# **Concurrent Programming Models**

CS 475, Spring 2019 **Concurrent & Distributed Systems** 



With material from Herlihy & Shavit, Art of Multiprocessor Programming





# **Review: Course Grained Locking**





### **Review: Fine-Grained Synchronization**

- Instead of using a single lock ...
- Split object into - Independently-synchronized components
- Methods conflict when they access - The same component ...
  - At the same time



# **Review: Fine Grained Locking List**







### **Review: Optimistic Synchronization**

- Search without locking ...
- If you find it, lock and check ...
  - OK: we are done
  - Oops: start over
- Evaluation
  - Usually cheaper than locking
  - Mistakes are expensive



### **Review: Optimistic List**





- Postpone hard work
- Removing components is tricky
  - Logical removal
  - Mark component to be deleted - Physical removal
    - Do what needs to be done

### **Review: Lazy Synchronization**



### Review: Lazy List

### C b a contains(b) Returns false, because b is marked. No need to validate or lock or re-traverse

Intuition: Now can judge if b is in list ONLY by looking at b, don't also need to look at a





## Today

- How do we increase performance with parallelism?
- How do we split up our program into concurrent sections effectively? Different models for parallel computation  $\bullet$
- Reading: H&S 16.1, 16.2



# Designing for Performance

- What factors can impact performance?
  - Limits imposed by physics
  - Limits imposed by technology
  - Limits imposed by economics
- These limits can force us to make tradeoffs
  - Smaller chips are faster, but harder to dissipate heat
  - Need to serve X clients, can only spend Y on CPUs



# Performance Metrics

- Capacity
  - Consistent measure of a service's size or amount of resources
- Utilization  $\bullet$ 
  - Percentage of that resource used for a workload
- Latency
  - How long it takes an input to propagate through a system and generate an output
- Throughput
  - Work done per time



11

Adjusted by buying

more

resou

hard about

oroble

3

Adjusted

### Latency

- and receiving response
- What contributes to latency?
  - Latency sending the message  $\bullet$
  - Latency processing the message
  - Latency sending the response
- Adding pipelined components -> latency is cumulative



In client/server model, latency is simply: time between client sending request



12

# Throughput

- Measure of the rate of useful work done for a given workload
- Example:  $\bullet$ 
  - Throughput is camera frames processed/second
- $\bullet$ • When adding multiple pipelined components -> throughput is the minimum value



Total throughput: 10fps



# Designing for Performance

- or latency)
- Measure each component to identify bottleneck
- Measure improvement
- Repeat

Measure system to find which aspect of performance is lacking (throughput

Identify if fixing that bottleneck will realistically improve system performance





# Improving Throughput



- Introduce concurrency into our pipeline
- Each stage runs in its own thread (or many threads, perhaps)
- If a stage completes its task, it can start processing the next request right away
  - E.g. our system will process multiple requests at the same time



# Improving Throughput



# Reducing Latency

- Often more challenging than increasing throughput
  - Examples:
    - Physical Speed of light (network transmissions over long distances) Algorithmic - Looking up an item in a hash table is limited by hash function • Economic - Adding more RAM gets expensive





- Buy low/sell high
- Most of skill is in knowing what a stock will do before your competitors





- Algorithmic trading -> computer programs look at various factors, place trades automatically
- Example:
  - President Trump tweets positively about a company -> price goes up
  - stock
  - Get in and out before it hits CNN!  $\bullet$
  - planet-money-s-stock-trading-twitter-bot

• Write a script to check twitter for company mentions, immediately buy/sell

https://www.npr.org/sections/money/2017/04/07/522897876/meet-botus-



- What if you set up this bot in Chicago, and I set one up in NYC?  $\bullet$ 
  - I would beet you to it, every time.

• This only works if you can make your trades **before** other people find out



- What is the speed of light?
  - ~300,000 km/sec
- How fast does your CPU execute an instruction?
  - 0.33 nanoseconds (say, 3Ghz CPU)
- How far does light travel in 1 CPU cycle?
  - 10 cm
- Chicago to NYC and back?
  - ~700 miles -> 7.4msec -> 22 million instructions
- $\bullet$ your stock order to NYC and get a response!

• How many instructions does your CPU execute in the time it takes light to travel from

Being in NYC would let me execute 22 million instructions in the time it took you to send





# **Reducing Latency with \$\$\$\$**

- $\bullet$ 
  - price of some stocks

People actually care a LOT about the latency between NYC and Chicago, because commodities are traded in Chicago and stocks are traded in NYC

• Changes to commodities prices (e.g. **ethanol**) can dramatically impact



# **Reducing Latency with \$\$\$\$**

- It's not quite as simple as 700 miles -> 7.4msec
- There are streams, mountains, etc... more like 1,000 miles
- Light is refracted in a fiber optic cable is ~31% slower
- What do we do if money is no object?





### **Reducing Latency with Billions of Dollars**



**ORIGINAL CABLE** Technology Buried fiber-optic cable

Completion Mid-1980s

Path length ~ 1.000 miles

Round-trip time for data 14.5

milliseconds and up

Approach Multiple routes followed the easiest rightsof-way-along rail lines. But that means time-sucking jogs and detours.

### SPREAD NETWORKS

Technology Buried fiber-optic cable

Completion August 2010

Path length 825 miles

Round-trip time for data

13.1 milliseconds

Approach Spread bought its own rights-of-way, avoiding a Philadelphia-ward dip in favor of a shorter path northwest through central Pennsylvania.

### 

### MCKAY BROTHERS

Technology Microwave beams through air

Completion July 4, 2012

Path length 744 miles

Round-trip time for data

9 milliseconds

### Approach

Microwaves generally move faster than photons in optical fiber, and McKay's network uses just 20 towers on a nearly perfect great circle.

### 

### TRADEWORX

Technology Microwave beams through air

Completion Winter 2012

Path length ~ 731 miles

Round-trip time for data

8.5 milliseconds (est.)

Approach Tradeworx is highly secretive, but the company is open about the price of a subscription: \$250,000 a year.

https://www.zerohedge.com/news/chicago-new-york-and-back-85-milliseconds



24

# **Reducing Latency without lots of \$\$\$**

- Approach: use **concurrency**
- Limited by serial section



GMU CS 475 Spring 2019



25

- These examples are at a very high level (components in a large server system)
- For this lecture, we'll focus on smaller, more concrete examples
- First: Matrix Multiplication

## **Exploiting Concurrency**

# $(C) = (A) \cdot (B)$



### $C_{ij} = \sum_{k=0}^{N-1} a_{ki} * b_{jk}$



```
class Worker extends Thread {
int row, col;
Worker(int row, int col) {
 this.row = row; this.col = col;
public void run() {
 double dotProduct = 0.0;
 for (int i = 0; i < n; i++)
  dotProduct += a[row][i] * b[i][col];
 c[row][col] = dotProduct;
}}}
```



class Worker extends Thread { int row, col, Worker(int row, int col) { this.row = row; this.col = c public void run() { double dotProduct = 0.0; for (int i = 0; i < n; i++) dotProduct += a[row][i] \* b[i][col]; c[row][col] = dotProduct; }}}





class Worker extends Thread { int row, col; Worker(int row, int col) { <del>rew = rew; th</del>is.col public void run() { double dotProduct = 0.0; for (int i = 0; i Which matrix entry to dotProduct += a[row][i] \* b[i][col]; c[row][col] = dotProduct, Compute }}}



class Worker extends Thread { int row, col; Worker(int row, Actual computation this.row = row; this.col = col;public void run() { double dotProduct = 0 TOP (IFIL I = 0, I < II, I++dotProduct += a[row][i] \* b[i][col]; c[row][col] = dotProduct; 





void multiply() { Worker[]] worker = new Worker[n][n]; for (int row ...) for (int col ...) worker[row][col] = new Worker(row,col); for (int row ...) for (int col ...) worker[row][col].start(); for (int row ...) for (int col ...) worker[row][col].join();



void multiply() { MarkarIII worker - new MarkarInIIn for (int row ...) for (int col ...) worker[row][col] = new Worker(row,col); for (int row ...) for (int col ...) worker[row][col].start(); for (int row ...) for (int col ...) worker[row][col].join();





void multiply() { Worker[]] worker = new Worker[n][n]; for (int row ...) for (int col ...) worker[row][col] for (int row ...) for (int col ...) worker[row][col].start(); for (int row ...) for (int col ...) worker[row][col].join();












- Threads Require resources
  - Memory for stacks
  - Setup, teardown
- Scheduler overhead
- Worse for short-lived threads

### Thread Overhead



## Thread Pools

- More sensible to keep a pool of long-lived threads
- Threads assigned short-lived tasks
  - Runs the task
  - Rejoins pool —
  - Waits for next assignment



- Insulate programmer from platform
  - Big machine, big pool
  - And vice-versa
- Portable code
  - Runs well on any platform \_\_\_\_\_
  - No need to mix algorithm/platform concerns \_\_\_\_

### Thread Pool = Abstraction



## **ExecutorService Interface**

- In java.util.concurrent
  - Task = **Runnable** object
    - If no result value expected
    - Calls **run()** method.
  - Task = Callable<T> object
    - If result value of type **T** expected
    - Calls **T call()** method.
    - Interesting question: how do you get the return value from call?



### Future<T>

Callable<T> task = ...; . . . Future<T> future = executor.submit(task); . . . T value = future.get();



### **Future<T>**

### Callable<T> task = ...;

Future<T> future = executor.submit(task);

T value = future.get();

. . .





### Future<T>

Callable<T> task = ...;

Future<T> future = executor.submit(task);

T value = future.get();

### The Future's get() method blocks until the value is available



Runnable task = ...; . . . Future<?> future = executor.submit(task); . . . future.get();

Future<?>



Runnable task = ...;

Future<?> future = executor.submit(task);

future.get();

**1 1 1** 

Future<?>





Runnable task = ...;

Future<?> future = executor.submit(task);

future.get();

### The Future's get() method blocks until the computation is complete

Future<?>



- Executor Service submissions
  - Like Maryland traffic signs
  - Are purely advisory in nature
- The executor
  - Like the Maryland driver
  - Is free to ignore any such advice —
  - And could execute tasks sequentially ...

### Note





### Matrix Addition

 $\begin{pmatrix} C_{00} & C_{00} \\ C_{10} & C_{10} \end{pmatrix} = \begin{pmatrix} A_{00} + B_{00} & B_{01} + A_{01} \\ A_{10} + B_{10} & A_{11} + B_{11} \end{pmatrix}$ 





### Matrix Addition

### 4 parallel additions



class AddTask implements Runnable { Matrix a, b; // multiply this! public void run() { if (a.dim == 1) { c[0][0] = a[0][0] + b[0][0]; // base case} else { (partition a, b into half-size matrices a<sub>ii</sub> and b<sub>ii</sub>) Future <?>  $f_{00} = exec.submit(add(a_{00}, b_{00}));$ 

Future <?>  $f_{11} = exec.submit(add(a_{11}, b_{11}));$ f<sub>00</sub>.get(); ...; f<sub>11</sub>.get();

}}

- This is not real Java



class AddTask implements Runnable { Matrix a, b; // multiply this! public void run() if (a.dim == 1) { c[0][0] = a[0][0] + b[0][0]; // base caseelse (partition a, b into half-size matrices a<sub>ii</sub> and b<sub>ii</sub>) Future <?>  $f_{00} = exec.submit(add(a_{00}, b_{00}));$ . . . Future <?>  $f_{11} = exec.submit(add(a_{11}, b_{11}));$ f<sub>00</sub>.get(); ...; f<sub>11</sub>.get(); Base case: add directly าา





class AddTask implements Runnable { Matrix a, b; // multiply this! public void run() { if (a.dim == 1) { c[0][0] = a[0][0] + b[0][0]; // base caseelse (partition a, b into half-size matrices a<sub>ii</sub> and b<sub>ii</sub>)  $> f_{00} = exec.submit(add(a_{00}, b_{00}))$ . . . Future<?> f<sub>11</sub> = exec.submit(add f<sub>00</sub>.get(); ...; f<sub>11</sub>.get();

**Constant-time operation** 



class AddTask implements Runnable { Matrix a, b; // multiply this! public void run() { if (a.dim == 1) { c[0][0] = a[0][0] + b[0][0]; // base case} else { (partition a, b into half-size matrices a<sub>ii</sub> and b<sub>ii</sub>) Future <?>  $f_{00} = exec.submit(add(a_{00}, b_{00}));$ Future <?>  $f_{11}$  = exec.submit(add( $a_{11}$ , $b_{11}$ )); f<sub>00</sub>.get(); ...; f<sub>11</sub>.get(); าา

Submit 4 tasks



class AddTask implements Runnable { Matrix a, b; // multiply this! public void run() { if (a.dim == 1) { c[0][0] = a[0][0] + b[0][0]; // base case} else { (partition a, b into half-size matrices a<sub>ii</sub> and b<sub>ii</sub>) Future <?>  $f_{00} = exec.submit(add(a_{00}, b_{00}));$ 



### Let them finish



## Dependencies

- Matrix example is not typical
- Tasks are independent
  - Don't need results of one task ... \_\_\_\_\_
  - To complete another
- Often tasks are not independent



### Fibonacci

# F(n) $\begin{cases} 1 \text{ if } n = 0 \text{ or } 1 \\ F(n-1) + F(n-2) \text{ otherwise} \end{cases}$

Note: potential parallelism, but subject to dependencies 



### Disclaimer

- This Fibonacci implementation is
  - Egregiously inefficient
    - So don't deploy it!
  - But illustrates our point
    - How to deal with dependencies



## Multithreaded Fibonacci

class FibTask implements Callable<Integer> { static ExecutorService exec = Executors.newCachedThreadPool(); int arg; public FibTask(int n) { arg = n;public Integer call() { if (arg > 2) { Future<Integer> left = exec.submit(new FibTask(arg-1)); Future<Integer> right = exec.submit(new FibTask(arg-2)); return left.get() + right.get(); } else { return 1; }}}



## Multithreaded Fibonacci

class FibTask implements Callable<Integer> { static ExecutorService exec = Executors.newCachedThreadPool(); int arg; public FibTask(int n) { Parallel calls arg = n;public Integer call() { if (arg > 2)Future<Integer> left = exec.submit(new FibTask(arg-1)); Future<Integer> right = exec.submit(new FibTask(arg-2)); return left.get() + right.get(); else { return 1; }}}



## Multithreaded Fibonacci

class FibTask implements Callable<Integer> { static ExecutorService exec = Executors.newCachedThreadPool(); int arg; public FibTask(int n) { Pick up & combine results public Integer call() { if (arg > 2) { Future<Integer> left = exec.submit(new FibTask(arg-1)); Future<Integer> right = exee.submit(new FibTask(arg-2)); return left.get() + right.get(); else return 1; }}}



## Dynamic Behavior

- Multithreaded program is
  - A directed acyclic graph (DAG)
  - That unfolds dynamically
- Each node is
  - A single unit of work





### Note inefficiency in this implementation: fib(2)'s result should be computed only once





### **Arrows Reflect Dependencies**



### Note inefficiency in this implementation: fib(2)'s result should





## How Parallel is That?

- Define work: - Total time on one processor
- Define critical-path length:
  - Longest dependency path
  - Can't beat that!













## Fib Critical Path





## Fib Critical Path





## Notation Watch

- $T_P = time on P processors$
- $T_1 = \text{work}$  (time on 1 processor)
- $T_{\infty}$  = critical path length (time on  $\infty$  processors)



## Simple Bounds

### • $T_P \ge T_1/P$

- In one step, can't do more than P work
- $T_P \ge T_\infty$ 
  - Can't beat infinite resources



## **More Notation Watch**

- Speedup on P processors - Ratio  $T_1/T_P$ 
  - How much faster with P processors
- Linear speedup  $-T_1/T_P = \Theta(P)$
- Max speedup (average parallelism)  $- T_1/T_{\infty}$





### Matrix Addition

 $\begin{pmatrix} C_{00} & C_{00} \\ C_{10} & C_{10} \end{pmatrix} = \begin{pmatrix} A_{00} + B_{00} & B_{01} + A_{01} \\ A_{10} + B_{10} & A_{11} + B_{11} \end{pmatrix}$ 




### Matrix Addition

### 4 parallel additions



### Addition

- Let  $A_P(n)$  be running time
  - For n x n matrix
  - on P processors
- For example
  - $A_1(n)$  is work
  - $A_{\infty}(n)$  is critical path length





Work is  $\bullet$ 

# Partition, synch, etc

# $A_1(n) = 4 A_1(n/2) + \Theta(1)$ 4 spawned additions

### Addition





Work is 

### $A_1(n) = 4 A_1(n/2) + \Theta(1)$ $= \Theta(n^2)$ Same as double-loop summation

### Addition





Critical Path length is

# $A_{\infty}(n) = A_{\infty}(n/2) - \Theta(1)$

### spawned additions in parallel

### Addition

### Partition, synch, etc





Critical Path length is

### $A_{\infty}(n) = A_{\infty}(n/2) + \Theta(1)$ $= \Theta(\log n)$

### Addition



### Matrix Multiplication Redux





### Matrix Multiplication Redux





 $\begin{pmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{21} & \mathbf{C}_{22} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} \end{pmatrix}$ 





### 8 multiplications



 $\begin{pmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{21} & \mathbf{C}_{22} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} \end{pmatrix}$ 

### Second Phase ...





Work is  $\bullet$ 

### $M_1(n) = 8 M_1(n/2) + A_1(n)$ 8 parallel multiplications

### Final addition



Work is 

### $M_1(n) = 8 M_1(n/2) + \Theta(n^2)$ $= \Theta(n^3)$

### Same as serial triple-nested loop



Critical path length is  $\bullet$ 

### Final addition $M_{\infty}(n) = M_{\infty}(n/2) + A_{\infty}(n)$ Half-size parallel multiplications



• Critical path length is

### $M_{\infty}(n) = M_{\infty}(n/2) + A_{\infty}(n)$ $= M_{\infty}(n/2) + \Theta(\log n)$ $= \Theta(\log^2 n)$



### Parallelism

- $M_1(n)/M_{\infty}(n) = \Theta(n^3/\log^2 n)$
- To multiply two 1000 x 1000 matrices  $-1000^{3}/10^{2}=10^{7}$
- Much more than number of processors on any real machine



### Shared-Memory Multiprocessors

- Parallel applications Do not have direct access to HW processors
- Mix of other jobs
  - All run together
  - Come & go dynamically
- point
- processors we actually get

### • Hence, we have **no control** over how many processors we get at any given

Instead, shoot for the best parallelism that we can get given however many



### Concurrent **Programming Models**



# Asynchronous Programming

- AKA event-driven programming
- A paradigm that lends itself well to scaling, especially in a multi-stage systems (like the example with Facebook)
- Allows us to think about what is done, abstract away how it is done
- We will discuss two asynchronous models: streams, and Promises, neither of which make you think about threads (or locks?)





### Streams

- Java 8 introduced the concept of Streams
- A stream is a sequence of objects
- non-interfering and stateless
  - Non-interfering: Does not modify the actual stream

### • Example:

IntStream.range(1, 1000000) //Generate a stream of all ints 1 - 1m .filter(x -> isPrime(x)) //Retain only values that pass some expensive isPrime function .forEach(System.out::println); //For each value returned by filter, print it

• Streams have functions that you can perform on them, which are (mostly)

• Stateless: Each time the function is called on the same data, get same result





### Sidebar: Lambdas

 I don't know if you have seen this before IntStream.range(1, 100000)

- .filter(x -> isPrime(x))
- .forEach(System.out::println);
- This line is called a lambda expression
- $\bullet$ was released in 2014
- Effectively, think of this as shorthand for:  $\bullet$ IntStream.range(1, 1000000) .filter(new IntPredicate() { @Override public boolean test(int x) { return isPrime(x); } })

.forEach(System.out::println);

 $\bullet$ right?)

We should have shown it to you before, because it's a core part of Java syntax since Java 8

In fact, javac generates exactly the long-hand code for that shorthand (but that println is cool,



### Streams

IntStream.range(1, 1000000) //Generate a stream of all ints 1 - 1m  $filter(x \rightarrow isPrime(x))$  //Retain only values that pass some expensive isPrime function .forEach(System.out::println); //For each value returned by filter, print it

- Why use the stream interface instead of for(int i = 1; i < 1000000; i++)</pre> if(isPrime(x)) System.out.println(x);
- Who wants to write the parallel version of this? IntStream.range(1, 1000000) //Generate a stream of all ints 1 - 1m  $filter(x \rightarrow isPrime(x))$  //Retain only values that pass some expensive isPrime function .parallel() //Do the filtering in parallel .forEach(System.out::println); //For each value returned by filter, print it
- The magic works as long as isPrime is stateless!



### Streams - what can't be parallelized

### Interference

```
List<String> list = new ArrayList<>(Arrays_asList("Luke", "Leia", "Han"));
list.stream()
.peek(name -> {
    if (name.equals("Han")) {
        list.add("Chewie"); // Adds to list that we are peeking into
    }
})
.forEach(i -> {});
  Stateful
boolean tooBusy = false;
public void isPrime(int x)
   if(tooBusy)
       return false;//don't bother running if another thread set tooBusy
   else
       //do a sieve of erasthenes
}
```

### • Side effects

```
List<Integer> list = new ArrayList<>(
Arrays.asList(1,3,5,7,9,11,13,15,17,19));
List<Integer> result = new ArrayList<>();
list.parallelStream()
.filter(x -> isPrime(x))
.forEach(x -> result.add(x)); //Changing external state, which may not (is not) thread safe
```





### Streams under the hood

- Just adding more parallel() doesn't always make it faster! (see: law of leaky abstractions)
- There is some overhead to how a parallel operation occurs Internally, Java keeps a pool of worker threads (rather than make new
- threads for each parallel task)
- Streams use a special kind of pool, called a ForkJoinPool



# Fork/Join Programming

- to combine the results
- Similar to map/reduce, but not distributed
- For streams:
  - **Fork** a task into subtasks for many threads to work on •
  - **Join** the results together •

Special kind of task - fork() defines how to create subtasks, join() defines how





### Fork/Join Programming

Obligatory array sum example  $\bullet$ 

}

class Sum extends RecursiveTask<Long> { static final int SEQUENTIAL\_THRESHOLD = 5000;

```
int low;
int high;
int[] array;
Sum(int[] arr, int lo, int hi) {
    array = arr;
    low = lo;
    high = hi;
}
protected Long compute() {
    if(high - low <= SEQUENTIAL_THRESHOLD) {</pre>
         long sum = 0;
        for(int i=low; i < high; ++i)</pre>
             sum += array[i];
        return sum;
    } else {
         int mid = low + (high - low) / 2;
         Sum left = new Sum(array, low, mid);
         Sum right = new Sum(array, mid, high);
         left.fork();
         long rightAns = right.compute();
         long leftAns = left.join();
         return leftAns + rightAns;
    }
}
static long sumArray(int[] array) {
     return ForkJoinPool.commonPool().invoke(new Sum(array,0,array.length));
}
```



### **Promises & CompleteableFutures**

- What if we want to run some task, and do stuff while we are waiting for it to be done?  $\bullet$  You COULD do it with a complicated combination of synchronized, wait, and notify You can use the **Promise** abstraction instead

- Called a **CompletableFuture** in Java 8 CompletableFuture<String> future = CompletableFuture.supplyAsync(() -> { try { TimeUnit.SECONDS.sleep(1); } catch (InterruptedException e) { throw new IllegalStateException(e); return "Result of the asynchronous computation"; });

```
// Block and get the result of the Future
String result = future.get();
System.out.println(result);
```

• Just like Future's from before, but supports *chaining* 







```
try {
        TimeUnit.SECONDS.sleep(1);
    } catch (InterruptedException e) {
        throw new IllegalStateException(e);
    }
    return "Jon";
});
// Chain on some more code to run when the future is done
    return "Hello, " + returnValue;
});
System.out.println(greetingFuture.get()); // Hello Jon
```

CompletableFuture<String> whatsYourNameFuture = CompletableFuture.supplyAsync(() -> {

- CompletableFuture<String> greetingFuture = whatsYourNameFuture.thenApply(returnValue -> {



```
try {
        TimeUnit.SECONDS.sleep(1);
    } catch (InterruptedException e) {
        throw new IllegalStateException(e);
    }
    return "Jon";
});
// Chain on some more code to run when the future is done
CompletableFuture<String> greetingFuture =
    return "Hello, " + returnValue;
});
System.out.println(greetingFuture.get()); // Hello Jon
```







### Task will return string "Jon" eventually

TimeUnit.SECONDS.sleep(1); } catch (InterruptedException e) { throw new IllegalStateException(e); return "Jon";

// Chain on some more code to run when the future is done return "Hello, " + returnValue; **}**); System.out.println(greetingFuture.get()); // Hello Jon

CompletableFuture.supplyAsync(() -> {

- CompletableFuture<String> greetingFuture = whatsYourNameFuture.thenApply(returnValue -> {



### Task will return string "Jon" eventually

TimeUnit.SECONDS.sleep(1); } catch (InterruptedException e) { throw new IllegalStateException(e); return "Jon"; // Chain on some more code to run when the future is done CompletableFuture<String> greetingFuture = whatsYourNameFuture.thenApply(returnValue -> { return "Hello, " + returnValue; }); System.out.println(greetingFuture.get()); // Hello Jon

### >> for the second second



```
CompletableFuture<String> whatsYourName
    try {
        TimeUnit.SECONDS.sleep(1);
    } catch (InterruptedException e) {
        throw new IllegalStateException(e);
    }
    return "Jon";
});
// Chain on some more code to run when the future is done
CompletableFuture<String> greetingFuture = whatsYourNameFuture.thenApply(returnValue -> {
    return "Hello, " + returnValue;
});
System.out.println(greetingFuture.get()); // Hello Jon
```

### Create ANOTHER future that is chained to the first





CompletableFuture<String> whatsYourNameFuture = CompletableFuture.supplyAsync(() -> { try { TimeUnit.SECONDS.sleep(1); } catch (InterruptedException e) { throw new IllegalStateException(e); Block the main thread for both futures to finish **}**); // Chain on some more code to run when the future is done CompletableFuture<String> greetingFuture = whatsYourNameFuture.thenApply(returnValue -> { return "Hello, " + returnValue; }); System.out.print n(greetingFuture.get()); // Hello Jon







# **CompleteableFuture Use-Cases**

- Asynchronous I/O
  - Read data from a web service
  - Then process it
  - Then save it to a file

• Any case where you need to have multiple things happen in the background, but care about the result, and care about them happening in some order



# CompletableFutures

- an exception occurs in any of those threads
- API: <u>https://docs.oracle.com/javase/8/docs/api/java/util/concurrent/</u> <u>CompletableFuture.html</u>

Catch errors by providing a callback function for exceptionally (called when



107

### This work is licensed under a Creative Commons Attribution-ShareAlike license

- This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International
- You are free to:
  - Share copy and redistribute the material in any medium or format
  - Adapt remix, transform, and build upon the material
  - for any purpose, even commercially.
- Under the following terms:
  - suggests the licensor endorses you or your use.
  - contributions under the same license as the original.
  - legally restrict others from doing anything the license permits.

License. To view a copy of this license, visit <u>http://creativecommons.org/licenses/by-sa/4.0/</u>

• Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that

• ShareAlike — If you remix, transform, or build upon the material, you must distribute your

No additional restrictions — You may not apply legal terms or technological measures that

