ethz.ch/~jwetzel)
- Meine Email: [jwetzel@ethz.ch](mailto:jwetzel@ethz.ch)
- Discord: @jonas.too

# Organisation

- Mein Name ist Jonas Wetzel
- Meine Website (Materialien und Inhalt der Übungen):  
n.ethz.ch/~jwetzel
- Meine Email: [jwetzel@ethz.ch](mailto:jwetzel@ethz.ch)
- Discord: @jonas.too
- Feedback zur Session: <https://forms.gle/qiDnqkfSP2NUQGvc9>

# Organisation

- Feedback zur Session: <https://forms.gle/qiDnqkfSP2NUQGvc9>
- Falls ihr Feedback möchten kommt bitte zu mir

# Organisation

- Wo sind wir jetzt?

Date	Title
Feb 17	Introduction & Course Overview
Feb 18	Java Recap and JVM Overview
Feb 24	Introduction to Threads and Synchronization (Part I)
Feb 25	Introduction to Threads and Synchronization (Part II)
Mar 3	Introduction to Threads and Synchronization (Part III)
Mar 4	Parallel Architectures: Parallelism on the Hardware Level
Mar 10	Basic Concepts in Parallelism
Mar 11	Divide & Conquer and Executor Service
Mar 17	DAG and ForkJoin Framework
Mar 18	Parallel Algorithms (Part I)
Mar 24	Parallel Algorithms (Part II)
Mar 25	Shared Memory Concurrency, Locks and Data Races
Mar 31	Virtual Threads
Apr 01	Exam Preparation (First Half)



# Plan für heute

- Organisation
- **Nachbesprechung Exercise 4**
- Theory Recap
- Intro Exercise 5
- Exam Questions
- Kahoot

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

a) total time if strictly sequential order?

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

a) total time if strictly sequential order?



$$50 + 90 + 15$$

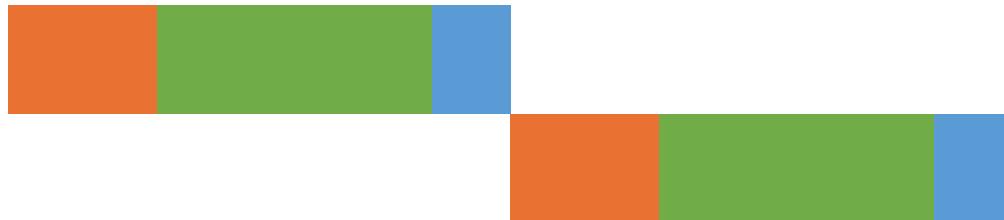


$$155$$

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

a) total time if strictly sequential order?



155      +      155

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

a) total time if strictly sequential order?



$$155 + 155 + 155 + 155 = 620$$

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

b) what would be a better (faster) strategy?



# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

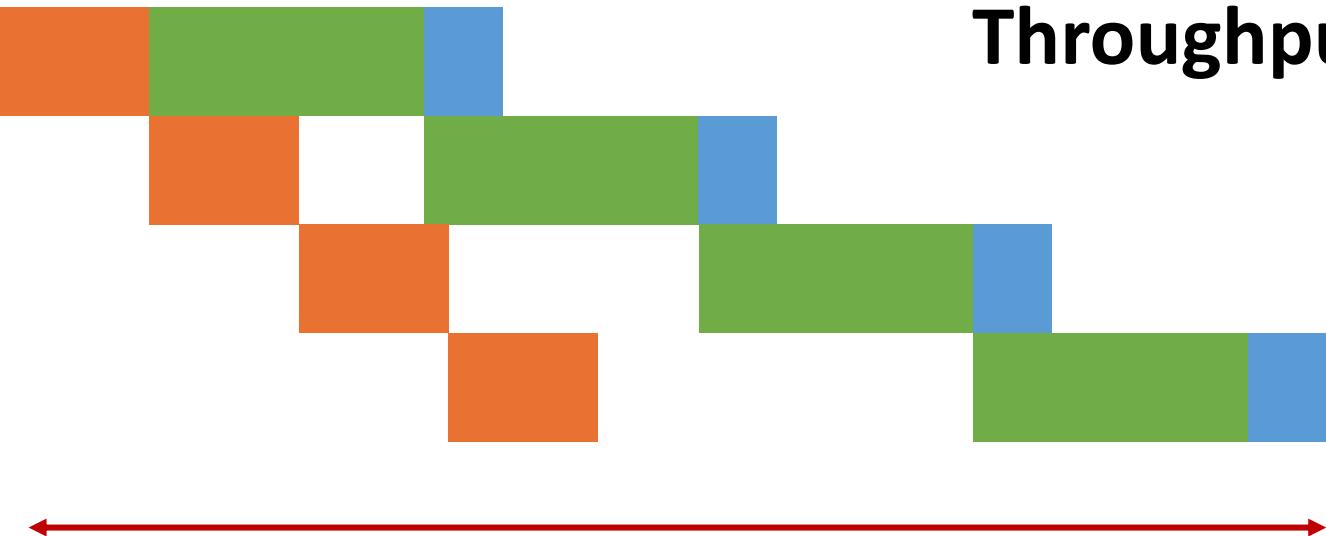
b) what would be a better (faster) strategy?



# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

b) what would be a better (faster) strategy?



**Throughput?**

With lead in & lead out  
(fixed number of students)

1 person per 106.25 min (425 min / 4)

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

b) what would be a better (faster) strategy?



**Throughput?**

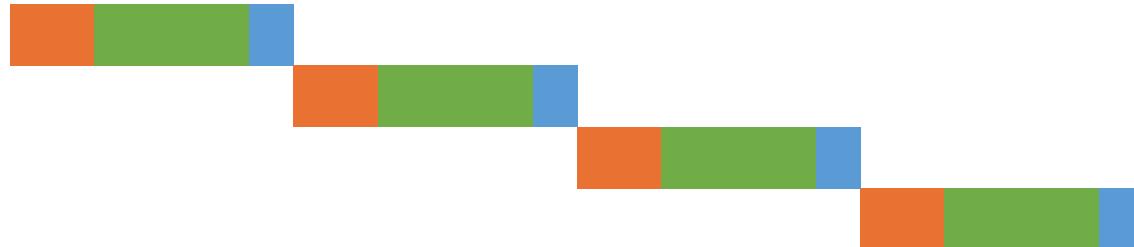
1 person per 90 min

Full utilization  
(indefinite number of students)

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

b) what is the achieved speedup?



Total time (sequential): 620 min



Total time (optimized): 425 min

$$\text{Speedup: } S = \frac{T_a}{T_b} = \frac{620}{425} \approx 1.46$$

# Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

Task b) what would be a better (faster) strategy?



**Throughput?**

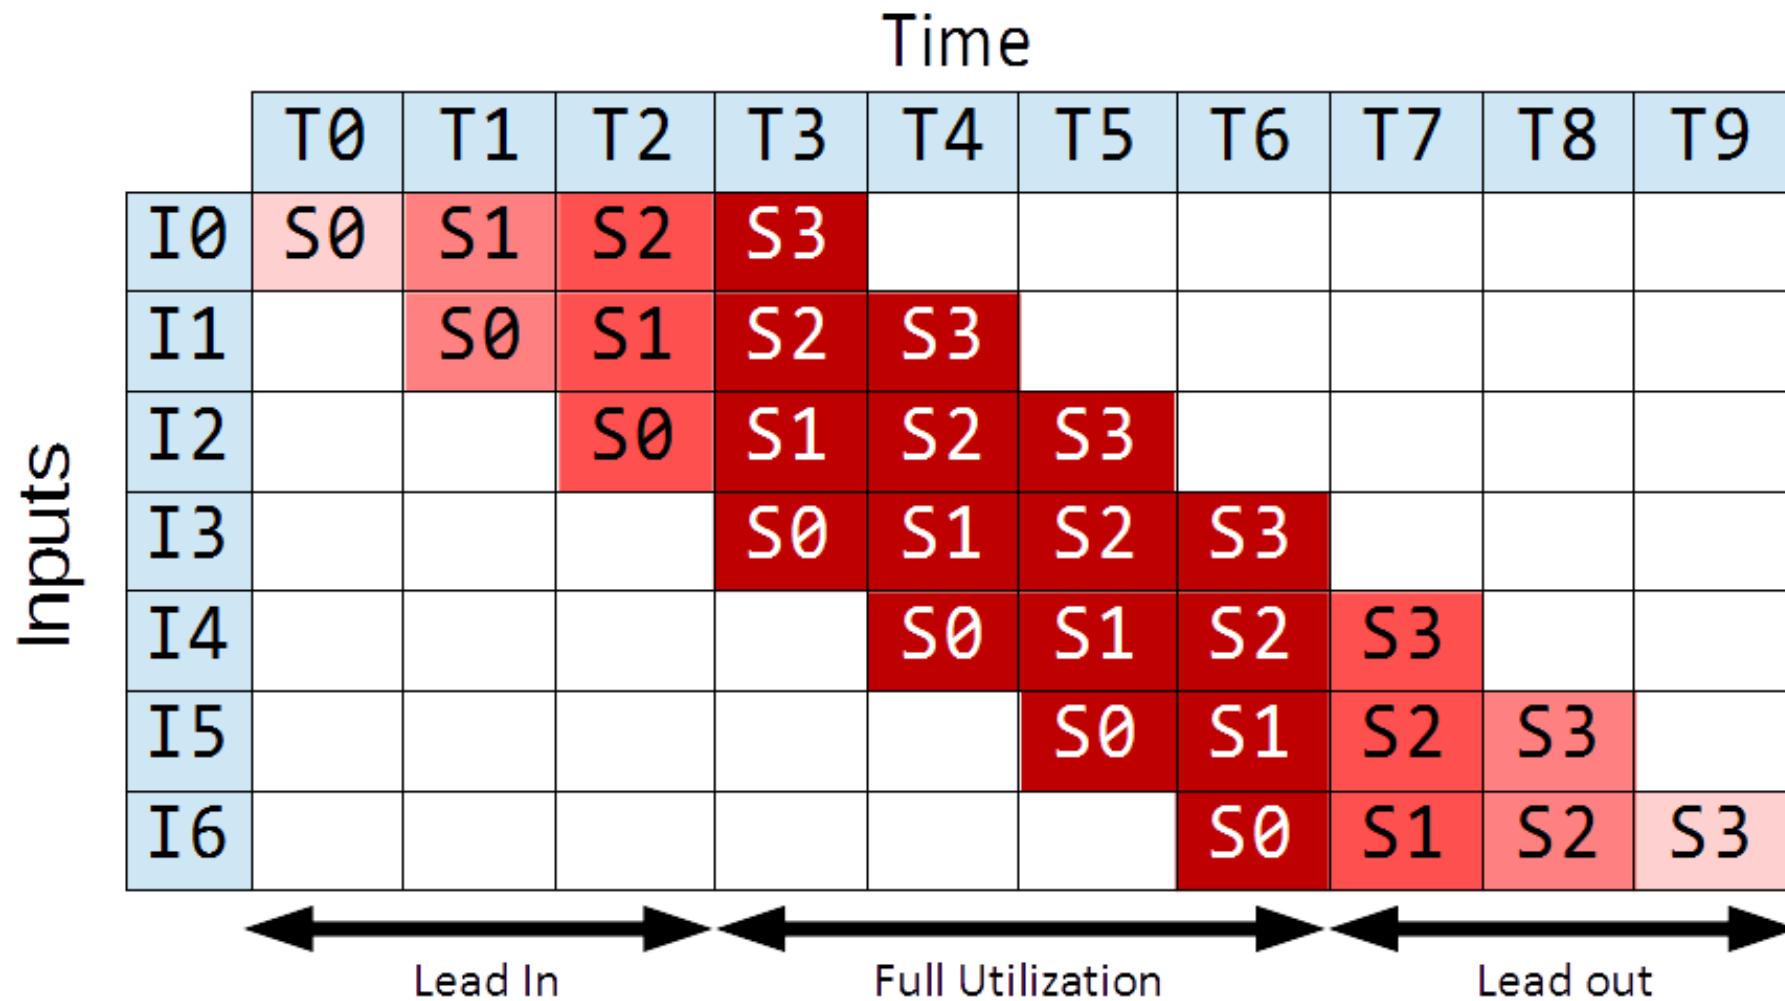
1 person per 90 min

How to compute throughput fast for pipelines with no duplicated stages?

$$\frac{1}{\max(\text{computation time}(stages))}$$

- Generalized formula:
- throughput =  $1 / \max(\text{pipeline stage time} / \text{execution units})$

# Pipelining



# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

c) what if they bought another **dryer**?



# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

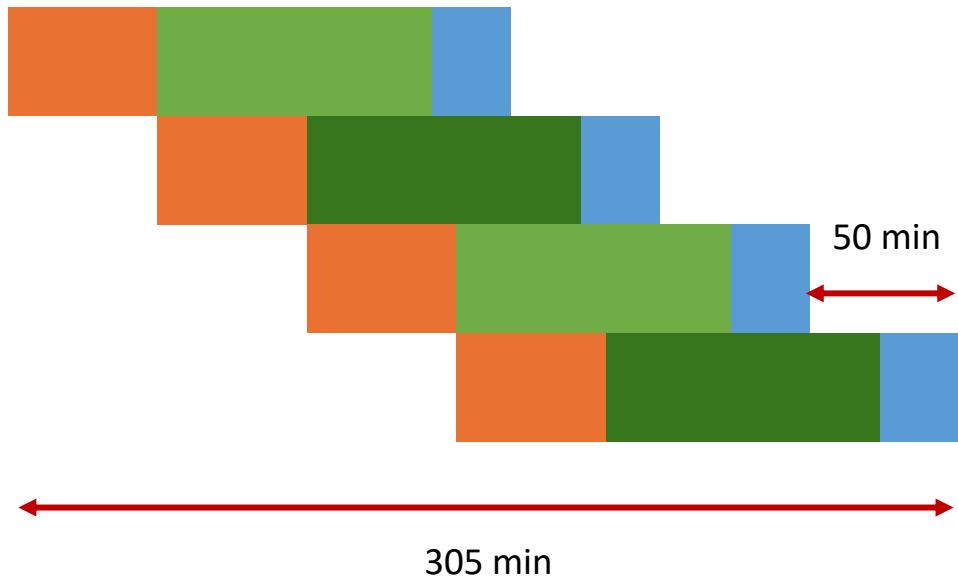
c) what if they bought another **dryer**?



# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

c) what if they bought another **dryer**?



**Latency?**

155 min

**Throughput?**

1 person per 50 min  
(full utilization)

1 person per 76.25 min  
(with lead in & lead out)

# Task 1: Pipelining

**Washing** - 50 min, **Dryer** - 90 min, **Iron** - 15 min

c) what if they bought another **dryer**?



Pipeline is not balanced as the stages  
do not take the same time.

# Task 2: Pipelining II

```
for (int i = 0; i < size; i++) {  
    data[i] = data[i] * data[i];  
}
```

```
for (int i = 0; i < size; i += 2) {  
    j = i + 1;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
}
```

```
for (int i = 0; i < size; i += 4) {  
    j = i + 1;  
    k = i + 2;  
    l = i + 3;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
    data[k] = data[k] * data[k];  
    data[l] = data[l] * data[l];  
}
```

# Task 2: Pipelining II

```
for (int i = 0; i < size; i++) {  
    data[i] = data[i] * data[i];  
}  
  
for (int i = 0; i < size; i += 2) {  
    j = i + 1;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
}  
  
for (int i = 0; i < size; i += 4) {  
    j = i + 1;  
    k = i + 2;  
    l = i + 3;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
    data[k] = data[k] * data[k];  
    data[l] = data[l] * data[l];  
}
```

## Assumptions:

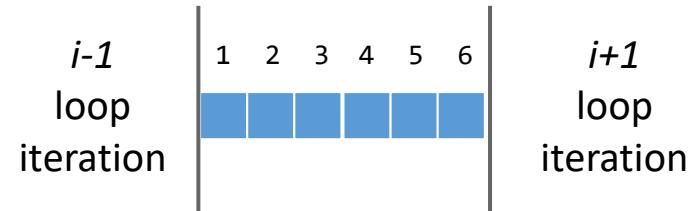
- Only one instruction can be issued per cycle.
- Loop body computation must be fully finished before next iteration starts
- Consider only arithmetic expressions in loop body and ignore all other operations (i.e., loads, stores, loop counter increment, etc.)
- Consider only *execute* step of an instruction (e.g., ignore fetch, decode, etc.)

How many cycles does the processor need to execute the following loops?

- Addition takes 3 cycles to execute
- Multiplication takes 6 cycles to execute

# Task 2: Pipelining II

```
for (int i = 0; i < size; i++) {  
    data[i] = data[i] * data[i];  
}
```

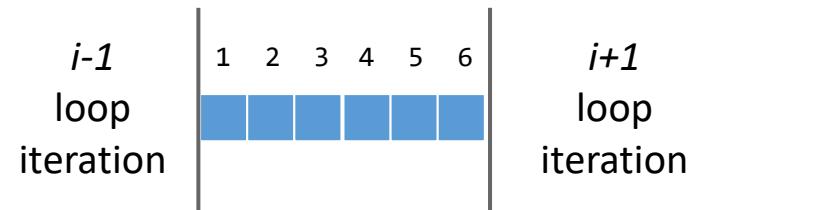


```
for (int i = 0; i < size; i += 2) {  
    j = i + 1;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
}
```

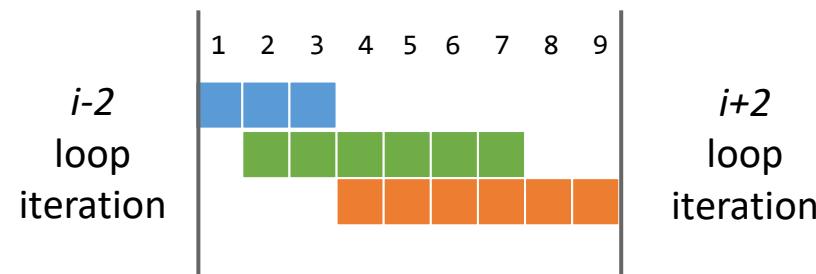
```
for (int i = 0; i < size; i += 4) {  
    j = i + 1;  
    k = i + 2;  
    l = i + 3;  
    data[i] = data[i] * data[i];  
    data[j] = data[j] * data[j];  
    data[k] = data[k] * data[k];  
    data[l] = data[l] * data[l];  
}
```

# Task 2: Pipelining II

```
for (int i = 0; i < size; i++) {
    data[i] = data[i] * data[i];
}
```



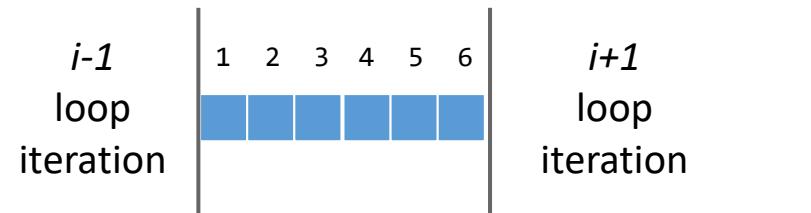
```
for (int i = 0; i < size; i += 2) {
    j = i + 1;
    data[i] = data[i] * data[i];
    data[j] = data[j] * data[j];
}
```



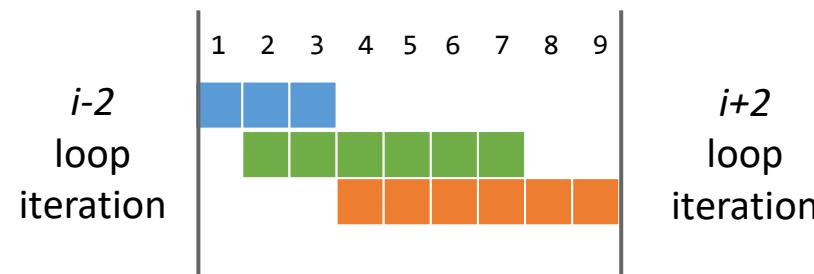
```
for (int i = 0; i < size; i += 4) {
    j = i + 1;
    k = i + 2;
    l = i + 3;
    data[i] = data[i] * data[i];
    data[j] = data[j] * data[j];
    data[k] = data[k] * data[k];
    data[l] = data[l] * data[l];
}
```

# Task 2: Pipelining II

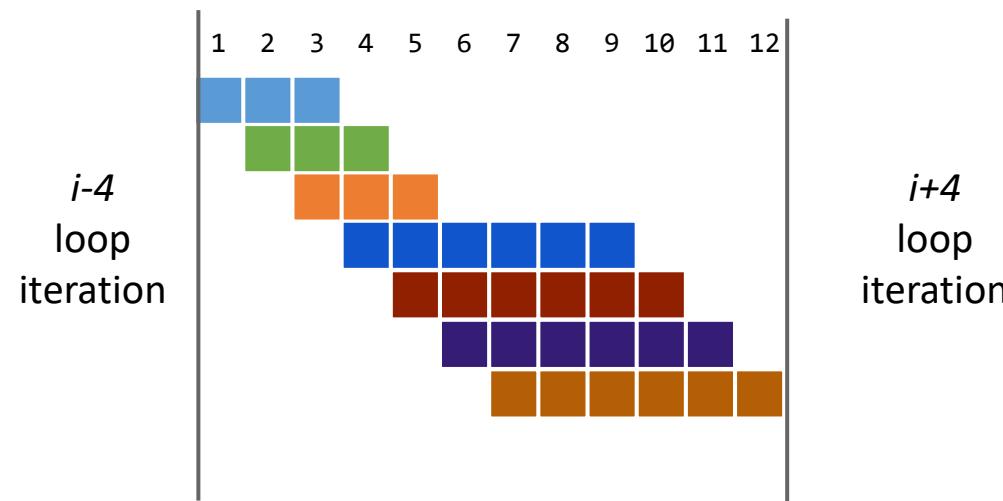
```
for (int i = 0; i < size; i++) {
    data[i] = data[i] * data[i];
}
```



```
for (int i = 0; i < size; i += 2) {
    j = i + 1;
    data[i] = data[i] * data[i];
    data[j] = data[j] * data[j];
}
```



```
for (int i = 0; i < size; i += 4) {
    j = i + 1;
    k = i + 2;
    l = i + 3;
    data[i] = data[i] * data[i];
    data[j] = data[j] * data[j];
    data[k] = data[k] * data[k];
    data[l] = data[l] * data[l];
}
```



# Task 3: Loop Parallelism

Can we parallelize the following loops?

```
for (int i=1; i<size; i++) { // for Loop: i from 1 to (size-1)
    if (data[i-1] > 0)           // If the previous value is positive
        data[i] = (-1)*data[i];   // change the sign of this value
}
```

# Task 3: Loop Parallelism

Can we parallelize the following loops?

```
for (int i=1; i<size; i++) { // for Loop: i from 1 to (size-1)
    if (data[i-1] > 0)           // If the previous value is positive
        data[i] = (-1)*data[i];   // change the sign of this value
}
                                // end for Loop
```

```
for (int i = 0; i < size; i++) {      // for Loop: i from 0 to (size-1)
    data[i] = Math.sin(data[i]);        // calculate sin() of the value
}
                                // end for Loop
```

# Task 3: Loop Parallelism

Can we parallelize the following loops?

```
for (int i=1; i<size; i++) { // for Loop: i from 1 to (size-1)
    if (data[i-1] > 0)          // If the previous value is positive
        data[i] = (-1)*data[i];   // change the sign of this value
}
                                // end for Loop
```

```
for (int i = 0; i < size; i++) {      // for Loop: i from 0 to (size-1)
    data[i] = Math.sin(data[i]);        // calculate sin() of the value
}
                                // end for Loop
```

# Example Task

# SOLA Stafette

- 14 runners (numbered from 0 to 13)
- each runner can start after the previous runner finished (except the first one).

Under which assumption does the following code work?

Initial value of y?

0

Reference to object y?

needs to be the same object shared across all the threads

```
public class RunnerThread extends Thread {  
  
    private AtomicInteger y;  
    private int id;  
  
    public RunnerThread(int id, AtomicInteger y) {  
        this.id = id;  
        this.y = y;  
    }  
  
    public void run(){  
        while (y.get() != this.id) {  
            // warten bis ich an der Reihe bin / wait until it is my turn  
        }  
  
        // Laufen / do running  
        y.incrementAndGet();  
    }  
}
```

# SOLA Stafette

- 14 runners (numbered from 0 to 13)
- each runner can start after the previous runner finished (except the first one).

Under which assumption does the following code work?

Initial value of y?

0

Reference to object y?

needs to be the same object shared across all the threads

```
public class RunnerThread extends Thread {  
  
    private AtomicInteger y;  
    private int id;  
  
    public RunnerThread(int id, AtomicInteger y) {  
        this.id = id;  
        this.y = y;  
    }  
  
    public void run(){  
        while (y.get() != this.id) {  
            // warten bis ich an der Reihe bin / wait until it is my turn  
        }  
  
        // Laufen / do running  
        y.incrementAndGet();  
    }  
}
```

Ensures that `y.get()` returns the updated value after calling `y.incrementAndGet();`

# SOLA Stafette

# Complete the implementation using wait/notify

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            .....  
            .....  
            .....  
            .....  
            .....  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            while (condition) {  
                ?.wait();  
            }  
        .....  
        .....  
        .....  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            while (condition) {  
                ?.wait();  
            }  
            // Laufen / do running  
            .....  
            .....  
            .....  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            while (condition) {  
                ?.wait();  
            }  
            // Laufen / do running  
            // update the condition  
        }  
    }  
}
```

# SOLA Stafette

## Complete the implementation using wait/notify

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

**This is a basic pattern that you will see in many tasks.  
Now, let's fill in the details.**

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            while (condition) {  
                ?.wait();  
            }  
            // Laufen / do running  
            // update the condition  
            ?.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. **What to synchronize on?**
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (?) {  
            while (condition) {  
                ?.wait();  
            }  
            // Laufen / do running  
            // update the condition  
            ?.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. **What to synchronize on?**
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (condition) {  
                x.wait();  
            }  
            // Laufen / do running  
            // update the condition  
            x.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. **Check that it's our turn, wait  
otherwise**
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private ..... x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, ..... x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (condition) {  
                x.wait();  
            }  
            // Laufen / do running  
            // update the condition  
            x.notify(); // or notifyAll()  
        }  
    }  
}
```

while (y.get() != this.id) {  
 original  
 condition  
 can we reuse it?  
 Yes!

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. **Check that it's our turn, wait  
otherwise**
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private AtomicInteger x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, AtomicInteger x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            // update the condition  
            x.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

## Complete the implementation using wait/notify

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. **Update the condition**
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private AtomicInteger x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, AtomicInteger x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            x.incrementAndGet(); // update the condition  
            x.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

## Complete the implementation using wait/notify

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. **Notify waiting threads**

## Should we use notify or notifyAll?

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private AtomicInteger x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, AtomicInteger x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            x.incrementAndGet(); // update the condition  
            x.notify(); // or notifyAll()  
        }  
    }  
}
```

# SOLA Stafette

## Complete the implementation using wait/notify

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. **Notify waiting threads**

## Should we use notify or notifyAll?

notifyAll because multiple threads can be waiting

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private AtomicInteger x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, AtomicInteger x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            x.incrementAndGet(); // update the condition  
            x.notifyAll();  
        }  
    }  
}
```

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private AtomicInteger x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, AtomicInteger x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public ..... void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            x.incrementAndGet(); // update the condition  
            x.notifyAll();  
        }  
    }  
}
```

**Can we replace  
synchronized (x)  
with synchronizing  
the method?**

No

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {

    private AtomicInteger x;
    private int id;

    public WaitNotifyRunnerThread(int id, AtomicInteger x) {
        this.id = id;
        this.x = x;
    }

    public void run(){
        synchronized (x) {
            while (x.get() != this.id) {
                x.wait();
            }
            // Laufen / do running
            x.incrementAndGet(); // update the condition
            x.notifyAll();
        }
    }
}
```

**Can we replace**  
AtomicInteger x  
**With**  
Integer x  
?

# SOLA Stafette

**Complete the implementation  
using wait/notify**

1. We'll definitely need synchronization
2. What to synchronize on?
3. Check that it's our turn, wait otherwise
4. Do work (run)
5. Update the condition
6. Notify waiting threads

```
public class WaitNotifyRunnerThread extends Thread {

    private Integer x;
    private int id;

    public WaitNotifyRunnerThread(int id, Integer x) {
        this.id = id;
        this.x = x;
    }

    public void run(){
        synchronized (x) {
            while (x.get() != this.id) {
                x.wait();
            }
            // Laufen / do running
            x += 1; // update the condition
            x.notifyAll();
        }
    }
}
```

**Can we replace**  
AtomicInteger x  
**With**  
Integer x  
?

# SOLA Stafette

Complete the implementation  
using wait/notify

Do these assumptions still hold?

Initial value of x?

0

Reference to object x?

needs to be the same object  
shared across all the threads

No, because  $x += 1$  in Java  
creates a new Object!

Integer is immutable!

```
public class WaitNotifyRunnerThread extends Thread {  
  
    private Integer x;  
    private int id;  
  
    public WaitNotifyRunnerThread(int id, Integer x) {  
        this.id = id;  
        this.x = x;  
    }  
  
    public void run(){  
        synchronized (x) {  
            while (x.get() != this.id) {  
                x.wait();  
            }  
            // Laufen / do running  
            x += 1; // update the condition  
            x.notifyAll();  
        }  
    }  
}
```

Can we replace  
AtomicInteger x

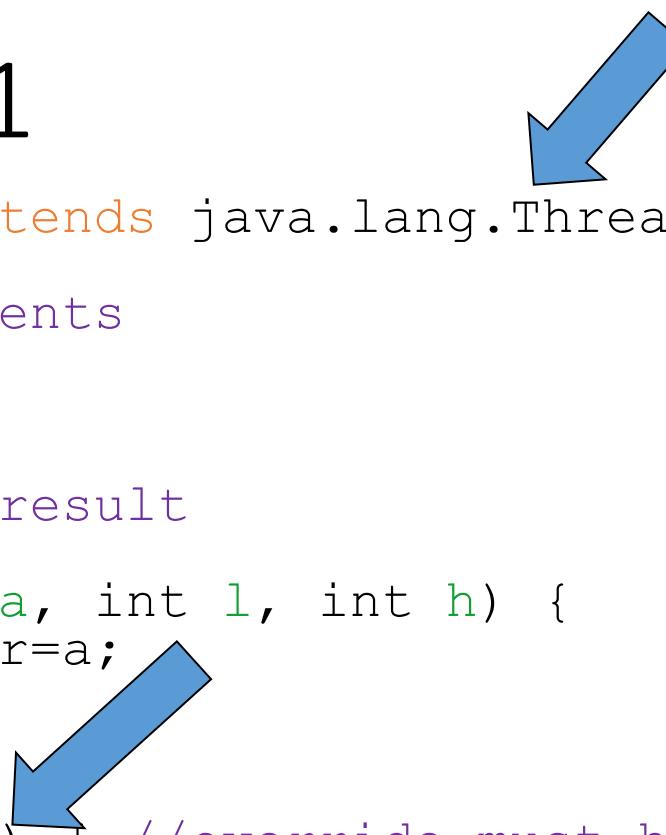
With  
Integer x  
?

# Plan für heute

- Organisation
- Nachbesprechung Exercise 4
- **Theory Recap**
- Intro Exercise 5
- Exam Questions
- Kahoot

# First attempt, part 1

```
class SumThread extends java.lang.Thread {  
  
    int lo; // arguments  
    int hi;  
    int[] arr;  
  
    int ans = 0; // result  
  
    SumThread(int[] a, int l, int h) {  
        lo=l; hi=h; arr=a;  
    }  
  
    public void run() { //override must have this type  
        for(int i=lo; i < hi; i++)  
            ans += arr[i];  
    }  
}
```



Because we must override a no-arguments/no-result run,  
we use fields to communicate across threads

# First attempt, continued (wrong)

```
class SumThread extends java.lang.Thread {  
    int lo, int hi, int[] arr; // arguments  
    int ans = 0; // result  
    SumThread(int[] a, int l, int h) { ... }  
    public void run(){ ... } // override  
}  
  
int sum(int[] arr){ // can be a static method  
    int len = arr.length;  
    int ans = 0;  
    SumThread[] ts = new SumThread[4];  
    for(int i=0; i < 4; i++) // do parallel computations  
        ts[i] = new SumThread(arr,i*len/4, (i+1)*len/4);  
  
    for(int i=0; i < 4; i++) // combine results  
        ans += ts[i].ans;  
    return ans;  
}
```

# Second attempt (still wrong)

```
class SumThread extends java.lang.Thread {  
    int lo, int hi, int[] arr; // arguments  
    int ans = 0; // result  
    SumThread(int[] a, int l, int h) { ... }  
    public void run(){ ... } // override  
}  
  
int sum(int[] arr){ // can be a static method  
    int len = arr.length;  
    int ans = 0;  
    SumThread[] ts = new SumThread[4];  
    for(int i=0; i < 4; i++){ // do parallel computations  
        ts[i] = new SumThread(arr,i*len/4,(i+1)*len/4);  
        ts[i].start(); // start actually runs the thread in parallel  
    }  
    for(int i=0; i < 4; i++) // combine results  
        ans += ts[i].ans;  
    return ans;  
}
```

# Third attempt (correct in spirit)

```
class SumThread extends java.lang.Thread {  
    int lo, int hi, int[] arr; // arguments  
    int ans = 0; // result  
    SumThread(int[] a, int l, int h) { ... }  
    public void run(){ ... } // override  
}  
  
int sum(int[] arr){ // can be a static method  
    int len = arr.length;  
    int ans = 0;  
    SumThread[] ts = new SumThread[4];  
    for(int i=0; i < 4; i++){ // do parallel computations  
        ts[i] = new SumThread(arr,i*len/4,(i+1)*len/4);  
        ts[i].start();  
    }  
    for(int i=0; i < 4; i++) { // combine results  
        ts[i].join(); // wait for helper to finish!  
        ans += ts[i].ans;  
    }  
    return ans;  
}
```

# Issues with this approach (and some workarounds)

Several reasons why this is a poor parallel algorithm

**Reason 1: want code to be reusable and efficient across platforms**

```
int sum(int[] arr){ // can be a static method
    int len = arr.length;
    int ans = 0;
    SumThread[] ts = new SumThread[4];
    for(int i=0; i < 4; i++){ // do parallel computations
        ts[i] = new SumThread(arr,i*len/4,(i+1)*len/4);
        ts[i].start();
    }
    for(int i=0; i < 4; i++) { // combine results
        ts[i].join(); // wait for helper to finish
        ans += ts[i].ans;
    }
    return ans;
}
```

# Issues with this approach (and some workarounds)

**Reason 2:** want to use (only) processors “available to you now”

- Not used by other programs or threads in your program
  - Maybe caller is also using parallelism
  - Available cores can change even while your threads run

**Reason 3:** Though unlikely for `sum`, in general subproblems may take significantly different amounts of time

# Divide and Conquer

Examples of divide-and-conquer algorithms are quicksort, mergesort, Strassen matrix multiplication, and many others.

Basic structure of a divide-and-conquer algorithm:

1. If problem is small enough, solve it directly
2. Otherwise
  - a. Break problem into subproblems
  - b. Solve subproblems recursively
  - c. Assemble solutions of subproblems into overall solution

# Adding Numbers from Vector: (Recursive Version)

```
public static int do_sum_rec(int[] xs, int l, int h) {  
    int size = h - l;  
    if (size == 1)  
        return xs[l];  
  
    int mid = size / 2;  
    int sum1 = do_sum_rec(xs, l, l + mid);  
    int sum2 = do_sum_rec(xs, l + mid, h);  
  
    return sum1 + sum2;  
}
```

# Parallel Version: Task Parallelism Model

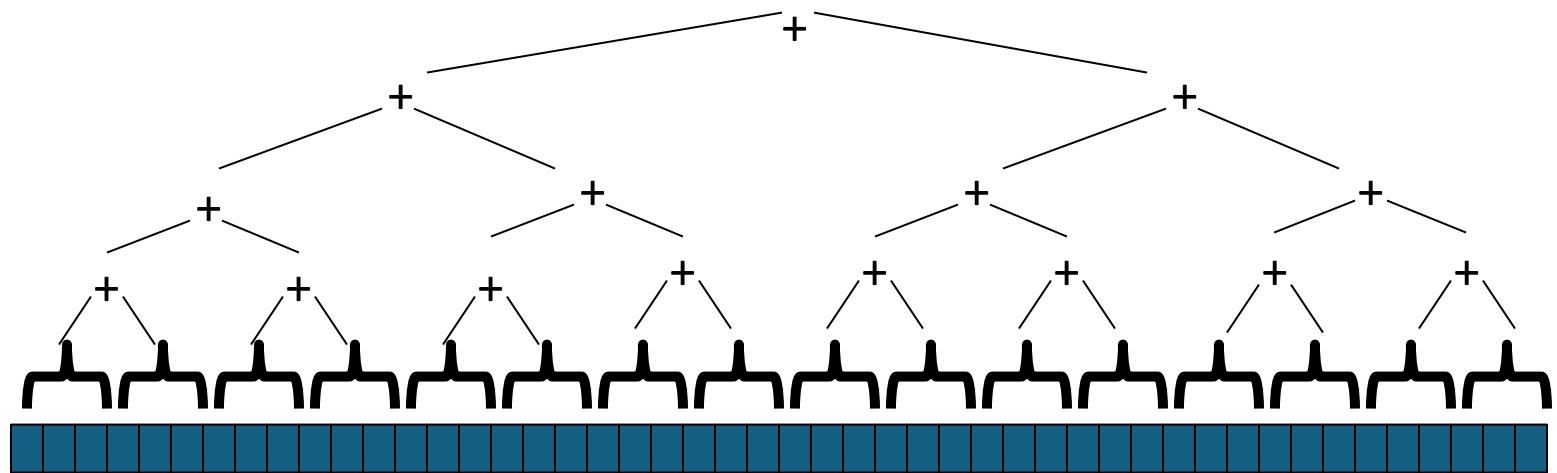
Basic flow of operations

- Create parallel tasks
- Wait for parallel tasks to complete

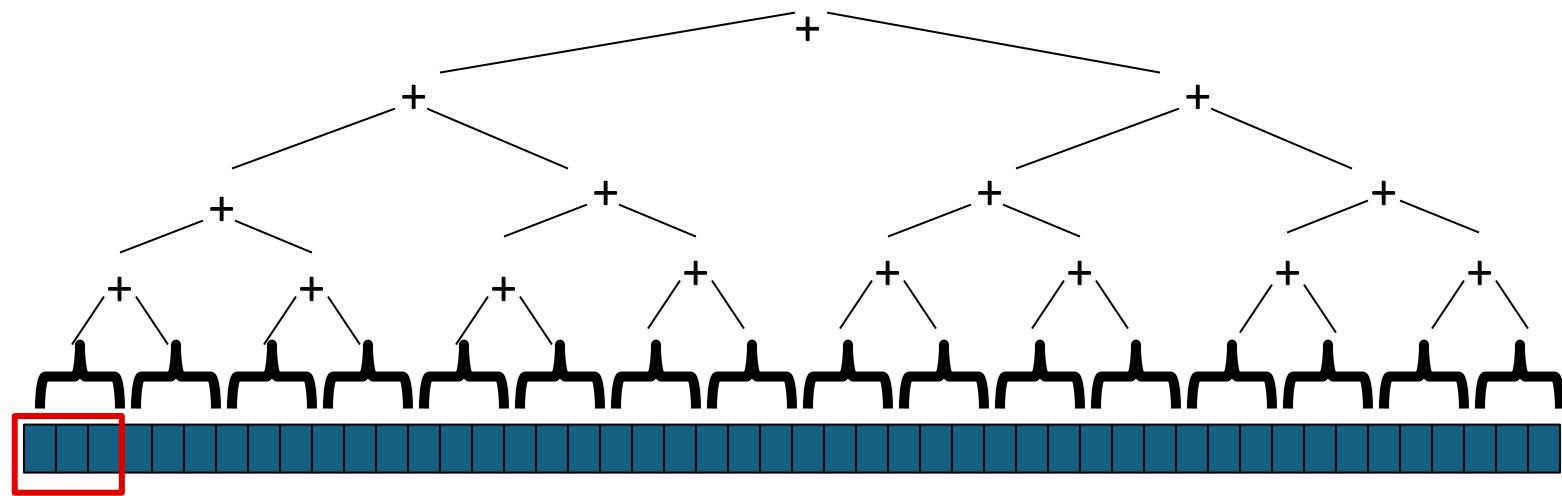
Parallel versions of divide-and-conquer on the task model are trivial:

- Create a task for the first and second part
- Wait for tasks to complete
- Combine their results

# Divide and Conquer



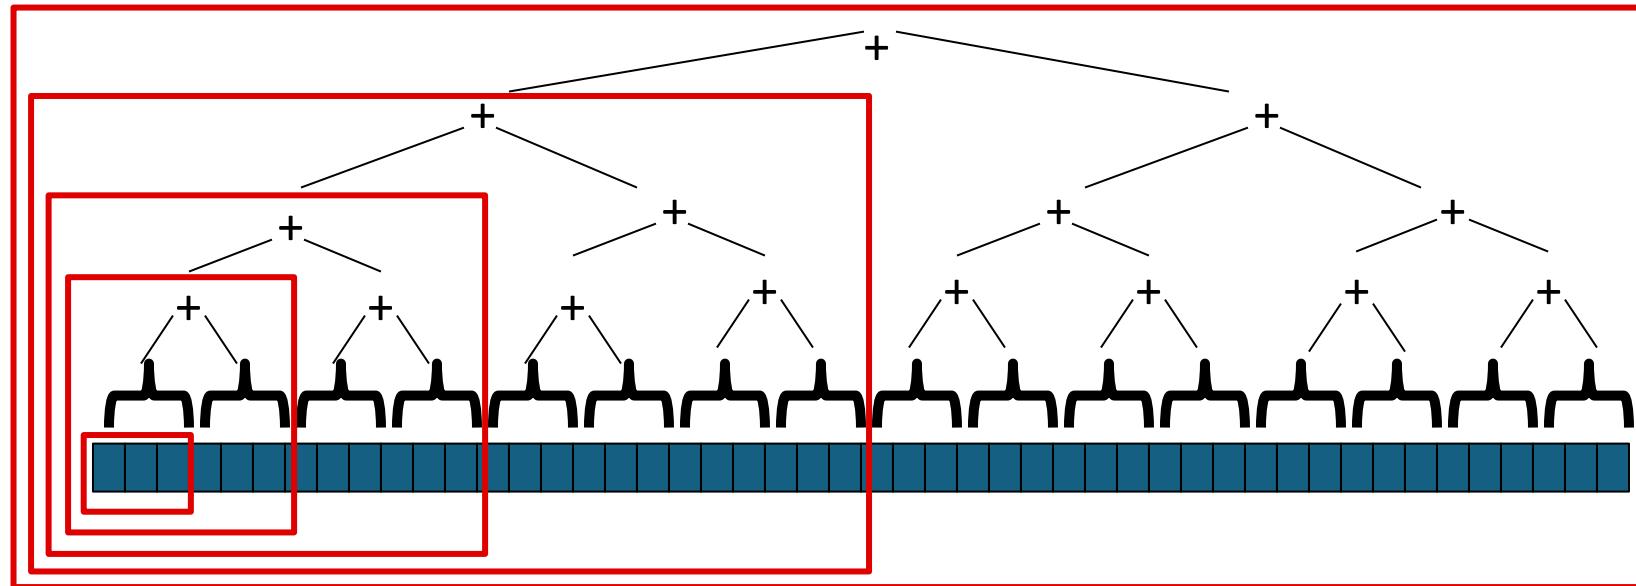
# Divide and Conquer



base case  
no further split

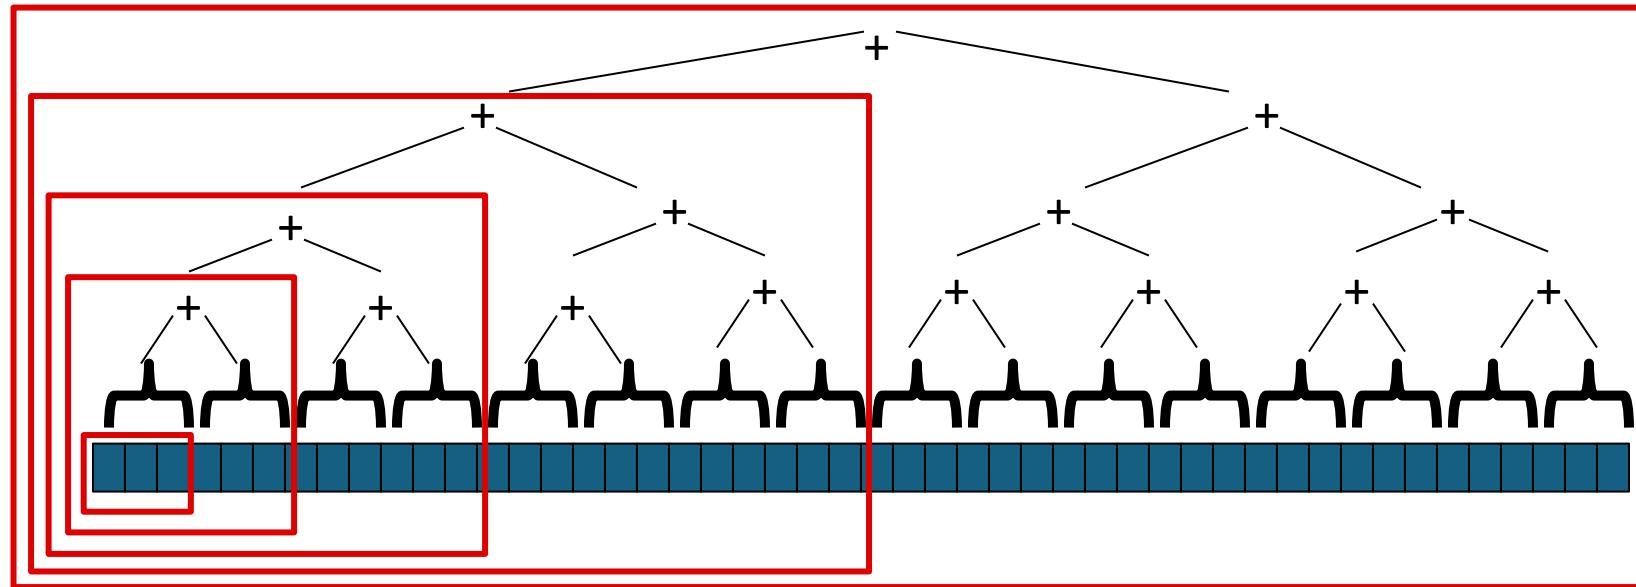
# Divide and Conquer

Tasks at different  
levels of granularity



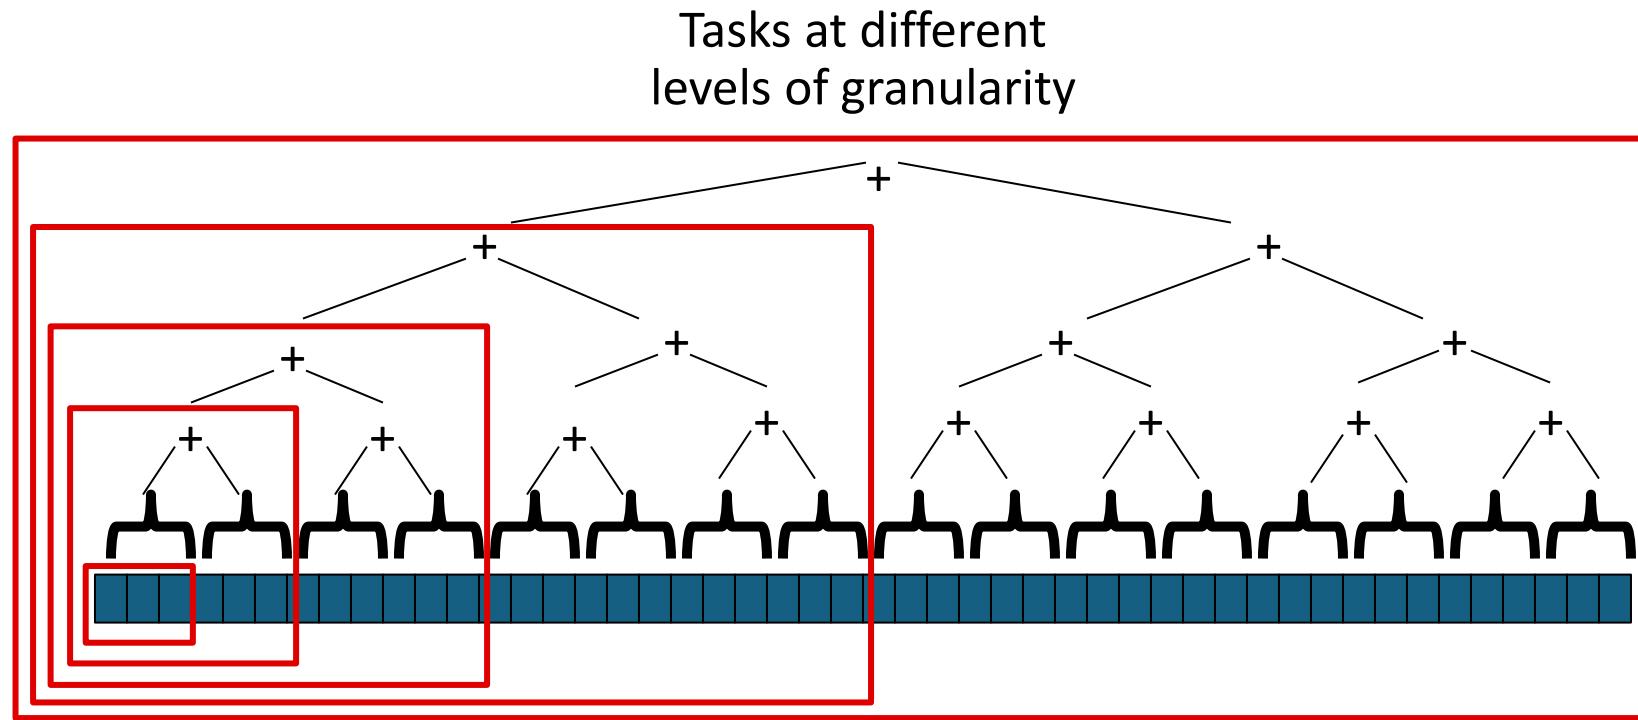
# Divide and Conquer

Tasks at different  
levels of granularity



What determines a task?

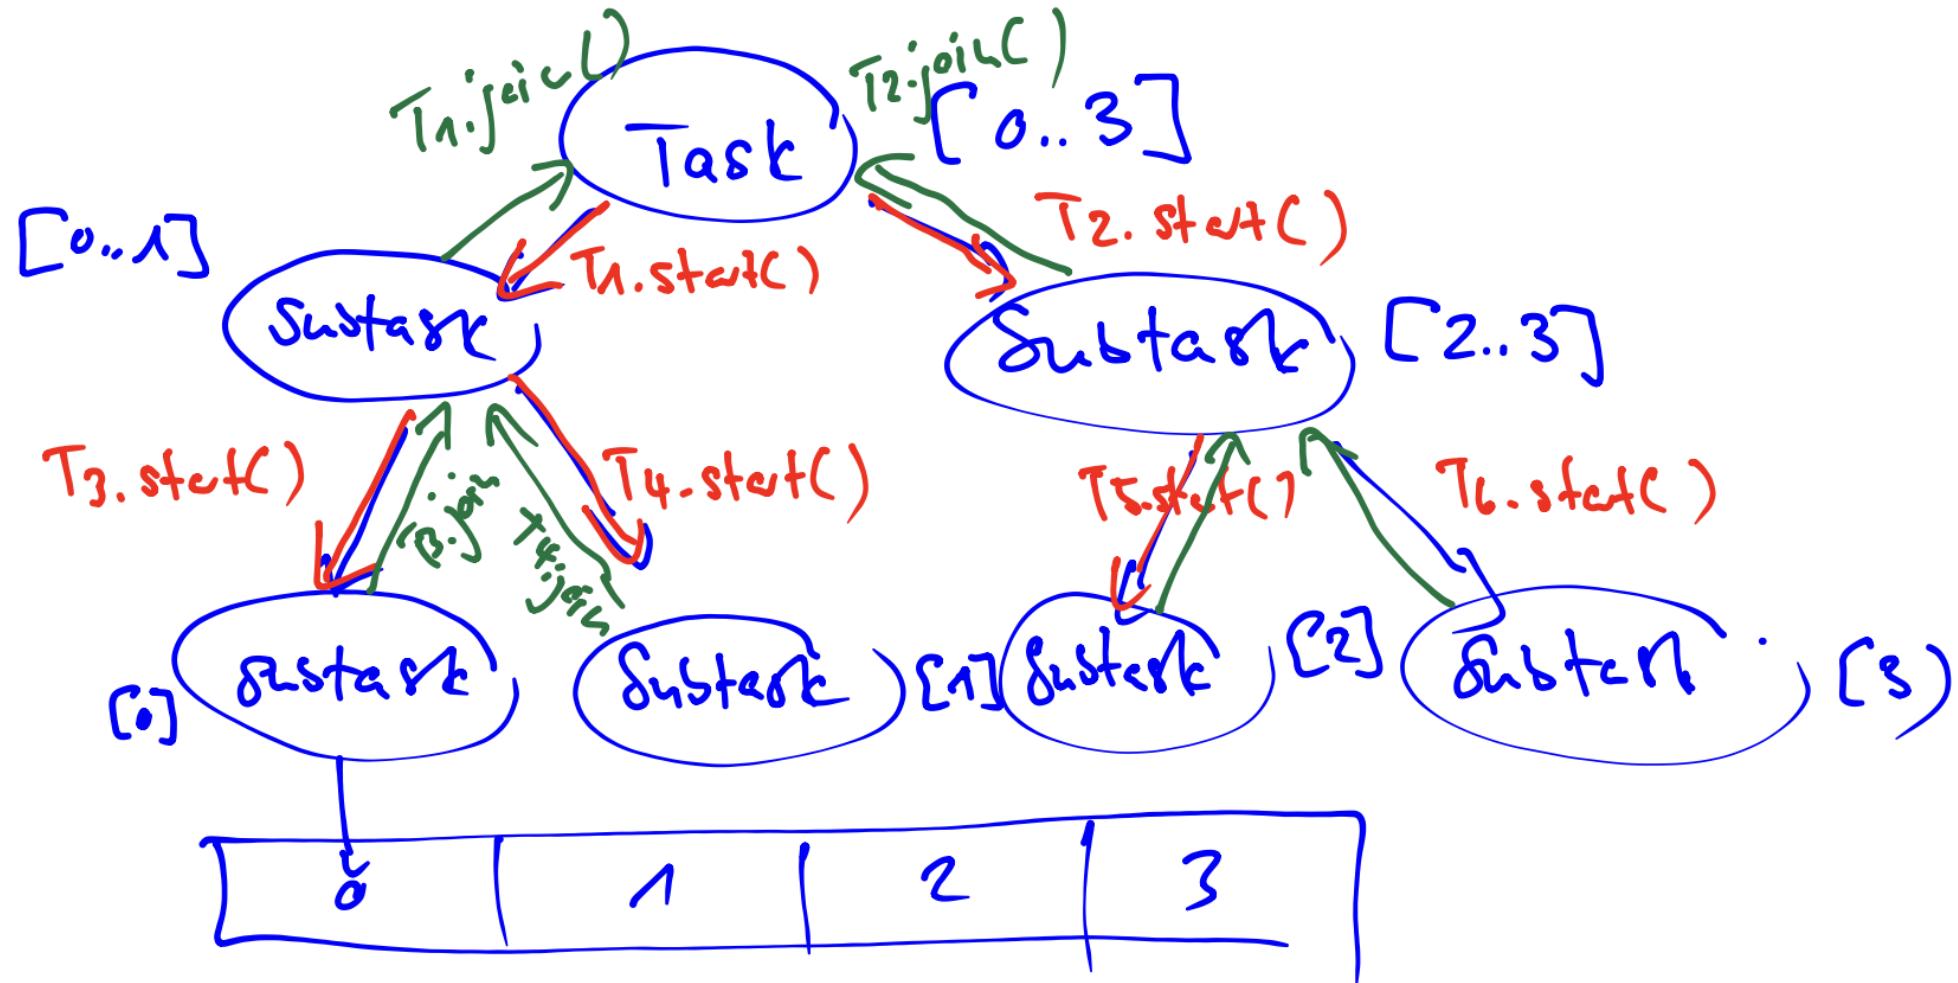
# Divide and Conquer



What determines a task?

- i) input array
- ii) start index
- iii) length/end index

These are fields we want to store in the task



# Parallel Recursive Sum (with Threads)

```
public void run(){
    int size = h-1;
    if (size == 1) {
        result = xs[1];
        return;
    }
    int mid = size / 2;
    SumThread t1 = new SumThread(xs, l, l + mid);
    SumThread t2 = new SumThread(xs, l + mid, h);

    t1.start();
    t1.join();
```

Is this OK?

```
t2.start();
t2.join();

result = t1.result + t2.result;
return;
}
```

# Parallel Recursive Sum (with Threads)

```
public void run(){
    int size = h-1;
    if (size == 1) {
        result = xs[1];
        return;
    }
    int mid = size / 2;
    SumThread t1 = new SumThread(xs, l, l + mid);
    SumThread t2 = new SumThread(xs, l + mid, h);

    t1.start();
    t2.start();

    t1.join(); ←
    t2.join();

    result = t1.result + t2.result;
    return;
}
```

Remark: This doesn't compile because  
join() can throw exceptions. In reality  
we need a try-catch block here.

# Result

**Java.lang.OutOfMemoryError: unable to create new native thread**

# Divide-and-conquer – with manual fixes (Pt. I)

```
public void run(){
    int size = h-l;
    if (size < SEQ_CUTOFF)
        for (int i=l; i<h; i++)
            result += xs[i];
    else {
        int mid = size / 2;
        SumThread t1 = new SumThread(xs, l, l + mid);
        SumThread t2 = new SumThread(xs, l + mid, h);
        t1.start();
        t2.start();
        t1.join();
        t2.join();
        result = t1.result + t2.result;
    }
}
```

# Half the threads

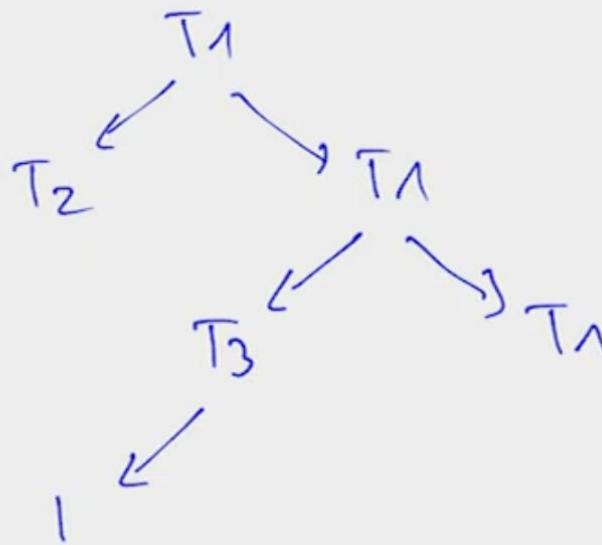
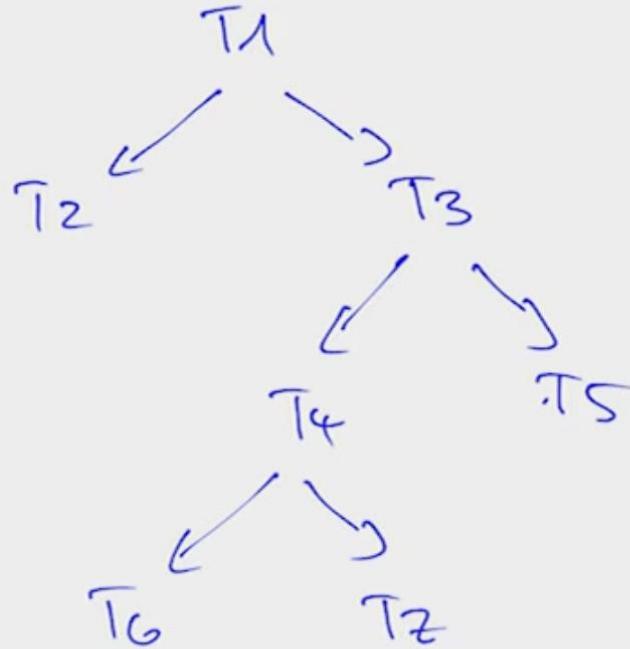
```
// wasteful: don't
SumThread t1 = ...
SumThread t2 = ...
t1.start();
t2.start();
t1.join();
t2.join();
result=t1.result+t2.result;
```

```
// better: do
// order of next 4 lines
// essential - why?
t1.start();
t2.run();
t1.join();
result=t1.result+t2.result;
```

If a *language* had built-in support for fork-join parallelism, we would expect this hand-optimization to be unnecessary

But the *library* we are using expects you to do it yourself (and the difference is surprisingly substantial)

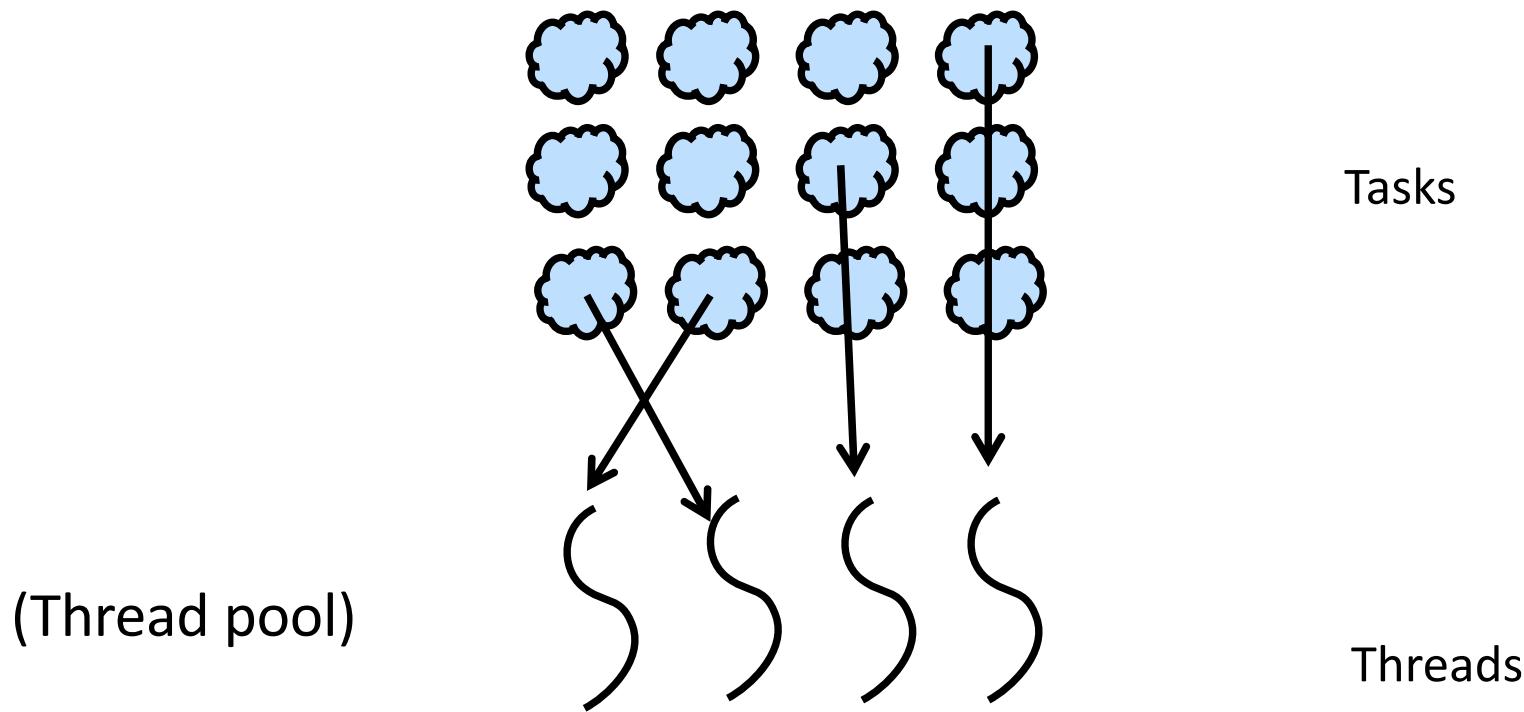
Again, no difference in theory



Java Threads are too  
heavyweight. What should we  
do?

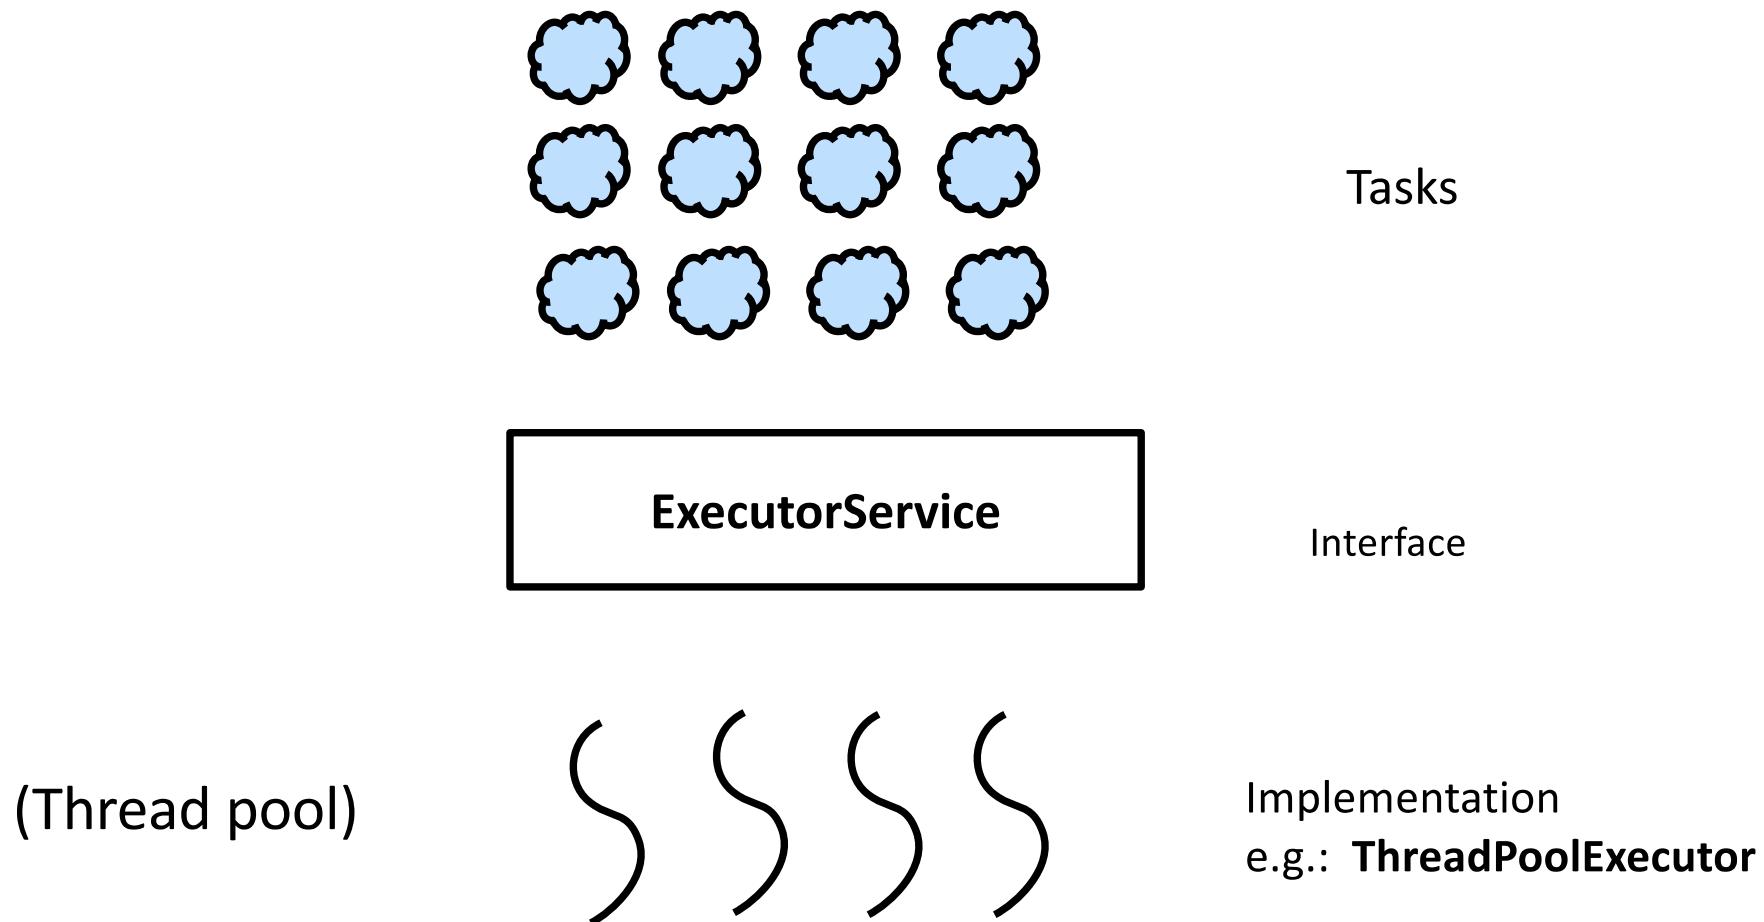
# ExecutorService!

# Alternative approach: schedule tasks on threads



How many threads would you use?

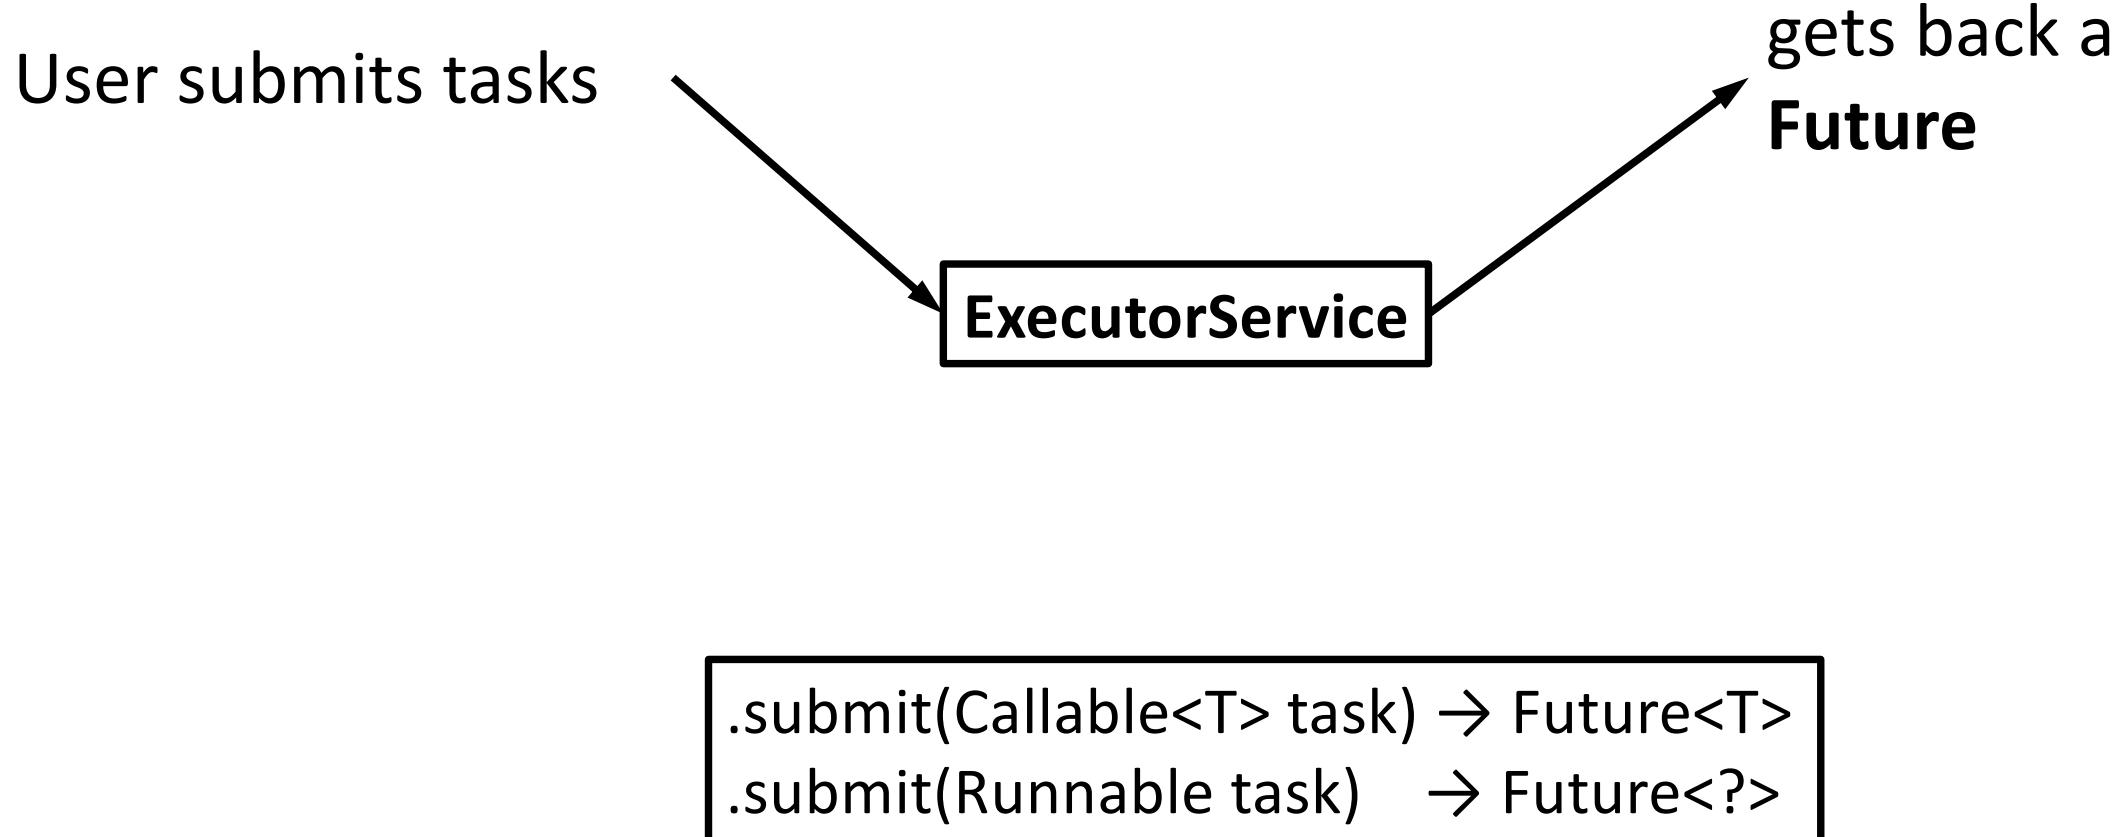
# Java's executor service: managing asynchronous tasks



# Why use ExecutorService?

- Before if we wanted to run multiple tasks in parallel, we had to create a new thread for each task
  - This is bad because OS threads are expensive
- Instead of creating a new thread per task, we use a pool of reusable threads
- Tasks are submitted to a task queue

# Java's executor service: managing asynchronous tasks



# Recursive Sum with ExecutorService

```
public Integer call() throws Exception {
    int size = h - 1;
    if (size == 1)
        return xs[1];

    int mid = size / 2;
    sumRecCall c1 = new sumRecCall(ex, xs, 1, 1 + mid);
    sumRecCall c2 = new sumRecCall(ex, xs, 1 + mid, h);

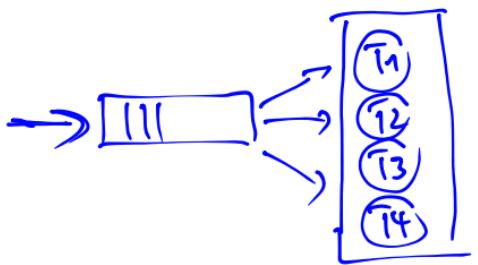
    Future<Integer> f1 = ex.submit(c1);
    Future<Integer> f2 = ex.submit(c2);

    return f1.get() + f2.get();
}
```

# Recursive Sum with ExecutorService

```
public Integer call() throws Exception {  
    int size = h - 1;  
    if (size == 1)  
        return xs[1];  
  
    int mid = size / 2;  
    sumRecCall c1 = new sumRecCall(ex, xs, 1, 1 + mid);  
    sumRecCall c2 = new sumRecCall(ex, xs, 1 + mid, h);  
  
    Future<Integer> f1 = ex.submit(c1);  
    Future<Integer> f2 = ex.submit(c2);  
  
    return f1.get() + f2.get();  
}
```

**Will this work?**



(T1)  $\text{sum}(0 \dots 100)$ :

$f_1 = \text{submit } \text{sum}(0 \dots 50)$   
 $f_2 = \text{submit } \text{sum}(50 \dots 100)$   
 return  $f_1.get() + f_2.get()$

(T2)  $\text{sum}(0 \dots 50)$ :

$f_1 = \text{submit } \text{sum}(0 \dots 25)$   
 $f_2 = \text{submit } \text{sum}(25 \dots 50)$   
 return  $f_1.get() + f_2.get()$

(T4)  $\text{sum}(0 \dots 25)$

$\text{sum}(50 \dots 100)$  · (T3)

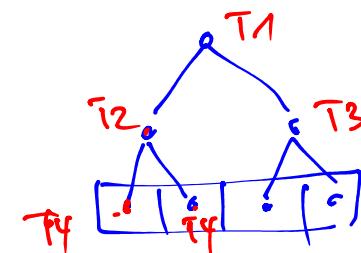
;

;

;

;

array size : 4

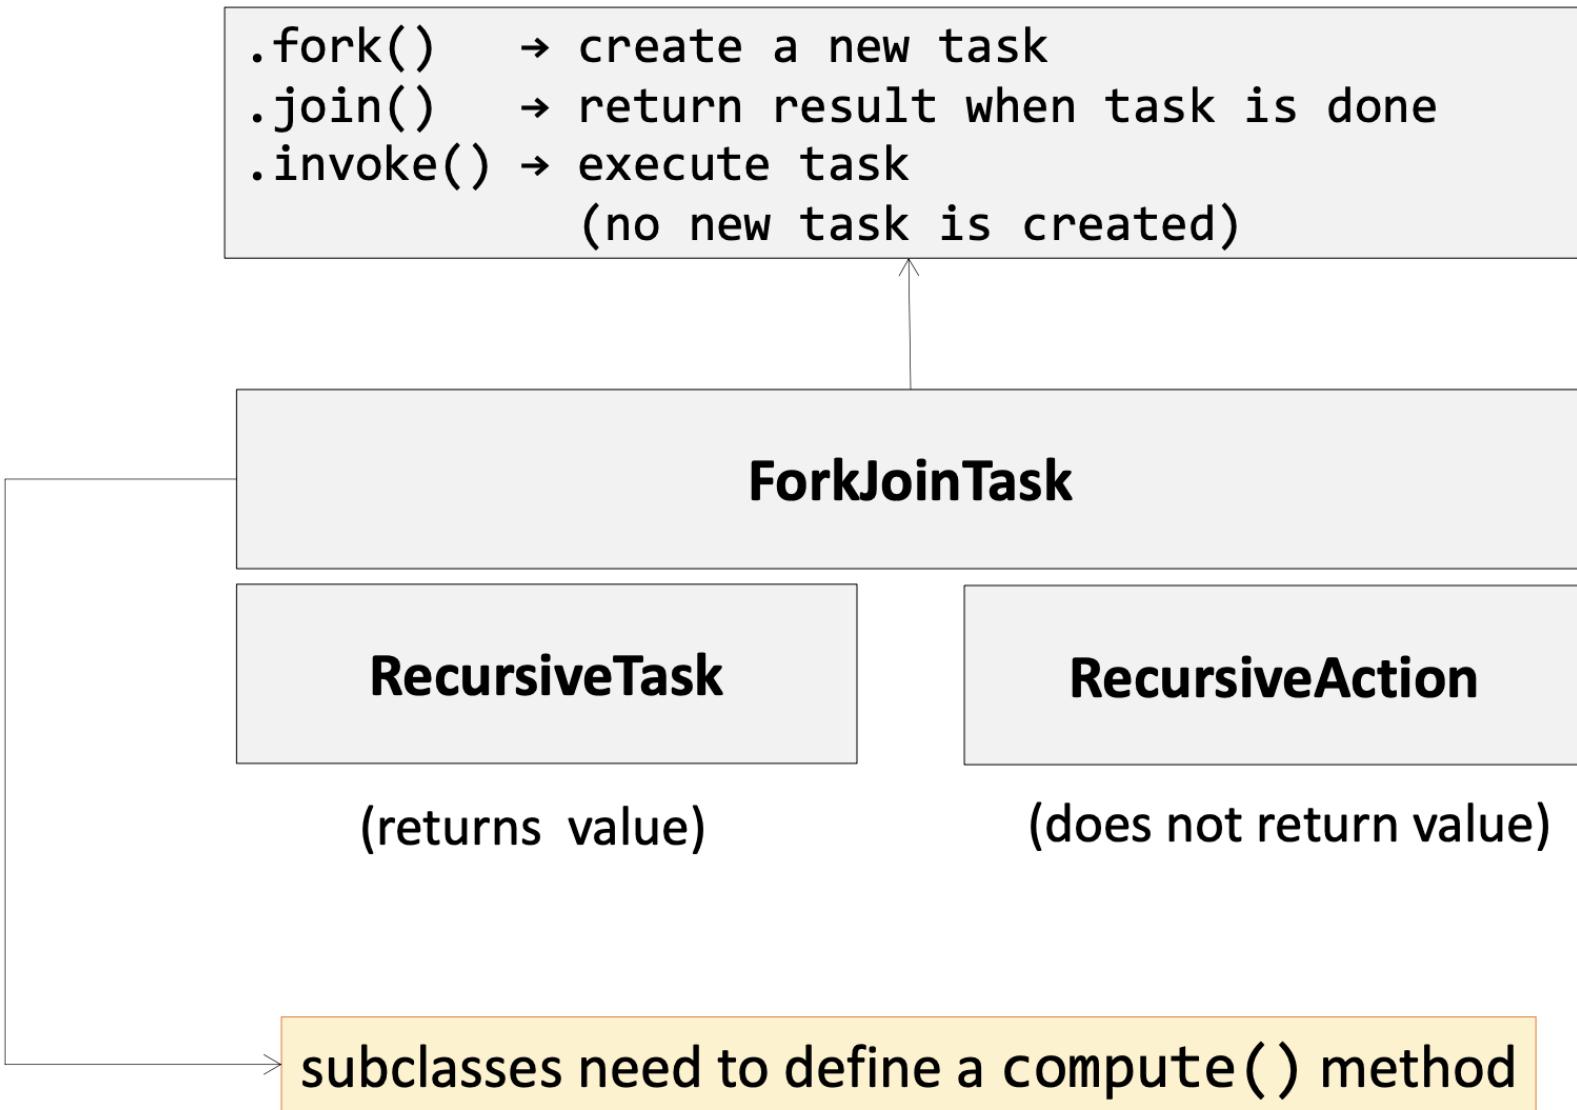


# Lesson:

ExecutorService is good for “flat” problems.

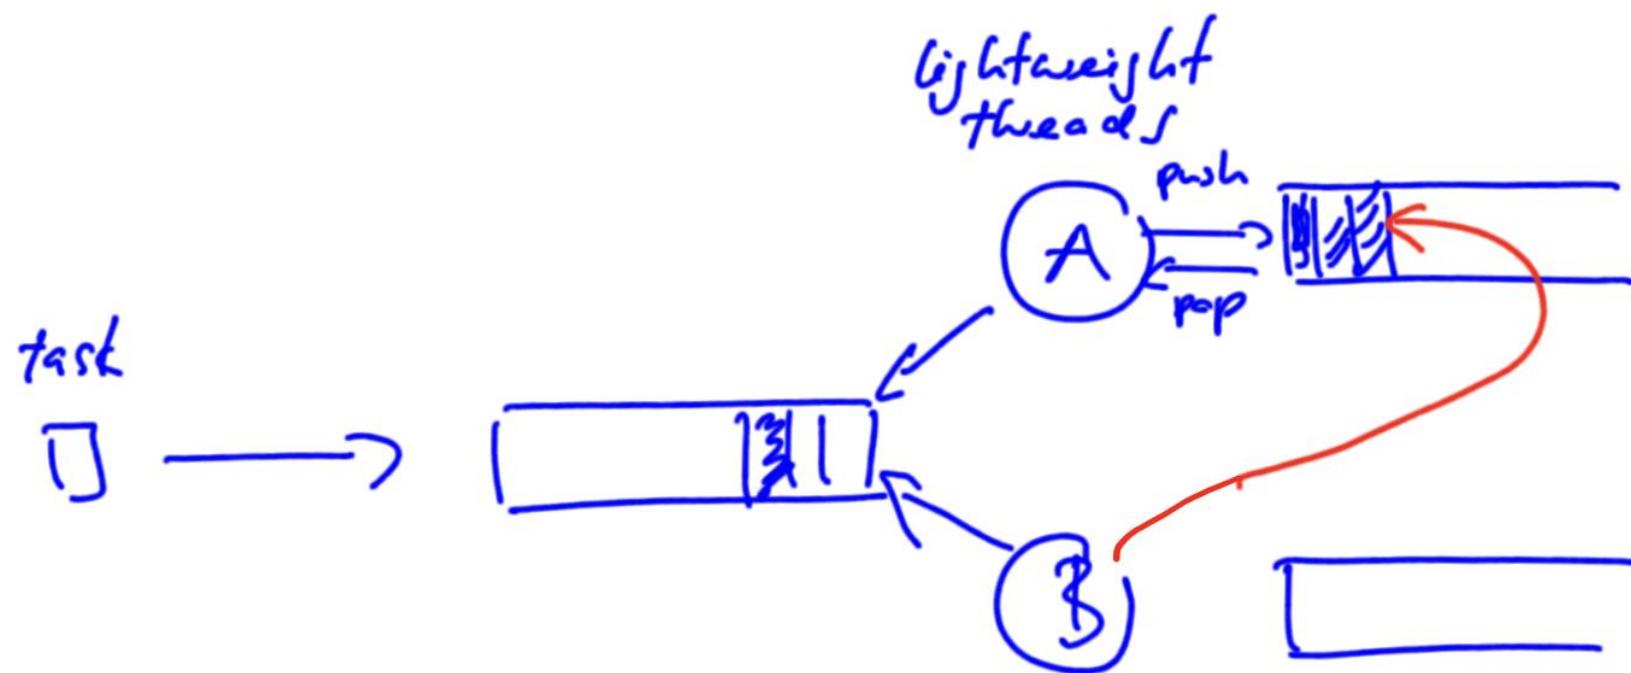
But for recursive problems (like divide-and-conquer) we need a more complex system: Work stealing

# Tasks in Fork/Join Framework



# Difference between ExecService and ForkJoin

- ExecService has one queue
- ForkJoin have work stealing
- each worker thread has its own deque (double-ended queue)
- Threads pick up tasks from their own deque
- If a thread runs out of work, it "steals" from another thread's deque



# Different terms, same basic idea

To use the ForkJoin Framework:

A little standard set-up code (e.g., create a **ForkJoinPool**)

Don't subclass **Thread**

Don't override **run**

Do not use an **ans** field

Don't call **start**

Don't just call **join**

Don't call **run** to hand-optimize

Don't have a topmost call to **run**

Do subclass **RecursiveTask<V>**

Do override **compute**

Do return a **V** from **compute**

Do call **fork**

Do call **join** which returns answer

Do call **compute** to hand-optimize

Do create a pool and call **invoke**

# Recursive sum with ForkJoin (compute)

```
protected Long compute() {  
    if(high - low <= 1)  
        return array[high];  
    else {  
        int mid = low + (high - low) / 2;  
        SumForkJoin left  = new SumForkJoin(array, low, mid);  
        SumForkJoin right = new SumForkJoin(array, mid, high);  
        left.fork();  
        right.fork();  
        return left.join() + right.join();  
    }  
}
```

# Fixes/Work-Around – cont'd

```
protected Long compute() {  
    if(high - low <= SEQUENTIAL_THRESHOLD) {  
        long sum = 0;  
        for(int i=low; i < high; ++i)  
            sum += array[i];  
        return sum;  
    } else {  
        int mid = low + (high - low) / 2;  
        SumForkJoin left = new SumForkJoin(array, low, mid);  
        SumForkJoin right = new SumForkJoin(array, mid, high);  
        left.fork();  
        long rightAns = right.compute();  
        long leftAns = left.join();  
        return leftAns + rightAns;  
    }  
}
```

Make sure each task has 'enough' to do!

# Why Fork/Join?

- one of the key reasons for using the Fork/Join Framework instead of just an ExecutorService is that tasks can be dependent on one another
- ExecutorService has no work stealing
- ExecutorService works with a fixed thread pool where each thread picks up tasks from a shared queue. If one thread finishes its work early, it remains idle while others are busy
- The Fork/Join framework uses **work-stealing**, meaning idle threads can "steal" unfinished tasks from busy threads, improving load balancing

# Why Fork/Join?

- See code examples

# Why Fork/Join?

- What happened?
- Thread Starvation:
  - Each  $\text{fib}(n)$  spawns two new tasks.
  - With only 4 threads, they quickly get exhausted.
  - Tasks are waiting on `Future.get()`, but no free thread is available to execute pending tasks.
  - Deadlock happens when all threads are blocked waiting.

# Plan für heute

- Organisation
- Nachbesprechung Exercise 4
- Theory Recap
- **Intro Exercise 5**
- Exam Questions
- Kahoot

# Assignment 5

## 1. *Parallel Search and Count*

Search an array of integers for a certain feature and count integers that have this feature.

- Light workload: count number of non-zero values.
- Heavy workload: count how many integers are prime numbers.

We will study single threaded and multi-threaded implementation of the problem.

**Explanation of given code**

# Assignment 5

## 1. *Parallel Search and Count*

Search an array of integers for a certain feature and count integers that have this feature.

- Light workload: count number of non-zero values.
- Heavy workload: count how many integers are prime numbers.

We will study single threaded and multi-threaded implementation of the problem.

## 2. *Amdahl's and Gustafson's Law*

## 3. *Amdahl's and Gustafson's Law II*

## 4. *Task Graph*

# Amdahl's & Gustafson's Law

Let  $f$  be the non-parallelizable serial fractions of the total work and  $p$  be the number of workers.

## Amdahl's Law

$$S_p \leq \frac{1}{f + \frac{1-f}{p}}$$

## Gustafson's Law

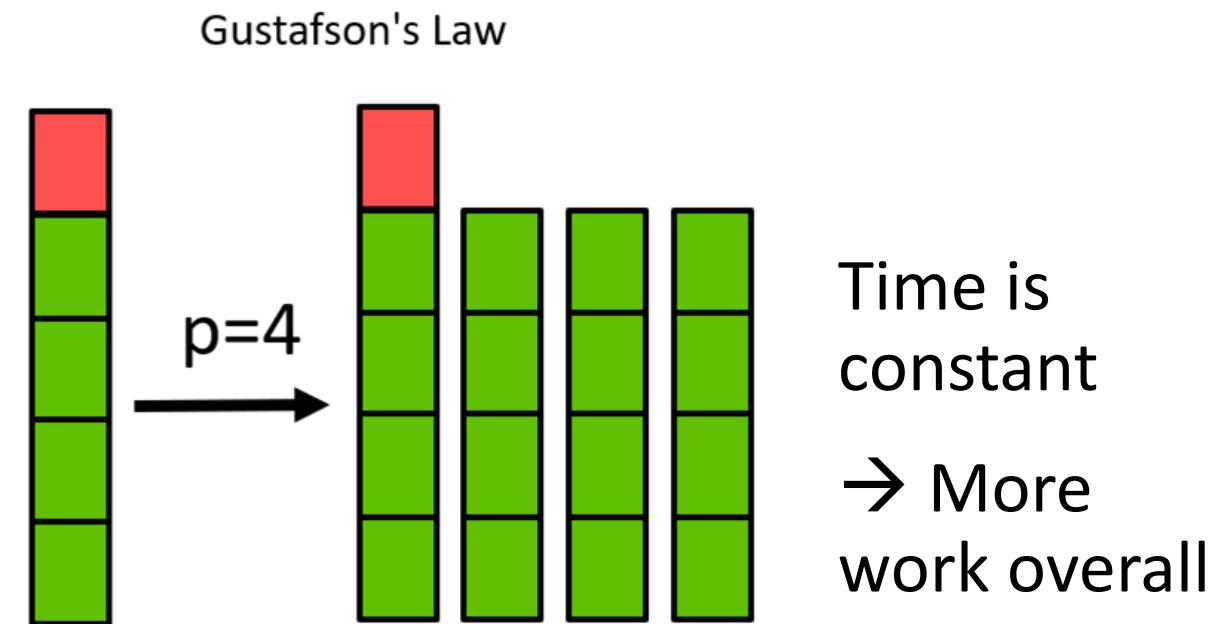
$$S_p = p - f(p - 1)$$

# Amdahl's & Gustafson's Law

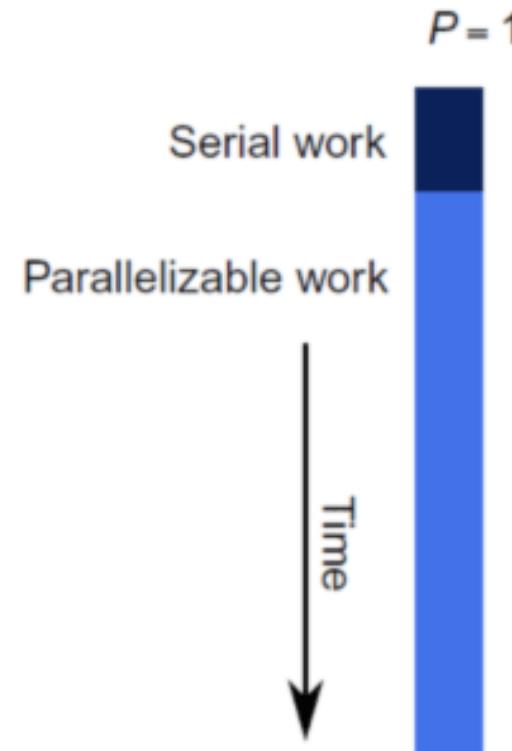
Assuming a program consists of 50% non-parallelizable code.

- a) Compute the speed-up when using 2 and 4 processors according to Amdahl's law.

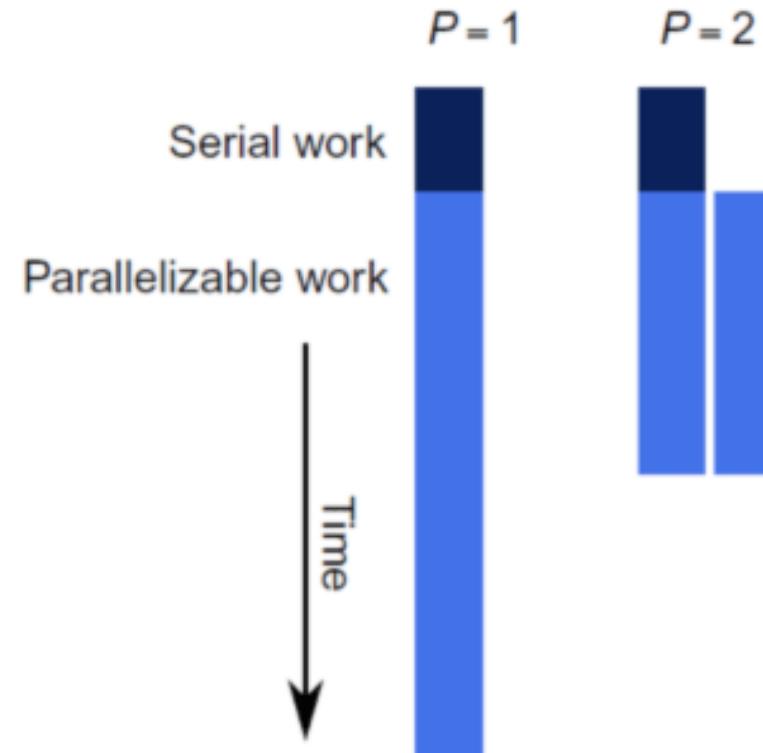
- b) Now assume that the parallel work per processor is fixed. Compute the speed-up when using 2 and 4 processors according to Gustafson's law.



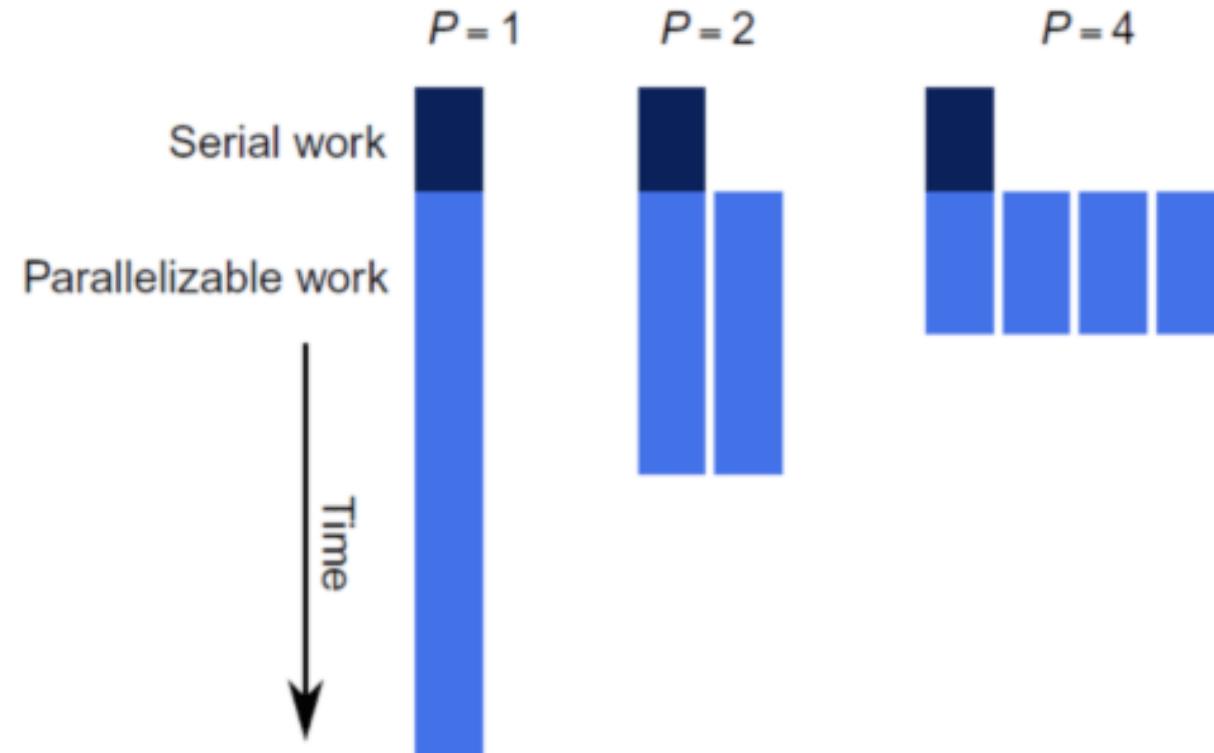
# Amdahl's Law Illustrated



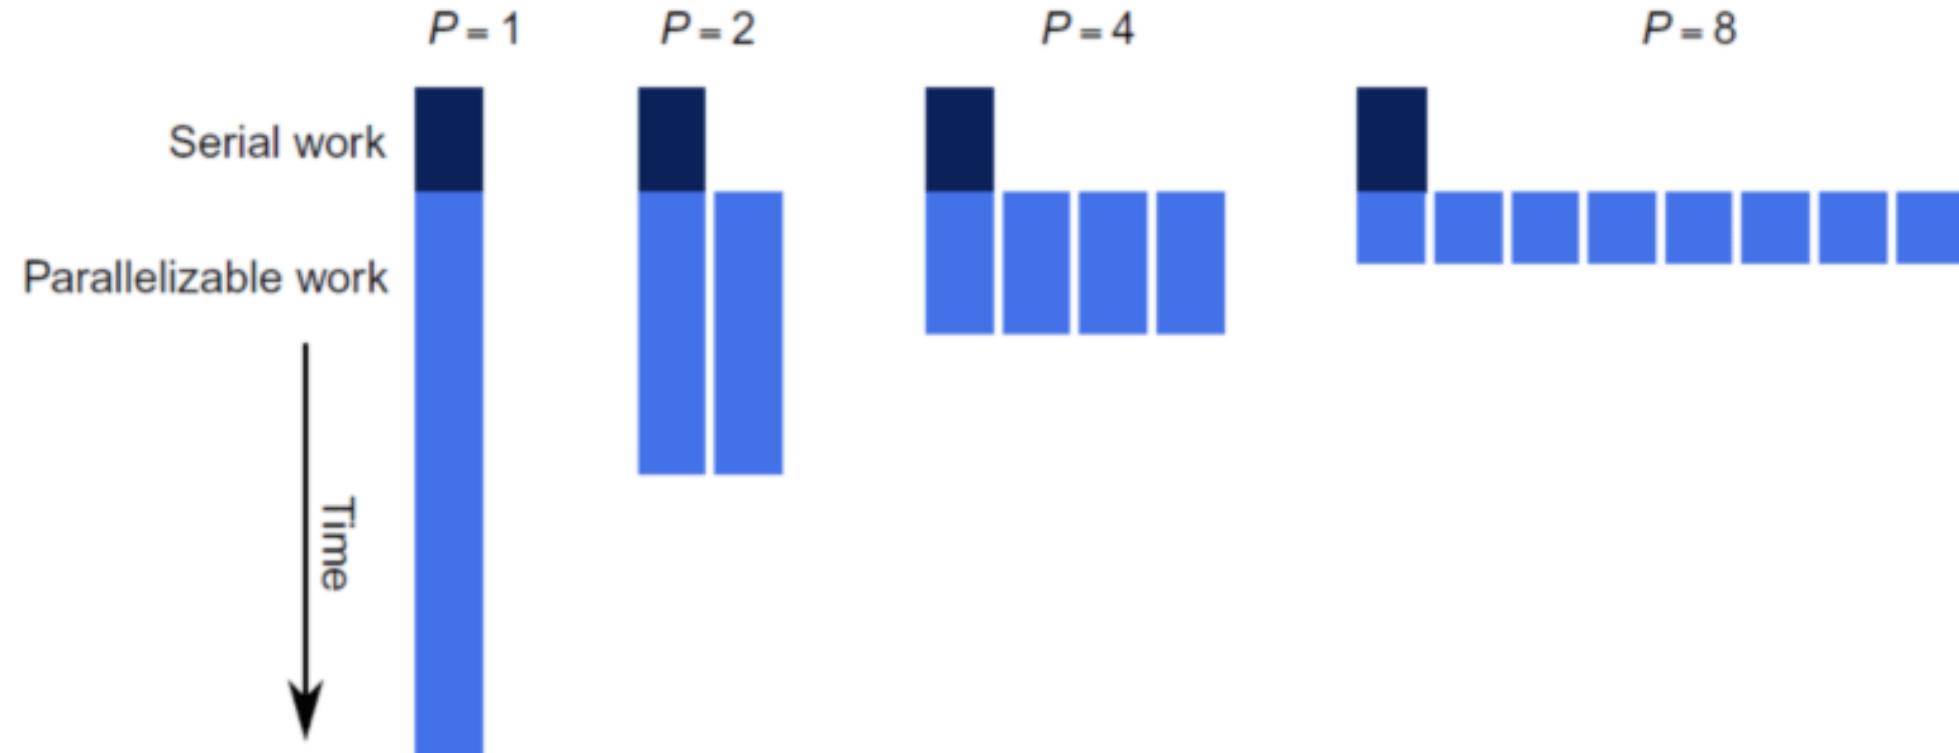
# Amdahl's Law Illustrated



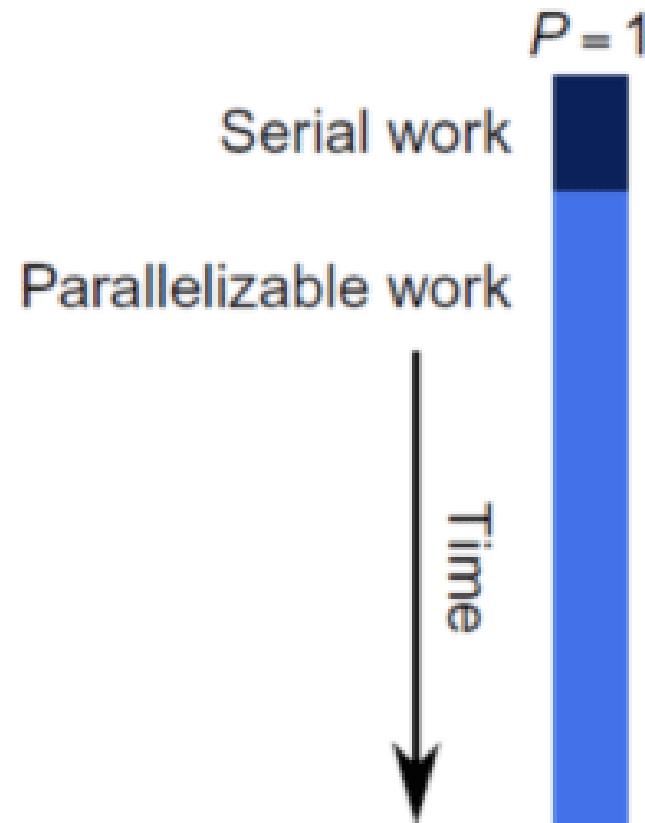
# Amdahl's Law Illustrated



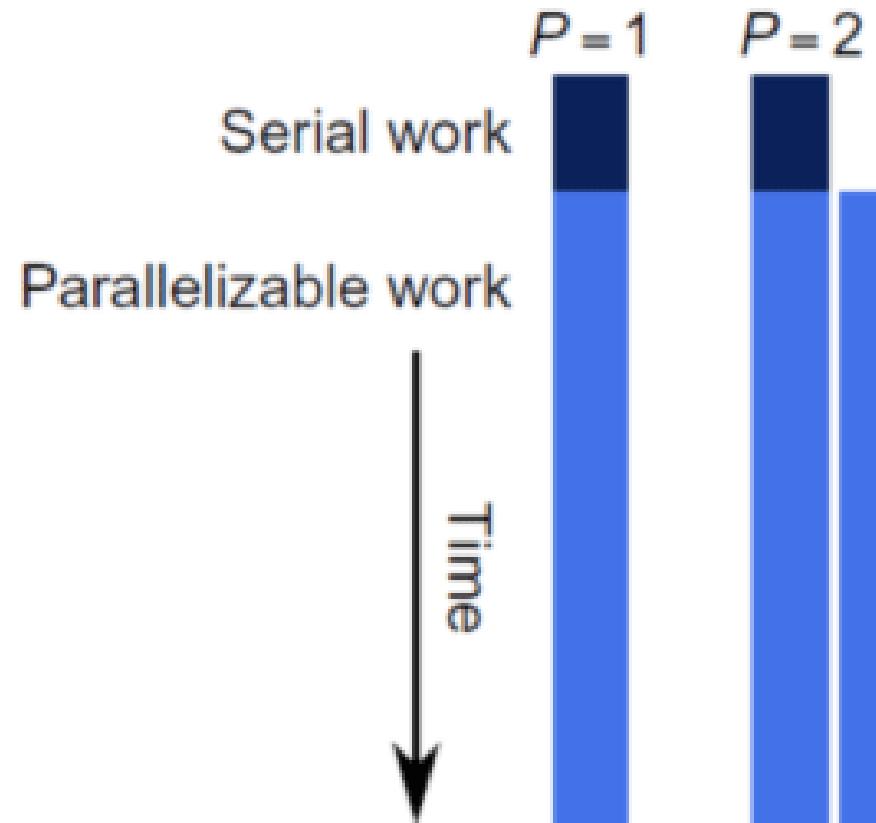
# Amdahl's Law Illustrated



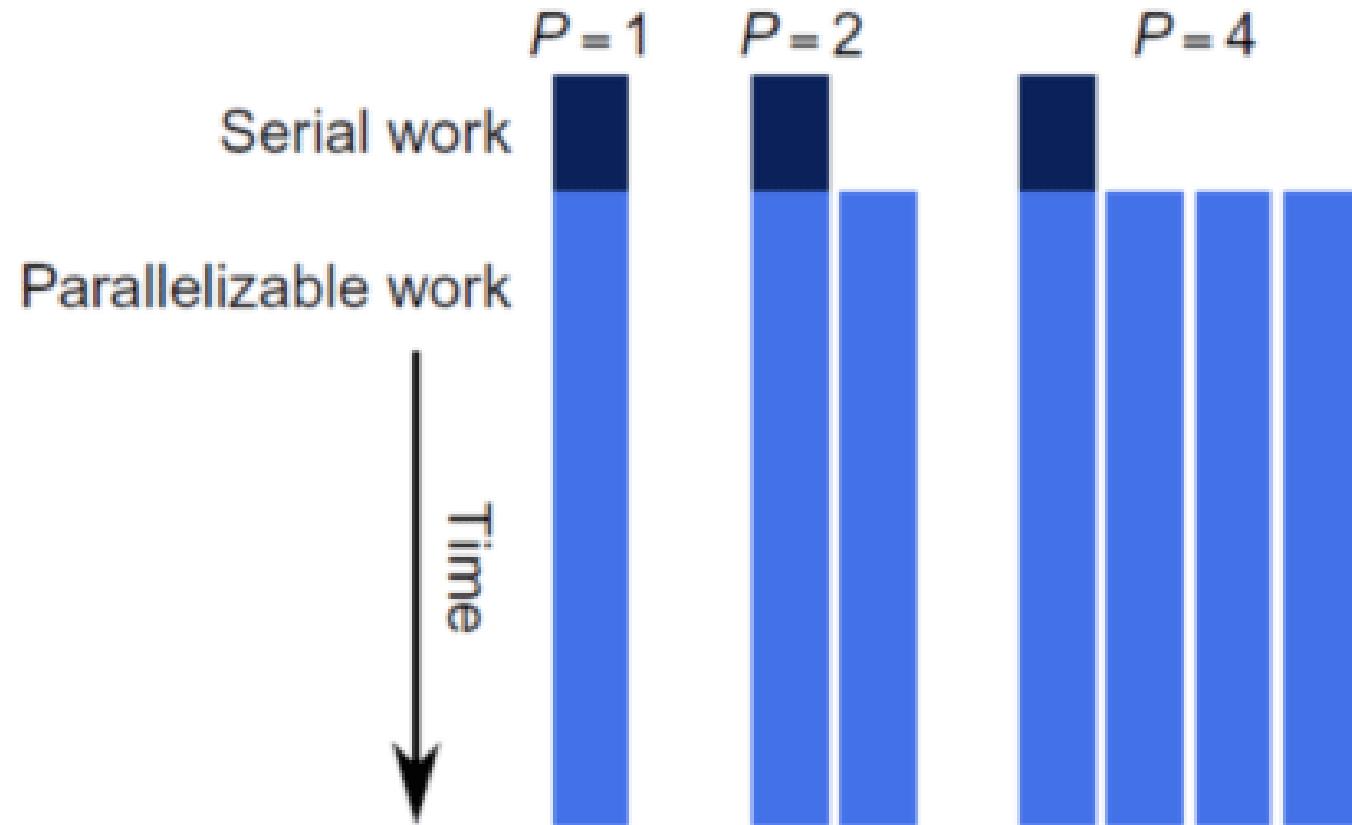
# Gustafson's Law



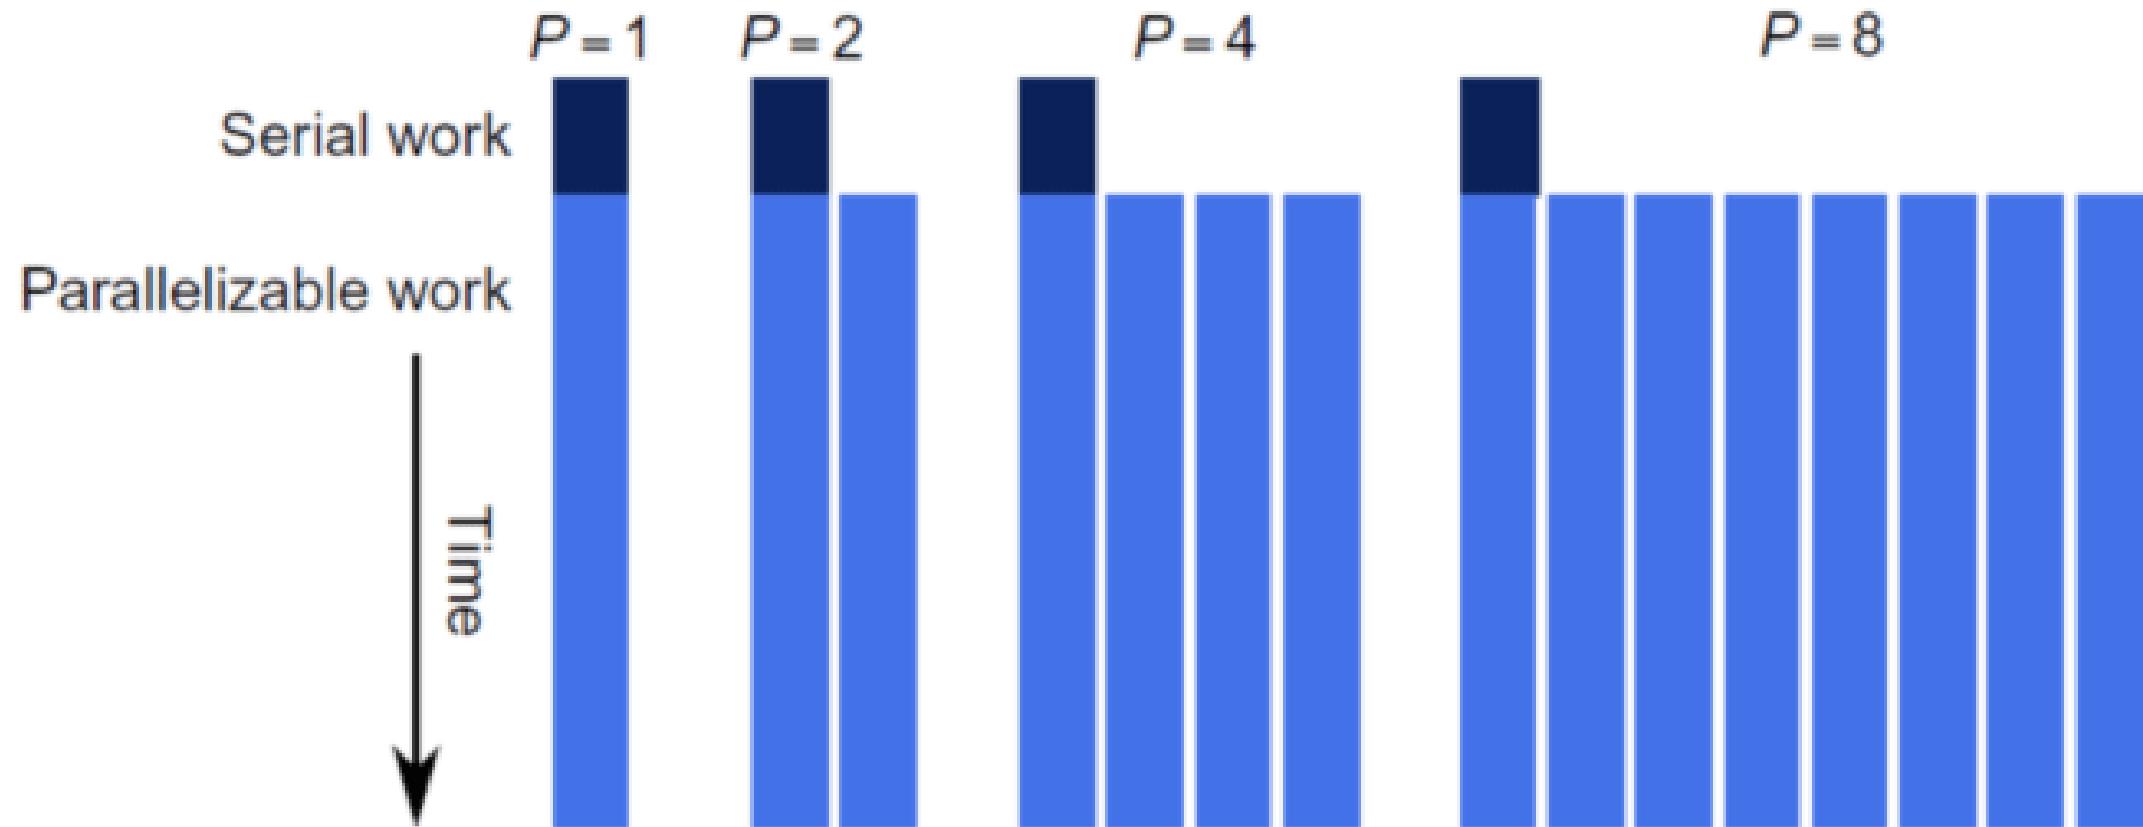
# Gustafson's Law



# Gustafson's Law



# Gustafson's Law



# Summary

$T_1$ : Time on one processor

$T_P$ : Time on  $P$  processors

$T_\infty$ : Time on "infinite" processors ( $\lim_{P \rightarrow \infty} T_P$ )

$S_P = \frac{T_1}{T_P}$ : Speedup generated using  $P$  processors

- **Amdahls Law** assumes a fixed workload in variable time it states:

$$S_P = \frac{T_1}{T_P} \leq \frac{1}{f + \frac{1-f}{P}}$$

- **Gustafsons Law** assumes a fixed time window, but measures the speedup in workload terms:

$$S_P \leq f + P(1 - f)$$

# Plan für heute

- Organisation
- Nachbesprechung Exercise 4
- Theory Recap
- Intro Exercise 5
- **Exam Questions**
- Kahoot

# Old Exam Task (HS20 – Task 1)

1. Nehmen Sie an ein Programm besteht zu 20% aus nicht-parallelisierbarer Arbeit. Wir wollen einen Speedup von 4 gemäss Gustafson's Gesetz erlangen. Wie viele Prozessoren sind notwendig? Geben Sie alle Rechenschritte an.

*Assume a program consisting of 20% non-parallelizable work. We want to achieve a speed up of 4 according to Gustafson's law. How many processors are necessary? Show all calculation steps.* (5)

Gustafson:  $S_p = p - f(p - 1)$ . From the text:

$$S_p = 4 \text{ and } f = 0.2$$

$$4 = p - 0.2(p - 1)$$

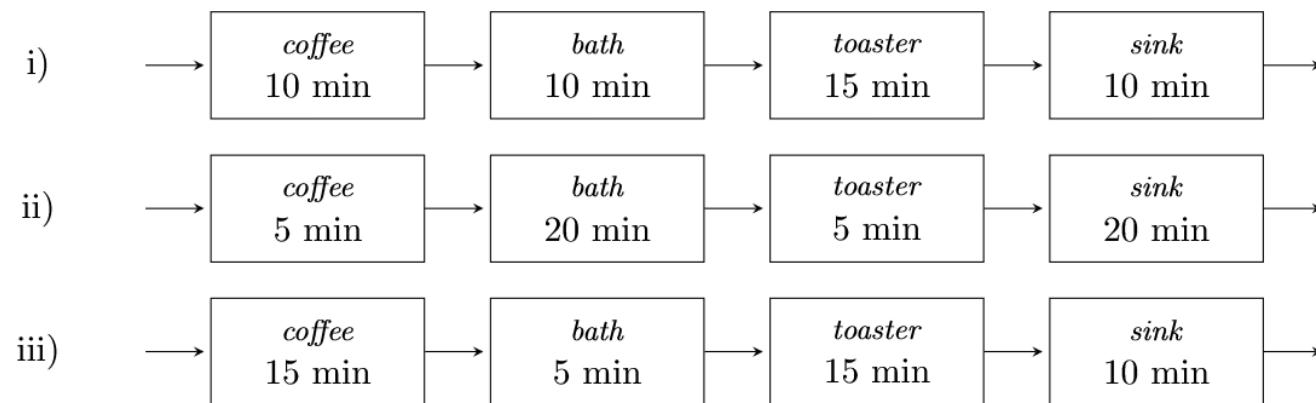
$$\Leftrightarrow 4 = 0.8p + 0.2$$

$$\Leftrightarrow 3.8 = 0.8p$$

$$\Leftrightarrow p = 4.75$$

→ 5 processors are required!

Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



- (a) Was ist der Durchsatz ("throughput") der einzelnen Pipelines pro Stunde in dieser WG?

*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

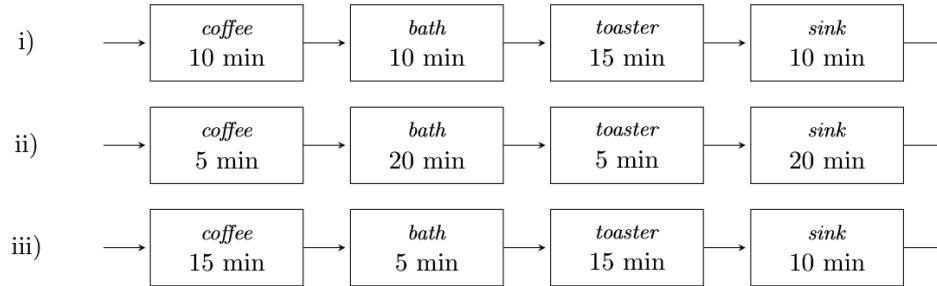
*What is the throughput of each pipeline (3) per hour in this living community?*

i)

ii)

iii)

Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

(a) Was ist der Durchsatz ("throughput") der einzelnen Pipelines pro Stunde in dieser WG?

*What is the throughput of each pipeline (3) per hour in this living community?*

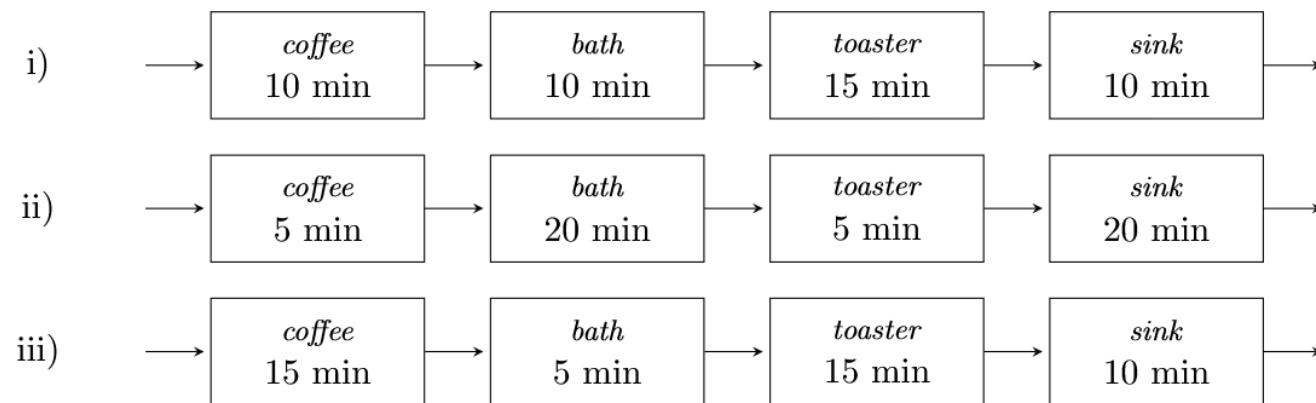
i)

ii)

iii)

- 1.) 1 Student /15 Minuten = 4 Studenten/h
- 2.) 1 Student/20Minuten = 3 Studenten/h
- 3.) 1 Student /15 Minuten = 4 Studenten/h

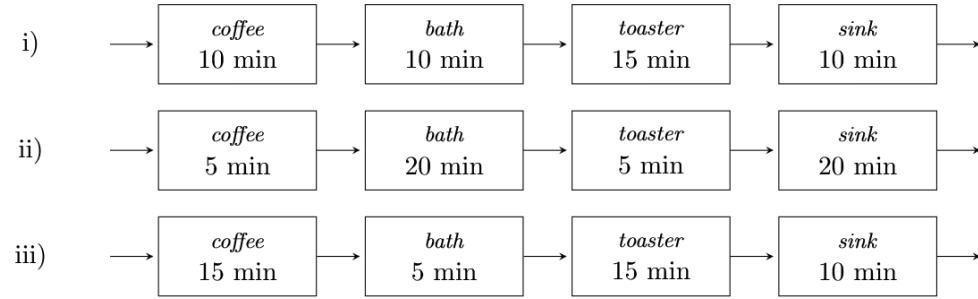
Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

- (b) Das Ziel ist, dass alle drei Studenten so schnell wie möglich fertig sind. Rechnen sie die Zeit aus, die es benötigt bis alle drei Studenten die Pipelines durchlaufen haben. Begründen Sie Ihre Antwort und geben Sie alle Rechenschritte an. Welche Pipeline(s) ist (sind) die schnellste(n)?

Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

- (b) Das Ziel ist, dass alle drei Studenten so schnell wie möglich fertig sind. Rechnen sie die Zeit aus, die es benötigt bis alle drei Studenten die Pipelines durchlaufen haben. Begründen Sie Ihre Antwort und geben Sie alle Rechenschritte an. Welche Pipeline(s) ist (sind) die schnellste(n)?

Zuerst berechnen wir die Zeit, welche die erste Person benötigt um fertig zu sein:

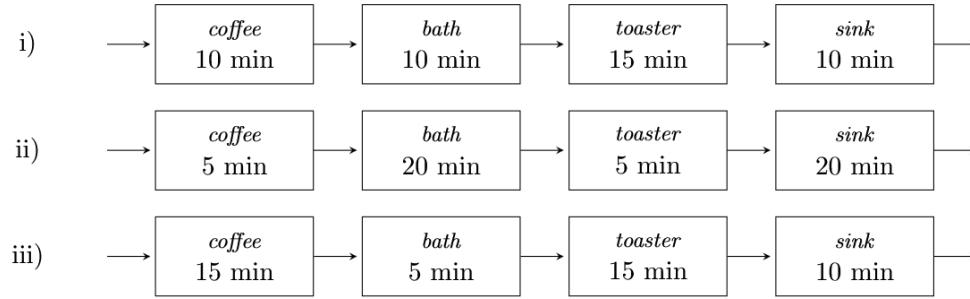
$$\begin{aligned}
 1. & 10 + 10 + 15 + 10 = 45 \\
 2. & 5 + 20 + 5 + 20 = 50 \\
 3. & 15 + 5 + 15 + 10 = 45
 \end{aligned}$$

Für jeden zusätzlichen Student, benötigen wir nun die längste Station (1 -> 15', 2 -> 20', 3 -> 15') mehr. Also:

$$\begin{aligned}
 1. & 45 + 2 * 15 = 75 \\
 2. & 50 + 2 * 20 = 90 \\
 3. & 45 + 2 * 15 = 75
 \end{aligned}$$

Somit sehen wir das (1) und (3) die schnellsten Pipelines sind.

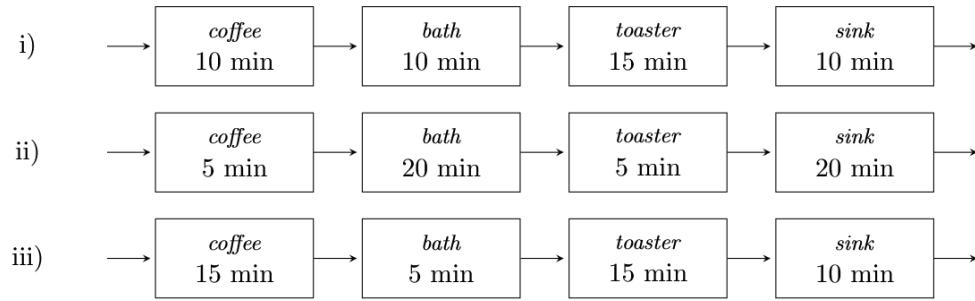
Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

- (d) Aufgrund von einer Pandemie ist es den Studenten leider nicht erlaubt eines der 4 gemeinsamen Objekte zu benutzen bevor der vorhergehende Student komplett fertig ist (d.h. mit dem "sink" Schritt abgeschlossen hat). Die Studenten haben jedoch die Möglichkeit, jeden Schritt um 10 Minuten zu verlängern um die Objekte zu desinfizieren, damit die nächste Person das Objekt unmittelbar nach der Desinfektion verwenden kann. Welche Pipeline ist am schnellsten mit dieser neuen Regelung und lohnt es sich für die Studenten diese zu implementieren? Aus Solidarität muss der letzte Student ebenso alles desinfizieren. Begründen Sie Ihre Antwort und geben Sie alle Rechenschritte an.

Pipelines. Nehmen Sie für alle nachfolgenden Fragen an, dass die drei Studenten kein Objekt gleichzeitig benutzen wollen. Die folgenden drei Zeilen zeigen drei mögliche separate Pipelines:



*For all of the following question, assume that the three students do not want to use an object at the same time. The following three lines show three possible separate pipelines:*

- (d) Aufgrund von einer Pandemie ist es den Studenten leider nicht erlaubt eines der 4 gemeinsamen Objekte zu benutzen bevor der vorhergehende Student komplett fertig ist (d.h. mit dem "sink" Schritt abgeschlossen hat). Die Studenten haben jedoch die Möglichkeit, jeden Schritt um 10 Minuten zu verlängern um die Objekte zu desinfizieren, damit die nächste Person das Objekt unmittelbar nach der Desinfektion verwenden kann. Welche Pipeline ist am schnellsten mit dieser neuen Regelung und lohnt es sich für die Studenten diese zu implementieren? Aus Solidarität muss der letzte Student ebenso alles desinfizieren. Begründen Sie Ihre Antwort und geben Sie alle Rechenschritte an.

Wenn wir alles sequentiell abarbeiten:

$$1. 10 + 10 + 15 + 10 = 45 \rightarrow 45 * 3 = 135$$

$$2. 5 + 20 + 5 + 20 = 50 \rightarrow 50 * 3 = 150$$

$$3. 15 + 5 + 15 + 10 = 45 \rightarrow 45 * 3 = 135$$

So sehen die Pipelines mit Desinfektion aus:

$$1. 20 + 20 + 25 + 20 = 85 \rightarrow 85 + 2 * 25 = 135$$

$$2. 15 + 30 + 15 + 30 = 90 \rightarrow 90 + 2 * 30 = 150$$

$$3. 25 + 15 + 25 + 20 = 85 \rightarrow 85 + 2 * 25 = 135$$

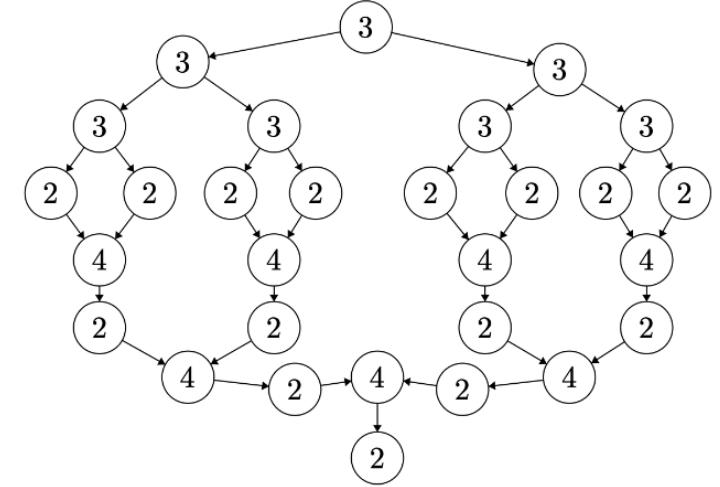
Da sich die Zeiten nicht unterscheiden, spielt es keine Rolle, ob die Regel implementiert ist oder nicht.

# Exam preparation (old exam tasks)

Thanks to Erxuan Li  
PProg FS25

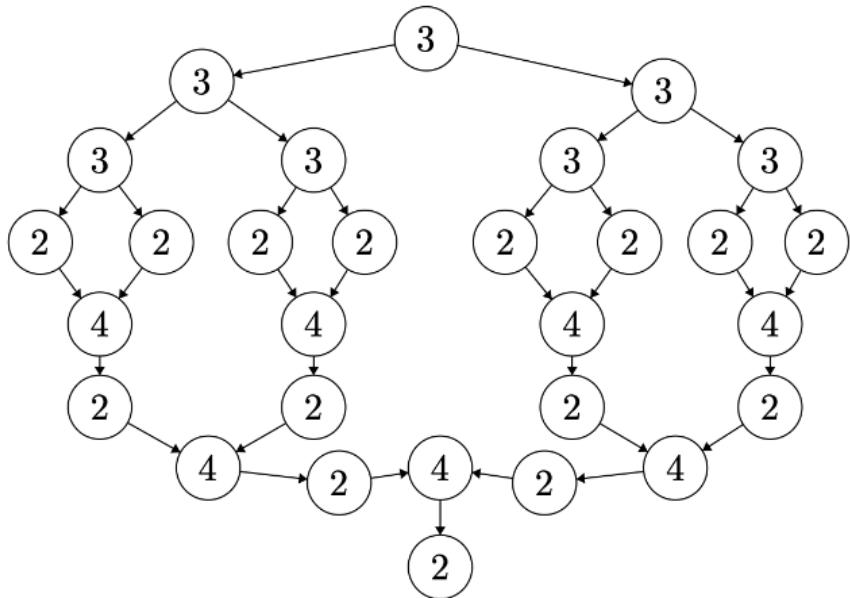
# Task Graph

- Way to model program execution
- DAG: directed acyclic graph
- Nodes are tasks / work labelled with execution time
- Edges show dependencies → data dependency along the edge



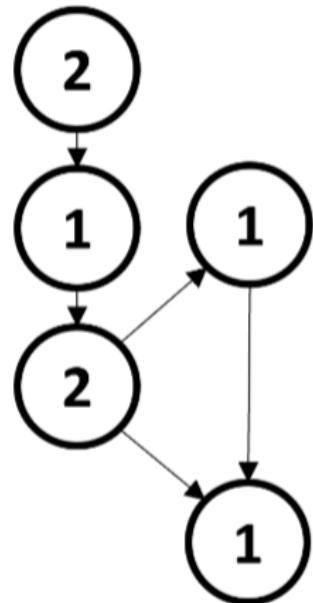
# Task Graph

- Used to calculate speed up
- $T_1 = \text{sum of all nodes}$
- $T_p = \text{sum of all nodes on critical (= longest) path}$



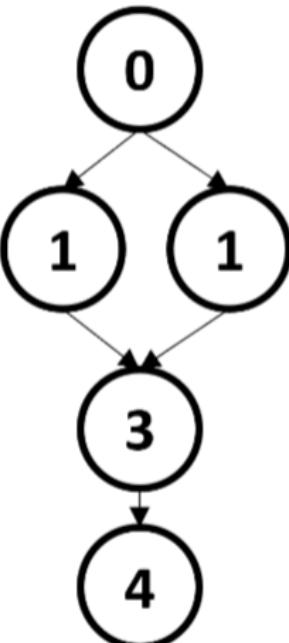
(e) Kreisen Sie in den folgenden Fragen die richtige Antwort ein.

- i. Welcher der folgenden Task Graphen ist **kein** gültiger Task Graph?



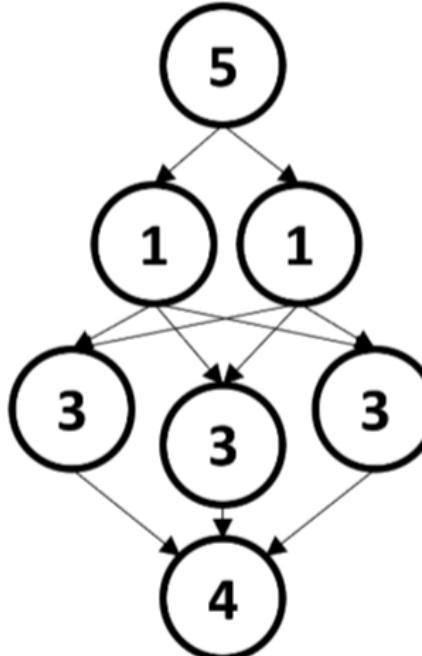
**A**

(A)



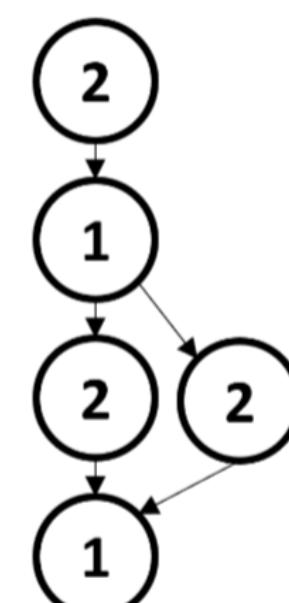
**B**

(B)



**C**

(C)



**D**

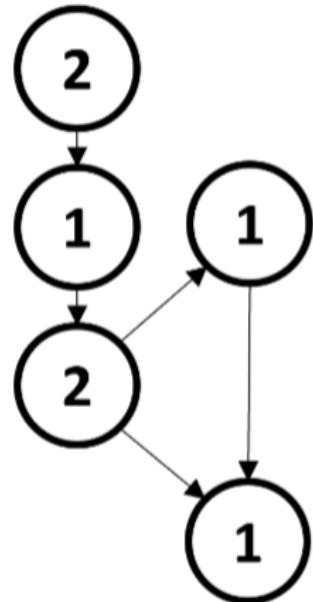
(D) (E) None of the above

*In the questions below, circle the correct answer.*

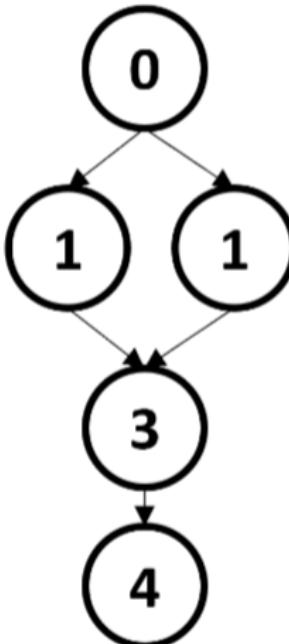
*Which of the following task graphs is (1) not a valid task graph?*

(e) Kreisen Sie in den folgenden Fragen die richtige Antwort ein.

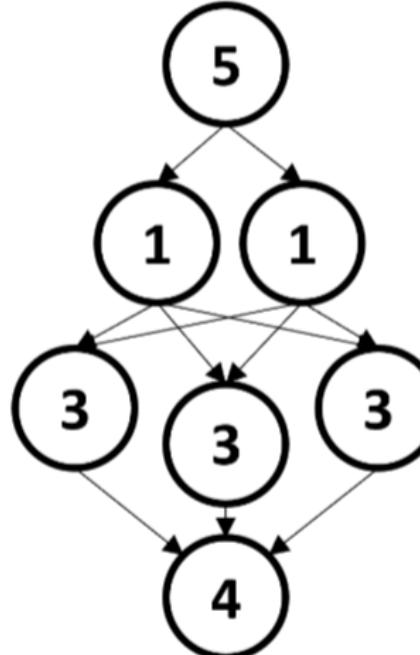
- i. Welcher der folgenden Task Graphen ist **kein** gültiger Task Graph?



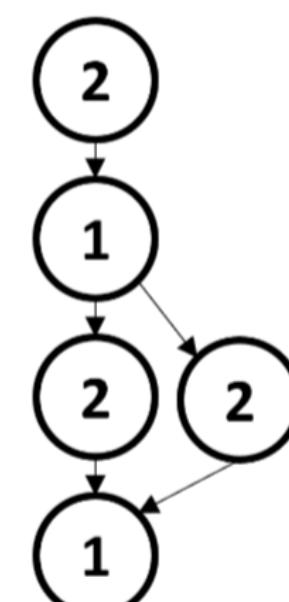
A



B



C



D

(A)

(B)

(C)

(D)

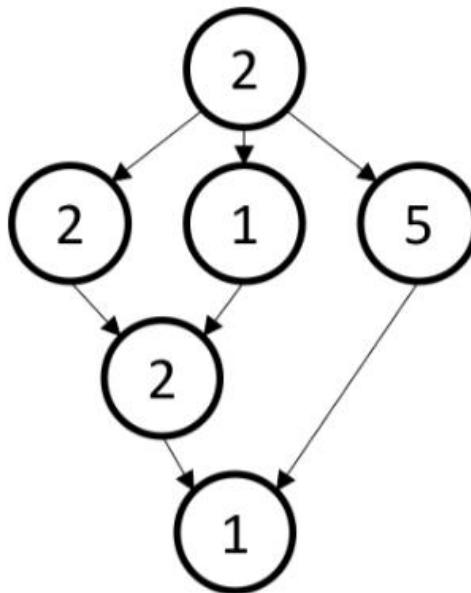
(E) None of the above

In the questions below, circle the correct answer.

Which of the following task graphs is (1) **not** a valid task graph?

ii. Was ist die Länge des kritischen Pfades des folgenden Task Graphen?

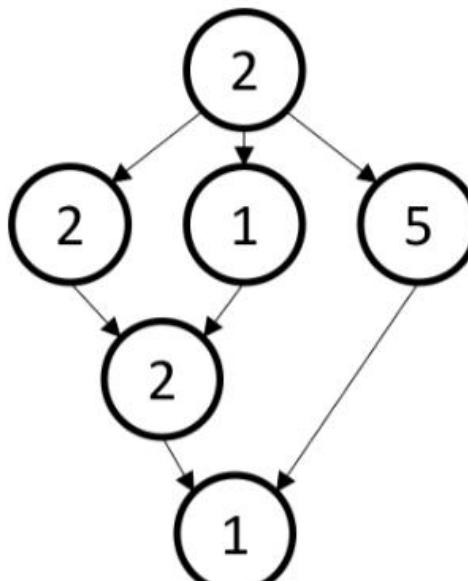
*What is the length of the critical path of the following task graph?*



- (A) 3      (B) 6      (C) 8      (D) 10      (E) 13

ii. Was ist die Länge des kritischen Pfades des folgenden Task Graphen?

*What is the length of the critical path of the following task graph?*



- (A) 3      (B) 6      (C) 8      (D) 10      (E) 13

 (C) 8

iii. Welcher Ausdruck gibt eine Untergrenze für  $T_p$  an, die Zeit, die zum Ausführen eines Task Graphen auf  $p$  Prozessoren benötigt wird?

*Which expression gives an lower bound on  $T_p$ , the time it takes to execute a task graph on  $p$  processors?*

(A)  $\frac{T_1}{p}$

(B)  $T_1 \cdot p$

(C)  $\frac{T_1 \cdot (p-1)}{p}$

(D)  $\frac{p-1}{T_1}$

(E)  $2 \cdot T_1$

---

iii. Welcher Ausdruck gibt eine Untergrenze für  $T_p$  an, die Zeit, die zum Ausführen eines Task Graphen auf  $p$  Prozessoren benötigt wird?

Which expression gives an lower bound on  $T_p$ , the time it takes to execute a task graph on  $p$  processors?

(A)  $\frac{T_1}{p}$

(B)  $T_1 \cdot p$

(C)  $\frac{T_1 \cdot (p-1)}{p}$

(D)  $\frac{p-1}{T_1}$

(E)  $2 \cdot T_1$

## Execution time $T_p$ on $p$ CPUs

- $T_p = T_1 / p$  (perfection)
- $T_p > T_1 / p$  (performance loss, what normally happens)
- $T_p < T_1 / p$  (sorcery!)

1. Nehmen Sie an ein Programm besteht zu 20% aus nicht-parallelisierbarer Arbeit. Wir wollen einen Speedup von 4 gemäss Gustafson's Gesetz erlangen. Wie viele Prozessoren sind notwendig? Geben Sie alle Rechenschritte an.

*Assume a program consisting of 20% non-parallelizable work. We want to achieve a speed up of 4 according to Gustafson's law. How many processors are necessary? Show all calculation steps.* (5)

Gustafson:  $S_p = p - f(p - 1)$ . From the text:

$$S_p = 4 \text{ and } f = 0.2$$

1. Nehmen Sie an ein Programm besteht zu 20% aus nicht-parallelisierbarer Arbeit. Wir wollen einen Speedup von 4 gemäss Gustafson's Gesetz erlangen. Wie viele Prozessoren sind notwendig? Geben Sie alle Rechenschritte an.

*Assume a program consisting of 20% non-parallelizable work. We want to achieve a speed up of 4 according to Gustafson's law. How many processors are necessary? Show all calculation steps.* (5)

Gustafson:  $S_p = p - f(p - 1)$ . From the text:

$$S_p = 4 \text{ and } f = 0.2$$

$$\begin{aligned} 4 &= p - 0.2(p - 1) \\ \Leftrightarrow 4 &= 0.8p + 0.2 \\ \Leftrightarrow 3.8 &= 0.8p \\ \Leftrightarrow p &= 4.75 \end{aligned}$$

→ 5 processors are required!

## Pipelines and Speedup (15 points)

2. (a) i. Welcher Speedup kann bei einem Task mit einem sequentiellen Anteil von 0.16 erreicht werden, wenn der Task gemäss dem Amdahl-Gesetz auf  $p = 126$  Prozessoren ausgeführt wird? *Hinweis:*  $126 \times 16 = 2016$

*Given a task with a sequential fraction (3) of 0.16, what speedup can be achieved when running the task on  $p = 126$  processors according to Amdahl's law?  
Hint:  $126 \times 16 = 2016$*

## Pipelines and Speedup (15 points)

2. (a) i. Welcher Speedup kann bei einem Task mit einem sequentiellen Anteil von 0.16 erreicht werden, wenn der Task gemäss dem Amdahl-Gesetz auf  $p = 126$  Prozessoren ausgeführt wird? *Hinweis:*  $126 \times 16 = 2016$

*Given a task with a sequential fraction (3) of 0.16, what speedup can be achieved when running the task on  $p = 126$  processors according to Amdahl's law?  
Hint:  $126 \times 16 = 2016$*

---

Tobias Steinbrecher @tsteinbreche · 8 months ago

▼ 11 ▲

---

We have

$$f = 0.16 := \text{sequential part of the program}$$

With Amdahl's law we obtain:

$$S_{126} = \frac{1}{0.16 + 0.84/126} = \frac{126}{20.16 + 0.84} = \frac{126}{21} = 6$$

+ Add Comment    ... More

ii. Kann eine solcher Speedup in der Praxis auf einem realen System erreicht werden? Wenn nicht, was verhindert den optimalen Speedup und unter welchen Umständen kann ein Speedup erreicht werden, der nahe am theoretischen Wert liegt?

*Can such a speedup be obtained in (3) practice on a real system? If not, what hinders the optimal speedup and under what circumstances can a speedup be achieved that is close to the theoretical value?*

- ii. Kann eine solcher Speedup in der Praxis auf einem realen System erreicht werden? Wenn nicht, was verhindert den optimalen Speedup und unter welchen Umständen kann ein Speedup erreicht werden, der nahe am theoretischen Wert liegt?

*Can such a speedup be obtained in (3) practice on a real system? If not, what hinders the optimal speedup and under what circumstances can a speedup be achieved that is close to the theoretical value?*

**Sven Glinz** @sglinz · 8 months ago

▼ 12 ▲

---

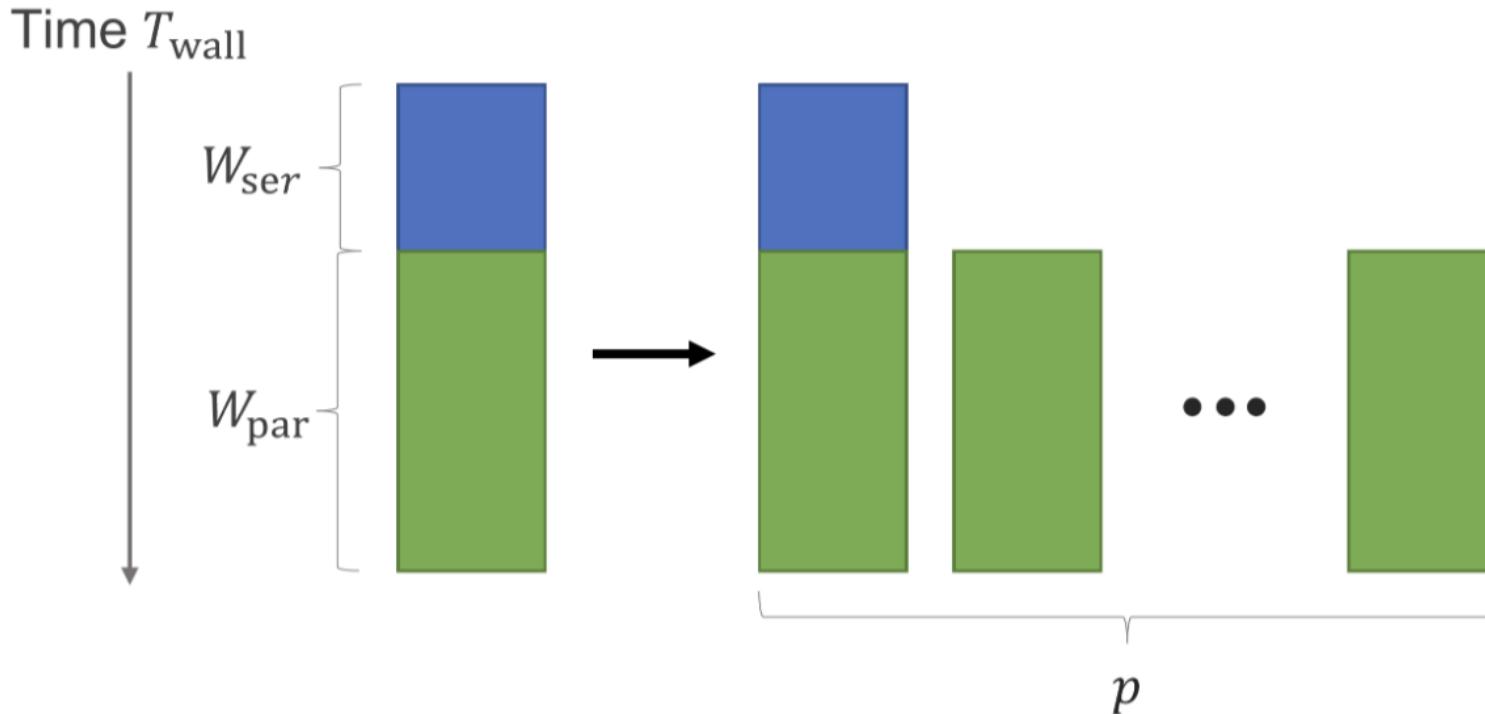
such a speedup could be reached if the needed communication between the individual threads is very low (eg. no locks, no need for cache coherency), each thread can optimally use the CPU and the number of threads is close to the number of logical cores such that only few context switches take place.

In practice, synchronization between threads, memory contention (multiple threads trying to access memory at the same time), load imbalance between different threads etc. could prevent an optimal speedup.

+ Add Comment    ... More

- iii. Ist Amdahl's Gesetz oder Gustafson's Gesetz in der folgenden Abbildung dargestellt? Benennen Sie das Gesetz, definieren Sie den seriellen Anteil  $f$  anhand der seriellen und parallelen Anteil der Arbeit,  $W_{ser}$  und  $W_{par}$ , und leiten Sie daraus die Formel für den Speedup  $S_p$  ab.

*Is Amdahl's Law or Gustafson's Law (4) demonstrated in the illustration below? Name the law, define the serial fraction  $f$  in terms of the serial and parallelizable fractions of the work,  $W_{ser}$  and  $W_{par}$ , and use it to derive the formula for speed up  $S_p$ .*



Tobias Steinbrecher @tsteinbreche · 8 months ago · edited 8 months ago

▼ 16 ▲

---

### Gustafson's Law

---

We have

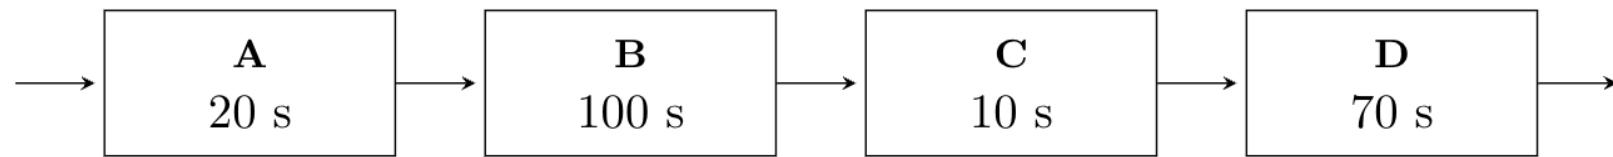
$$f = \frac{W_{ser}}{W_{ser} + W_{par}} := \text{sequential part}$$

Thus, for the speedup we obtain:

$$S_p = \frac{W_{ser} + pW_{par}}{W_{ser} + W_{par}} = \frac{W_{ser}}{W_{ser} + W_{par}} + \frac{pW_{par}}{W_{ser} + W_{par}} = f + p \cdot \frac{W_{par} + W_{ser} - W_{ser}}{W_{ser} + W_{par}} = f + p(1 - f)$$

+ Add Comment ⋮ More

- (b) Betrachten Sie die unten gezeigte Pipeline. Es gibt mehrere parallele Eingänge und A, B, C und D sind unterschiedliche Funktions-einheiten.

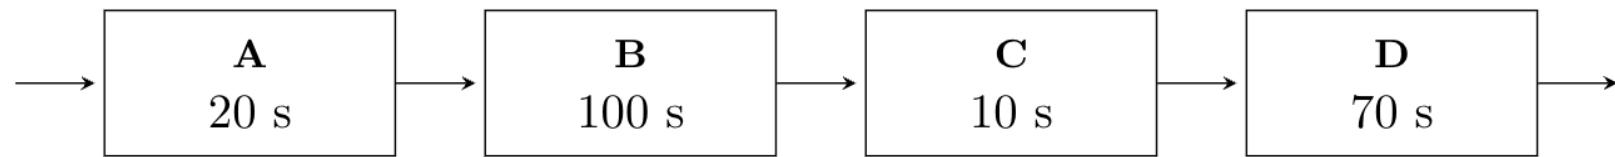


- i. Ist die Pipeline balanced? Begründen Sie Ihre Antwort. Sie müssen keine Zahl oder Berechnung angeben.

Consider the pipeline shown below. There are multiple parallel inputs and A, B, C, and D are different functional units.

Is the pipeline balanced? Justify your (2) answer. You do not need to provide a number or calculation.

- (b) Betrachten Sie die unten gezeigte Pipeline. Es gibt mehrere parallele Eingänge und A, B, C und D sind unterschiedliche Funktions-einheiten.



- i. Ist die Pipeline balanced? Begründen Sie Ihre Antwort. Sie müssen keine Zahl oder Berechnung angeben.

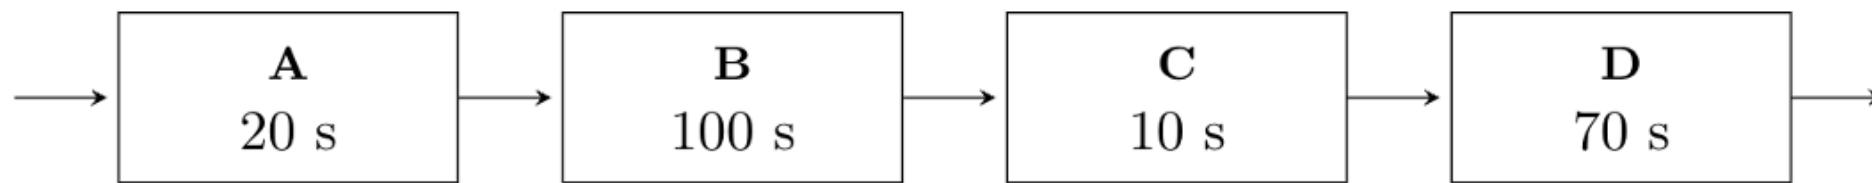
Consider the pipeline shown below. There are multiple parallel inputs and A, B, C, and D are different functional units.

Is the pipeline balanced? Justify your (2) answer. You do not need to provide a number or calculation.

No, since the stages don't take the same amount of time

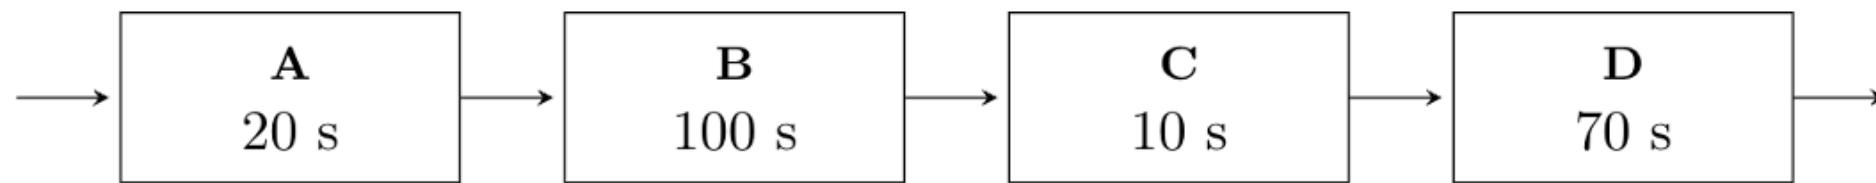
- ii. Wie lange wird es dauern, die gesamte Pipeline (d.h. die Stufen A-B-C-D) 15 Mal auszuführen? Gehen Sie davon aus, dass wir die Pipeline so parallelisieren, dass jede Stufe parallel zu den anderen Stufen laufen kann, jedoch nicht zu sich selbst. Achten Sie darauf, Ihre Antwort in den entsprechenden Einheiten anzugeben.

*How long will it take to run the whole pipeline (i.e., stages A-B-C-D) 15 times? Assume that we parallelize the pipeline such that each stage can run in parallel with the other stages, but not with itself. Make sure to leave your answer in the appropriate units.* (3)



- ii. Wie lange wird es dauern, die gesamte Pipeline (d.h. die Stufen A-B-C-D) 15 Mal auszuführen? Gehen Sie davon aus, dass wir die Pipeline so parallelisieren, dass jede Stufe parallel zu den anderen Stufen laufen kann, jedoch nicht zu sich selbst. Achten Sie darauf, Ihre Antwort in den entsprechenden Einheiten anzugeben.

*How long will it take to run the whole pipeline (i.e., stages A-B-C-D) 15 times? Assume that we parallelize the pipeline such that each stage can run in parallel with the other stages, but not with itself. Make sure to leave your answer in the appropriate units.*



$$t = 15 * (100\text{s}) + 20\text{s} + 10\text{s} + 70\text{s} = 1600\text{s}$$

```

1 public class Main {
2     public static Thread CreateThread(int start) {
3         return new Thread(new Runnable() {
4
5             @Override
6             public void run() {
7                 for (int i = start; i < 7; i+=2) {
8                     System.out.println("Number " + i);
9                 }
10            }
11        });
12    }
13
14    public static void main(String[] args) throws InterruptedException {
15        CreateThread(1).start();
16        CreateThread(2).start();
17    }
18 }
```

Markieren Sie alle Ausgaben, welche durch den Codeausschnitt ausgegeben werden können.

*Mark all the print sequences that can be produced by running the program shown above.*





1 Number 1  
2 Number 2  
3 Number 3  
4 Number 4  
5 Number 5  
6 Number 6

1 Number 1  
2 Number 6  
3 Number 3  
4 Number 4  
5 Number 5  
6 Number 2

1 Number 6  
2 Number 5  
3 Number 4  
4 Number 3  
5 Number 2  
6 Number 1

1 Number 2  
2 Number 4  
3 Number 6  
4 Number 1  
5 Number 3  
6 Number 5

```

1 public class Main {
2     public static Thread CreateThread(int start) {
3         return new Thread(new Runnable() {
4
5             @Override
6             public void run() {
7                 for (int i = start; i < 7; i+=2) {
8                     System.out.println("Number " + i);
9                 }
10            }
11        });
12    }
13
14    public static void main(String[] args) throws InterruptedException {
15        CreateThread(1).start();
16        CreateThread(2).start();
17    }
18 }
```

Markieren Sie alle Ausgaben, welche durch den Codeausschnitt ausgegeben werden können.

- 1 Number 1
- 2 Number 2
- 3 Number 3
- 4 Number 4
- 5 Number 5
- 6 Number 6

- 1 Number 1
- 2 Number 6
- 3 Number 3
- 4 Number 4
- 5 Number 5
- 6 Number 2

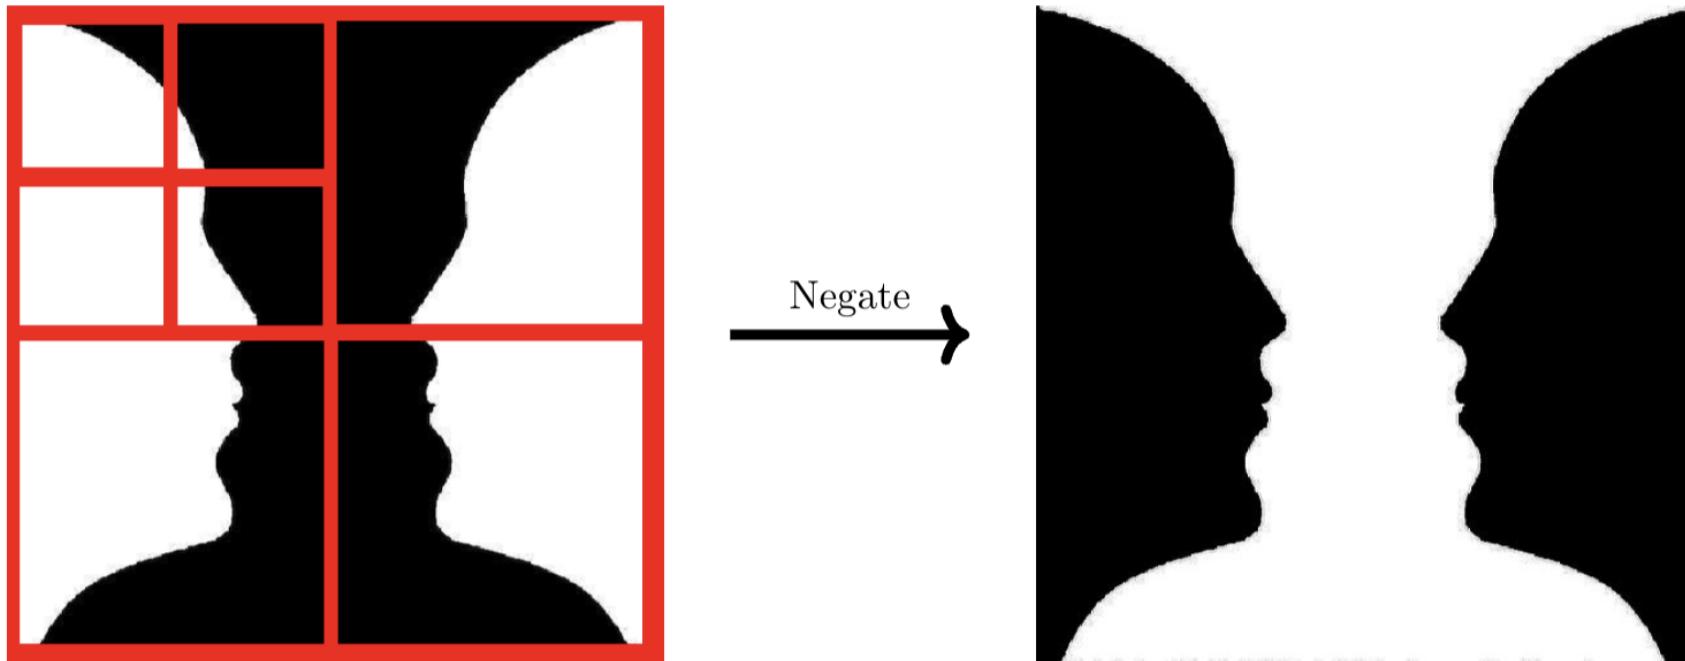
- 1 Number 6
- 2 Number 5
- 3 Number 4
- 4 Number 3
- 5 Number 2
- 6 Number 1

- 1 Number 2
- 2 Number 4
- 3 Number 6
- 4 Number 1
- 5 Number 3
- 6 Number 5

*Mark all the print sequences that can be produced by running the program shown above.*

## Fork/Join Framework (16 points)

3. Der folgende Code zielt darauf ab, ein Bild zu negieren, indem es mithilfe des Fork/Join-Frameworks rekursiv in mehrere Unterfenster (vier pro Rekursionsschritt) unterteilt wird. Die Unterfenster können dann parallel negiert werden. Das folgende Beispiel verdeutlicht die Unterteilung des Bildes und die Negierung der einzelnen Unterfenster.



Bitte lesen Sie den Code sorgfältig durch und beantworten Sie dann die Fragen zum Code:

*The following code aims to negate an image by recursively subdividing it into multiple subwindows (four per recursion step) using the Fork/Join framework. The subwindows can then be negated in parallel. The example below illustrates the subdivision of the image and negation of the individual subwindows.*

*Please read the code carefully and then answer the questions regarding the code:*

```
public class ImageNegationFJ extends RecursiveAction {
    final static int CUTOFF = 32;
    double[][] image, invertedImage;
    int startx, starty;
    int length;

    public ImageNegationFJ(double[][] image, double[][] invertedImage,
                          int startx, int starty, int length) {
        this.image = image;
        this.invertedImage = invertedImage;
        this.startx = startx;
        this.starty = starty;
        this.length = length;
    }

    @Override
    protected void compute() {
```

```
@Override
protected void compute() {
    if (this.length <= CUTOFF) {
        for (int offsetX = 0; offsetX < this.length; offsetX++) {
            for (int offsetY = 0; offsetY < this.length; offsetY++) {
                this.invertedImage[this.startx + offsetX][this.starty + offsetY] =
                    - this.image[this.startx + offsetX][this.starty + offsetY];
            }
        }
    } else {
        int halfSize = (this.length) / 2;
        ImageNegationFJ upperLeft = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx, this.starty, halfSize);
        ImageNegationFJ upperRight = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx + halfSize, this.starty,
            halfSize);
        ImageNegationFJ lowerLeft = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx, this.starty + halfSize,
            halfSize);
        ImageNegationFJ lowerRight = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx + halfSize,
            this.starty + halfSize, halfSize);
        upperLeft.fork();
        upperLeft.join();
        upperRight.fork();
        upperRight.join();
        lowerLeft.fork();
        lowerLeft.join();
        lowerRight.compute();
    }
}
```

```
final static int CUTOFF = 32;
```

```
double[][] image, invertedImage;
```

```
@Override  
protected void compute() {  
    if (this.length <= CUTOFF) {  
        for (int offsetX = 0; offsetX < this.length; offsetX++) {  
            for (int offsetY = 0; offsetY < this.length; offsetY++) {  
                this.invertedImage[this.startx + offsetX][this.starty + offsetY] = 1  
                    - this.image[this.startx + offsetX][this.starty + offsetY];  
            }  
        }  
    } else {  
        int halfSize = (this.length) / 2;  
        ImageNegationFJ upperLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty, halfSize);  
        ImageNegationFJ upperRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize, this.starty,  
            halfSize);  
        ImageNegationFJ lowerLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty + halfSize,  
            halfSize);  
        ImageNegationFJ lowerRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize,  
            this.starty + halfSize, halfSize);  
        upperLeft.fork();  
        upperLeft.join();  
        upperRight.fork();  
        upperRight.join();  
        lowerLeft.fork();  
        lowerLeft.join();  
        lowerRight.compute();  
    }  
}
```

- (a) Welche Annahme trifft der Code bezüglich der Abmessungen des Arrays, das das Eingabebild darstellt?

*What assumption does the code make (2) concerning the dimensions of the array representing the input image?*

**Tobias Steinbrecher** @tsteinbreche · 8 months ago · edited 8 months ago

▼ 13 ▲

---

The image should be square  $s \times s$  and we should have  $s = d2^k$ , where  $d \leq 32$ . This is necessary, because we want `length` to be divisible by 2 in the case `length > 32`. If this would not be the case, we would do floor division and leave pixels unprocessed.

+ Add Comment ⋮ More

```

@Override
protected void compute() {
    if (this.length <= CUTOFF) {
        for (int offsetX = 0; offsetX < this.length; offsetX++) {
            for (int offsetY = 0; offsetY < this.length; offsetY++) {
                this.invertedImage[this.startx + offsetX][this.starty + offsetY] =
                    - this.image[this.startx + offsetX][this.starty + offsetY];
            }
        }
    } else {
        int halfSize = (this.length) / 2;
        ImageNegationFJ upperLeft = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx, this.starty, halfSize);
        ImageNegationFJ upperRight = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx + halfSize, this.starty,
            halfSize);
        ImageNegationFJ lowerLeft = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx, this.starty + halfSize,
            halfSize);
        ImageNegationFJ lowerRight = new ImageNegationFJ(this.image,
            this.invertedImage, this.startx + halfSize,
            this.starty + halfSize, halfSize);
        upperLeft.fork();
        upperLeft.join();
        upperRight.fork();
        upperRight.join();
        lowerLeft.fork();
        lowerLeft.join();
        lowerRight.compute();
    }
}

```

- (b) Parallelisiert der Code die beabsichtigte Aufgabe korrekt oder gibt es weitere Optimierungsmöglichkeiten? Wenn ja, welche Optimierung würden Sie vorschlagen und warum?
- 2024-07-30T13:53:26.664575+00:00

*Does the code correctly parallelize the intended task or is there further optimization that could be done? If so, which optimization would you propose and why?* (4)

Not good:

```
upperLeft.fork();
upperLeft.join();
upperRight.fork();
upperRight.join();
lowerLeft.fork();
lowerLeft.join();
lowerRight.compute();
```

Tobias Steinbrecher @tsteinbreche · 8 months ago · edited 7 months ago

▼ 8 ▲

No, the parallelization is incorrect, as we have subsequent `fork()` and `join()` calls, which means that we wait for the corresponding subproblem to be finished, before calling `fork()` on the next one. To fix this, we should do the following:

```
upperLeft.fork();
upperRight.fork();
lowerLeft.fork();
lowerRight.compute();
upperLeft.join();
upperRight.join();
lowerLeft.join();
```

```
public class ImageNegationFJ extends RecursiveAction {
    final static int CUTOFF = 32;
    double[][] image, invertedImage;
    int startx, starty;
    int length;

    public ImageNegationFJ(double[][] image, double[][] invertedImage,
        int startx, int starty, int length) {
        this.image = image;
        this.invertedImage = invertedImage;
        this.startx = startx;
        this.starty = starty;
        this.length = length;
    }

    @Override
    protected void compute() {
```

- (c) Vervollständigen Sie das folgende Codegerüst, indem Sie die oben implementierte ImageNegationFJ Klasse und die ForkJoinPool Klasse verwenden, um die Variable negatedImage mit den negierten Werten zu füllen.

```
double[][] image = {{0, 1}, {1, 0}};
int imageSize = image.length;
double[][] negatedImage = new double[imageSize][imageSize];
.....
.....
```

*Complete the following code skeleton by using the above implemented ImageNegationFJ class and the ForkJoinPool class to fill the variable negatedImage with the negated values.* (4)

**Tobias Steinbrecher** @tsteinbreche · 8 months ago

▼ 10 ▲

```
double[][] image = {{0,1}, {1,0}};
int imageSize = image.length;
double[][] negatedImage = new double[imageSize][imageSize];
ForkJoinPool fjp = new ForkJoinPool();
ForkJoinTask t = new ImageNegationFJ(image, negatedImage, 0, 0 ,imageSize);
fjp.invoke(t);
```

+ Add Comment ⋮ More

- (d) Unter der Annahme, dass die Klasse `ImageNegationFJ` korrekt parallelisiert ist, wie viele Threads verwendet der `ForkJoinPool` effektiv, um das  $2 \times 2$  `negatedImage` Array aus Aufgabe 3c) zu füllen?

Assuming that the `ImageNegationFJ` (2) class is correctly parallelized, how many threads does the `ForkJoinPool` effectively use to fill the  $2 \times 2$  `negatedImage` array from task 3c)?

```
double[][] image = {{0,1}, {1,0}};
```

```
@Override  
protected void compute() {  
    if (this.length <= CUTOFF) {  
        for (int offsetX = 0; offsetX < this.length; offsetX++) {  
            for (int offsetY = 0; offsetY < this.length; offsetY++) {  
                this.invertedImage[this.startx + offsetX][this.starty + offsetY] = 1  
                    - this.image[this.startx + offsetX][this.starty + offsetY];  
            }  
        }  
    } else {  
        int halfSize = (this.length) / 2;  
        ImageNegationFJ upperLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty, halfSize);  
        ImageNegationFJ upperRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize, this.starty,  
            halfSize);  
        ImageNegationFJ lowerLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty + halfSize,  
            halfSize);  
        ImageNegationFJ lowerRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize,  
            this.starty + halfSize, halfSize);  
        upperLeft.fork();  
        upperLeft.join();  
        upperRight.fork();  
        upperRight.join();  
        lowerLeft.fork();  
        lowerLeft.join();  
        lowerRight.compute();  
    }  
}
```

```
final static int CUTOFF = 32;
```

- (d) Unter der Annahme, dass die Klasse `ImageNegationFJ` korrekt parallelisiert ist, wie viele Threads verwendet der `ForkJoinPool` effektiv, um das  $2 \times 2$  `negatedImage` Array aus Aufgabe 3c) zu füllen?

Assuming that the `ImageNegationFJ` (2) class is correctly parallelized, how many threads does the `ForkJoinPool` effectively use to fill the  $2 \times 2$  `negatedImage` array from task 3c)?

```
double[][] image = {{0,1}, {1,0}};
```

```
@Override  
protected void compute() {  
    if (this.length <= CUTOFF) {  
        for (int offsetX = 0; offsetX < this.length; offsetX++) {  
            for (int offsetY = 0; offsetY < this.length; offsetY++) {  
                this.invertedImage[this.startx + offsetX][this.starty + offsetY] = 1  
                    - this.image[this.startx + offsetX][this.starty + offsetY];  
            }  
        }  
    } else {  
        int halfSize = (this.length) / 2;  
        ImageNegationFJ upperLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty, halfSize);  
        ImageNegationFJ upperRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize, this.starty,  
            halfSize);  
        ImageNegationFJ lowerLeft = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx, this.starty + halfSize,  
            halfSize);  
        ImageNegationFJ lowerRight = new ImageNegationFJ(this.image,  
            this.invertedImage, this.startx + halfSize,  
            this.starty + halfSize, halfSize);  
        upperLeft.fork();  
        upperLeft.join();  
        upperRight.fork();  
        upperRight.join();  
        lowerLeft.fork();  
        lowerLeft.join();  
        lowerRight.compute();  
    }  
}
```

```
final static int CUTOFF = 32;
```

Tobias Steinbrecher @tsteinbreche · 8 months ago · edited 8 months ago

Because of the sequential cutoff, only **one** Thread would be used *effectively*.

- (e) Gehen Sie von einem konstanten Overhead von  $16\text{ MB} = 2^4\text{ MB}$  pro Thread aus und dass pro Split immer vier neue Threads erstellt werden. Dies bedeutet, dass die Anzahl der Threads nicht durch den ForkJoinPool festgelegt wird, sodass kein Thread wieder verwendet wird und es zu keinem Work Stealing zwischen den Threads kommt. Was ist der niedrigste Wert für CUTOFF, wenn Sie ein Bild der Größe  $4000 \times 4000$  eingeben, bevor Ihnen bei einem RAM der Größe 10 GB der Speicher ausgeht? Hinweis:  $1\text{ GB} = 2^{10}\text{ MB}$ .

*Assume a fixed overhead of  $16\text{ MB} = 2^4\text{ MB}$  per thread and that there are always four new threads created per split. This means that the number of threads is not fixed by the ForkJoinPool, so no thread is re-used and there is no work stealing among the threads. What is the lowest value for CUTOFF if you input an image of size  $4000 \times 4000$  before you run out of memory using a RAM of size 10 GB? Hint:  $1\text{ GB} = 2^{10}\text{ MB}$ .*

- (e) Gehen Sie von einem konstanten Overhead von  $16\text{ MB} = 2^4\text{ MB}$  pro Thread aus und dass pro Split immer vier neue Threads erstellt werden. Dies bedeutet, dass die Anzahl der Threads nicht durch den ForkJoinPool festgelegt wird, sodass kein Thread wieder verwendet wird und es zu keinem Work Stealing zwischen den Threads kommt. Was ist der niedrigste Wert für CUTOFF, wenn Sie ein Bild der Größe  $4000 \times 4000$  eingeben, bevor Ihnen bei einem RAM der Größe 10 GB der Speicher ausgeht? Hinweis:  $1\text{ GB} = 2^{10}\text{ MB}$ .

Assume a fixed overhead of  $16\text{ MB} = 2^4\text{ MB}$  per thread and that there are always four new threads created per split. This means that the number of threads is not fixed by the ForkJoinPool, so no thread is re-used and there is no work stealing among the threads. What is the lowest value for CUTOFF if you input an image of size  $4000 \times 4000$  before you run out of memory using a RAM of size 10 GB? Hint:  $1\text{ GB} = 2^{10}\text{ MB}$ .

Tobias Steinbrecher @tsteinbreche · 8 months ago · edited 8 months ago

▼ 14 ▲

---

Number of threads, which we can use:

$$N = \frac{10 \cdot 2^{10}}{2^4} = 10 \cdot 2^6 = 10 \cdot 4^3$$

In each recursive call, we will use 4 new threads (under given assumptions). Thereby, we have the constraint ( $i$  := number of divisions)

$$4^i \leq 10 \cdot 4^3 \iff i \leq \log_4(10) + 3 \iff i \leq 4$$

and the smallest possible value is  $\text{CUTOFF} = 4000/2^4 = 250$  to avoid a fifth division.

# Plan für heute

- Organisation
- Nachbesprechung Exercise 4
- Theory Recap
- Intro Exercise 5
- Exam Questions
- **Kahoot**

The logo consists of the word "QUIZ" in red, bold, sans-serif letters. Each letter is contained within a separate white rectangular box with a black double-line border. The boxes are slightly overlapping and angled towards the top right, creating a dynamic feel.

# Feedback

- Falls ihr Feedback möchten sagt mir bitte Bescheid!
- Schreibt mir eine Mail oder auf Discord

# Danke

- Bis nächste Woche!