Practical Concurrent and Parallel Programming 5

Peter Sestoft
IT University of Copenhagen

Friday 2014-09-26*

Plan for today

- Tasks and the Java executor framework
  - Executors, Runnables, Callables, Futures
- The states of a task
- Task creation overhead
- Using tasks to count prime numbers
- Java versus the .NET Task Parallel Library
- Producer-consumer pipelines
- Bounded queues, thread wait() and notify()
- The states of a thread

Exercises ...

- The 1200-1400 time slot is in **3A18** again
- Hand-ins: Submit a zip-file containing
  - explanation as a text file answers.txt
  - a subdirectory src/ with source code
  - graphs as image files *.jpg or *.png
  - NO Netbeans projects, Eclipse workspaces or other junk

Prefer executors and tasks to threads

- We have used *threads* to parallelize work
  - But creating many threads takes time and memory
- Better divide work into small *tasks*
  - Then submit the tasks to an executor
  - This uses a pool of (few) threads to run the tasks
- Goetz chapters 6-8 and Bloch item 68

should generally refrain from working directly with threads. The key abstraction is no longer `Thread`, which served as both the unit of work and the mechanism for executing it. Now the unit of work and mechanism are separate. The key abstraction is the unit of work, which is called a *task*. There are two kinds of tasks: `Runnable` and its close cousin, `Callable` (which is like `Runnable`, except that it returns a value). The general mechanism for executing tasks is the *executor ser-*

Bloch item 68

Executors and tasks

- A task is just a Runnable or Callable<T>
- Submitting it to an executor gives a Future

```
Future<?> fut
  = executor.submit(new Runnable() { public void run() {
      System.out.println("Task ran!");
}});
```

TestTaskSubmit.java

- The executor has a pool of threads and uses one of them to run the task
- Use the Future to wait for task completion

```
try { fut.get(); }
catch (InterruptedException exn) { System.out.println(exn); }
catch (ExecutionException exn) { throw new RuntimeException(exn);
```

A task that produces a result

- Make the task from a Callable<T>

Future's result type

... same a Callable's

```
Future<String> fut
  = executor.submit(new Callable<String>() {
      public String call() throws IOException {
        return getPage("http://www.wikipedia.org", 10);
}});
```

TestTaskSubmit.java

- Use the Future to get the task's result:

```
try {
  String webpage = fut.get();
  System.out.println(webpage);
} catch (InterruptedException exn) { System.out.println(exn); }
  catch (ExecutionException exn) { throw new RuntimeExcep...; }
```

Task rules

- Different tasks may run on different threads
  - Objects accessed from tasks must be thread-safe
- A thread running a task can be interrupted
  - So a task can be interrupted
  - So **`fut.get()`** can throw InterruptedException

- Creating a task is fast, takes little memory
- Creating a thread is slow, takes much mem.

The states of a task

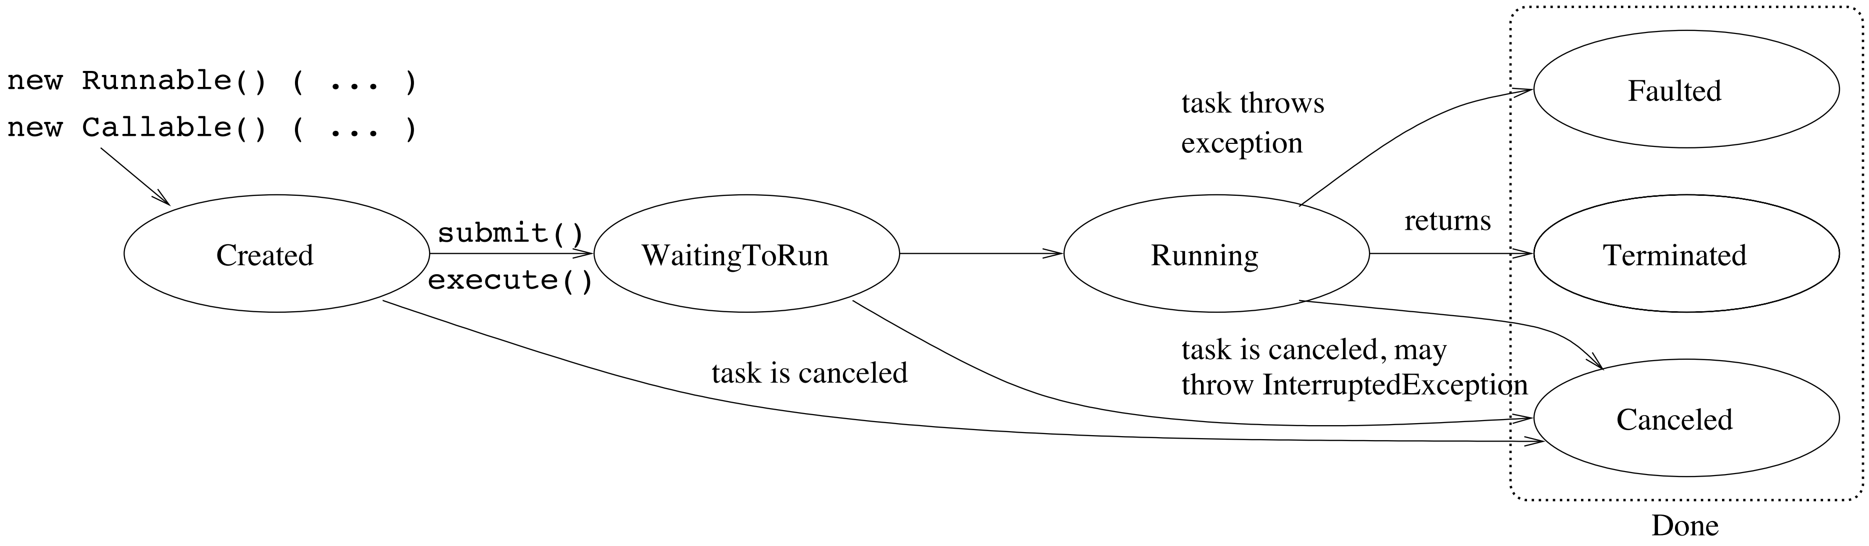

- After **`submit`** or **`execute`**
  - a task may be running immediately or much later
  - depending on the executor and available threads

Thread creation vs task creation

- Task creation is faster than thread creation

	Thread	Task
Work	6581 ns	6612 ns
Create	1030 ns	77 ns
Create+start/(submit+cancel)	48929 ns	835 ns
Create+(start/submit)+complete	72759 ns	21226 ns

Intel i7 2.4 GHz JVM 1.8

- A task also uses much less memory

Various Java executors

- In class java.util.concurrent.Executors:
- newFixedThreadPool(n)
  - Fixed number n of threads; automatic restart
- newCachedThreadPool()
  - Dynamically adapted number of threads, no bound
- newSingleThreadExecutor()
  - A single thread; so tasks need not be thread-safe
- newScheduledThreadPool()
  - Delayed and periodic tasks; eg clean-up, reporting
- newWorkStealingPool() — New in Java 8. Use it
  - Adapts thread pool to number of processors, uses multiple queues; therefore better scalability

Plan for today

- Tasks and the Java executor framework
  - Executors, Runnables, Callables, Futures
- The states of a task
- Task creation overhead
- **Using tasks to count prime numbers**
- Java versus the .NET Task Parallel Library
- Producer-consumer pipelines
- Bounded queues, thread wait and notify
- The states of a thread

Week 2 flashback: counting primes in multiple threads

```
final LongCounter lc = new LongCounter();
Thread[] threads = new Thread[threadCount];
for (int t=0; t<threadCount; t++) {
  final int from = perThread * t,
    to = (t+1==threadCount) ? range : perThread * (t+1);
  threads[t] = new Thread(new Runnable() { public void run()
    for (int i=from; i<to; i++)
      if (isPrime(i))
        lc.increment();
    }});
}
for (int t=0; t<threadCount; t++)
  threads[t].start();
```

Last thread has **to==range**

Thread t processes segment [from,to)

TestCountPrimes.java

- Creates one thread for each segment

Counting primes in multiple tasks

T1

```
final LongCounter lc = new LongCounter();
List<Future<?>> futures = new ArrayList<Future<?>>();
for (int t=0; t<taskCount; t++) {
  final int from = perTask * t,
    to = (t+1 == taskCount) ? range : perTask * (t+1);
  futures.add(executor.submit(new Runnable() { public void r
      for (int i=from; i<to; i++)
        if (isPrime(i))
          lc.increment();
  }}));
}
try {
 for (Future<?> fut : futures)
   fut.get();
} catch (...) { ... }
```

Create task, submit to executor, save a future

Add to shared

Wait for all tasks to complete

TestCountPrimesTasks.java

- Creates a task for each segment
- The tasks execute on a thread pool

Tasks that return task-local counts

T2

```
List<Callable<Long>> tasks = new ArrayList<Callable<Long>>();
for (int t=0; t<taskCount; t++) {
  final int from = perTask * t,
    to = (t+1 == taskCount) ? range : perTask * (t+1);
  tasks.add(new Callable<Long>() { public Long call() {
    long count = 0;
    for (int i=from; i<to; i++)
      if (isPrime(i))
        count++;
    return count;
  }});
}
long result = 0;
try {
  List<Future<Long>> futures = executor.invokeAll(tasks);
  for (Future<Long> fut : futures)
    result += fut.get();
} catch (...) { ... }
```

Create task

Add to local

Submit tasks, wait for all to complete, get futures

Add local task results

TestCountPrimesTasks.java

Callable<Void> is like Runnable

T3

```
final LongCounter lc = new LongCounter();
List<Callable<Void>> tasks = new ArrayList<Callable<Void>>()
for (int t=0; t<taskCount; t++) {
  final int from = perTask * t,
    to = (t+1 == taskCount) ? range : perTask * (t+1);
  tasks.add(new Callable<Void>() { public Void call() {
      for (int i=from; i<to; i++)
        if (isPrime(i))
          lc.increment();
      return null;
  }});
}
try {
  executor.invokeAll(tasks);
} catch (...) { ... }
```

Create task

Add to shared

Submit tasks, wait for all to complete

TestCountPrimesTasks.java

Type parameters <Void> and <?>

- The type java.lang.Void contains only **`null`**
- Callable<Void> requires **`Void call() {...}`**
  - Similar to Runnable's **`void run() { ... }`**
  - With Future<Void> the **`get()`** returns null

Not same

- Future<?> has an unknown type of value
  - With Future<?> the **`get()`** returns null also
- Java's type system is a somewhat muddled
  - Will not allow this assignment:

Type Future<?>

```
Future<Void> future;
future = executor.submit(new Runnable() { ... });
```

Overhead of creating many threads

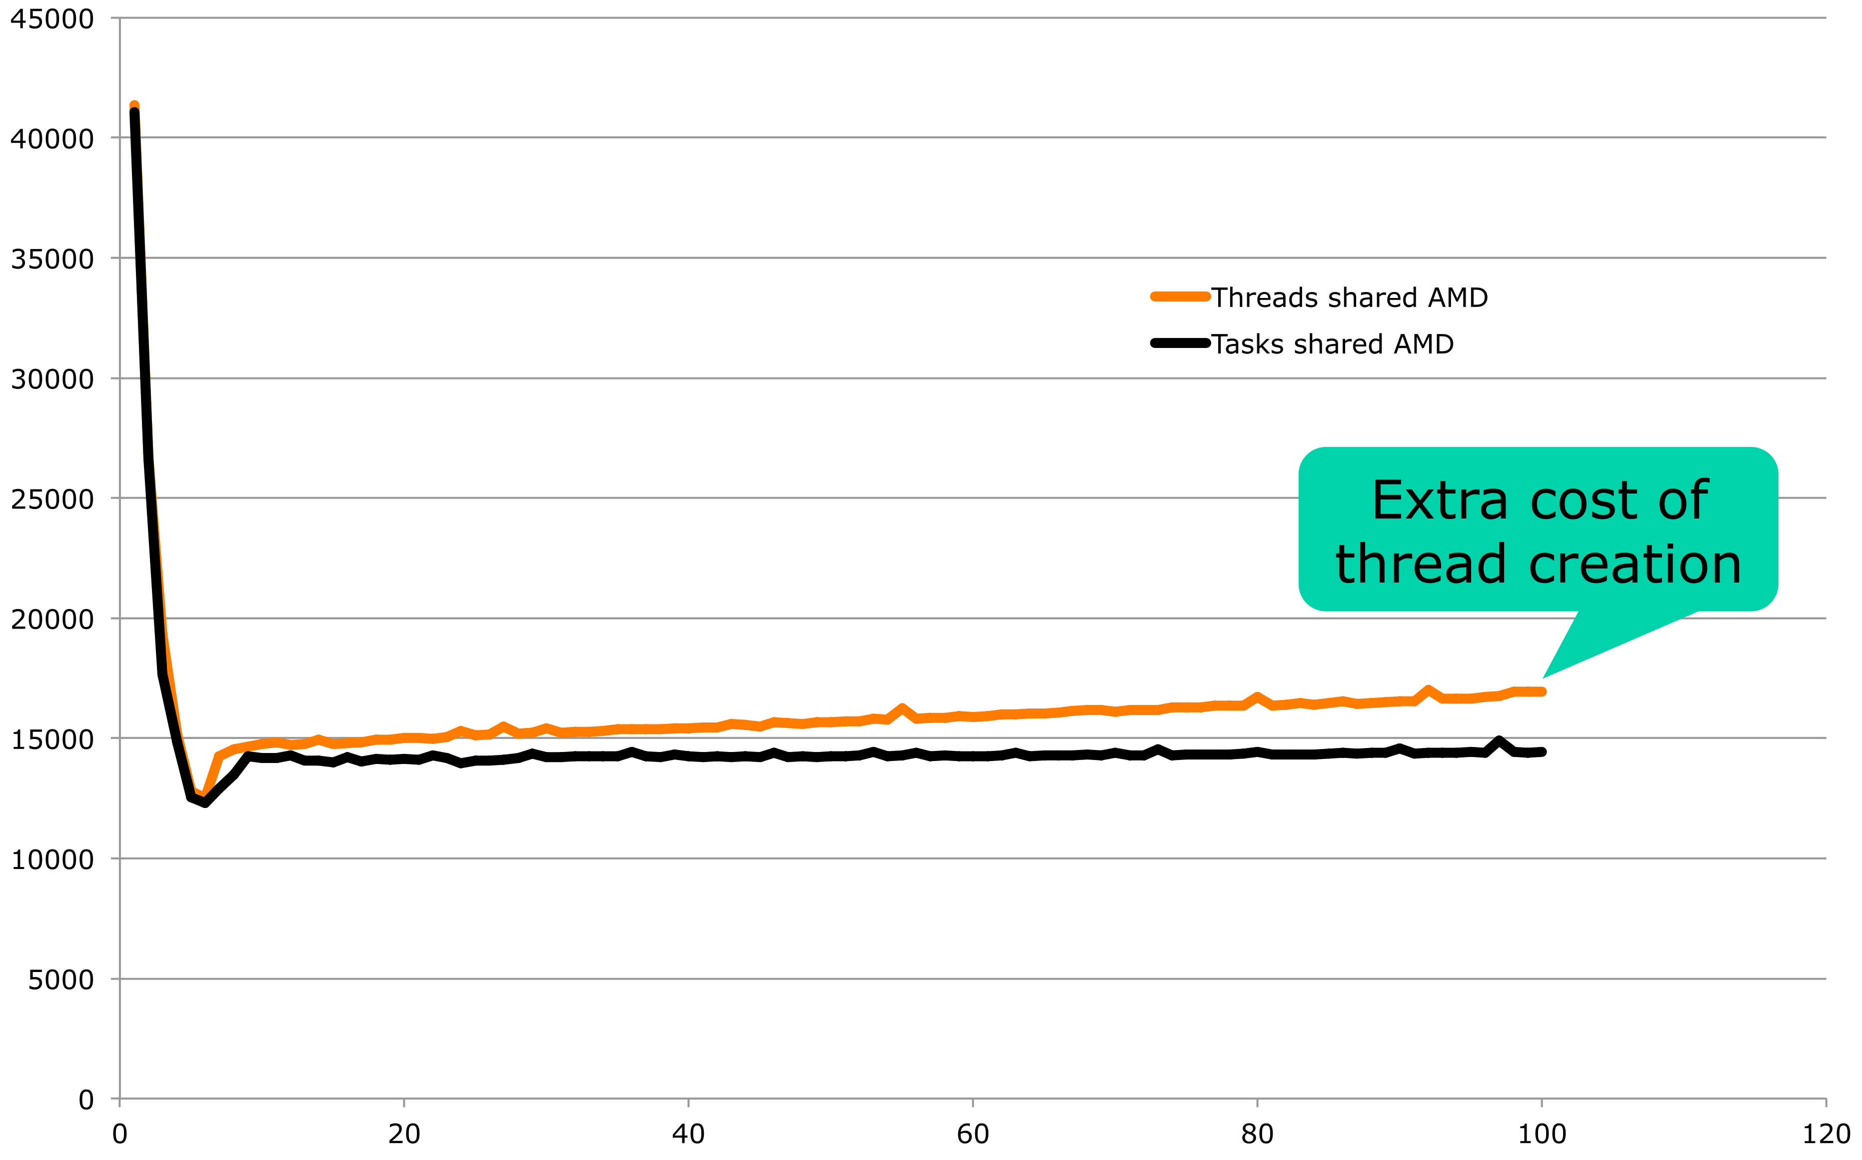

Shared counter vs local counter

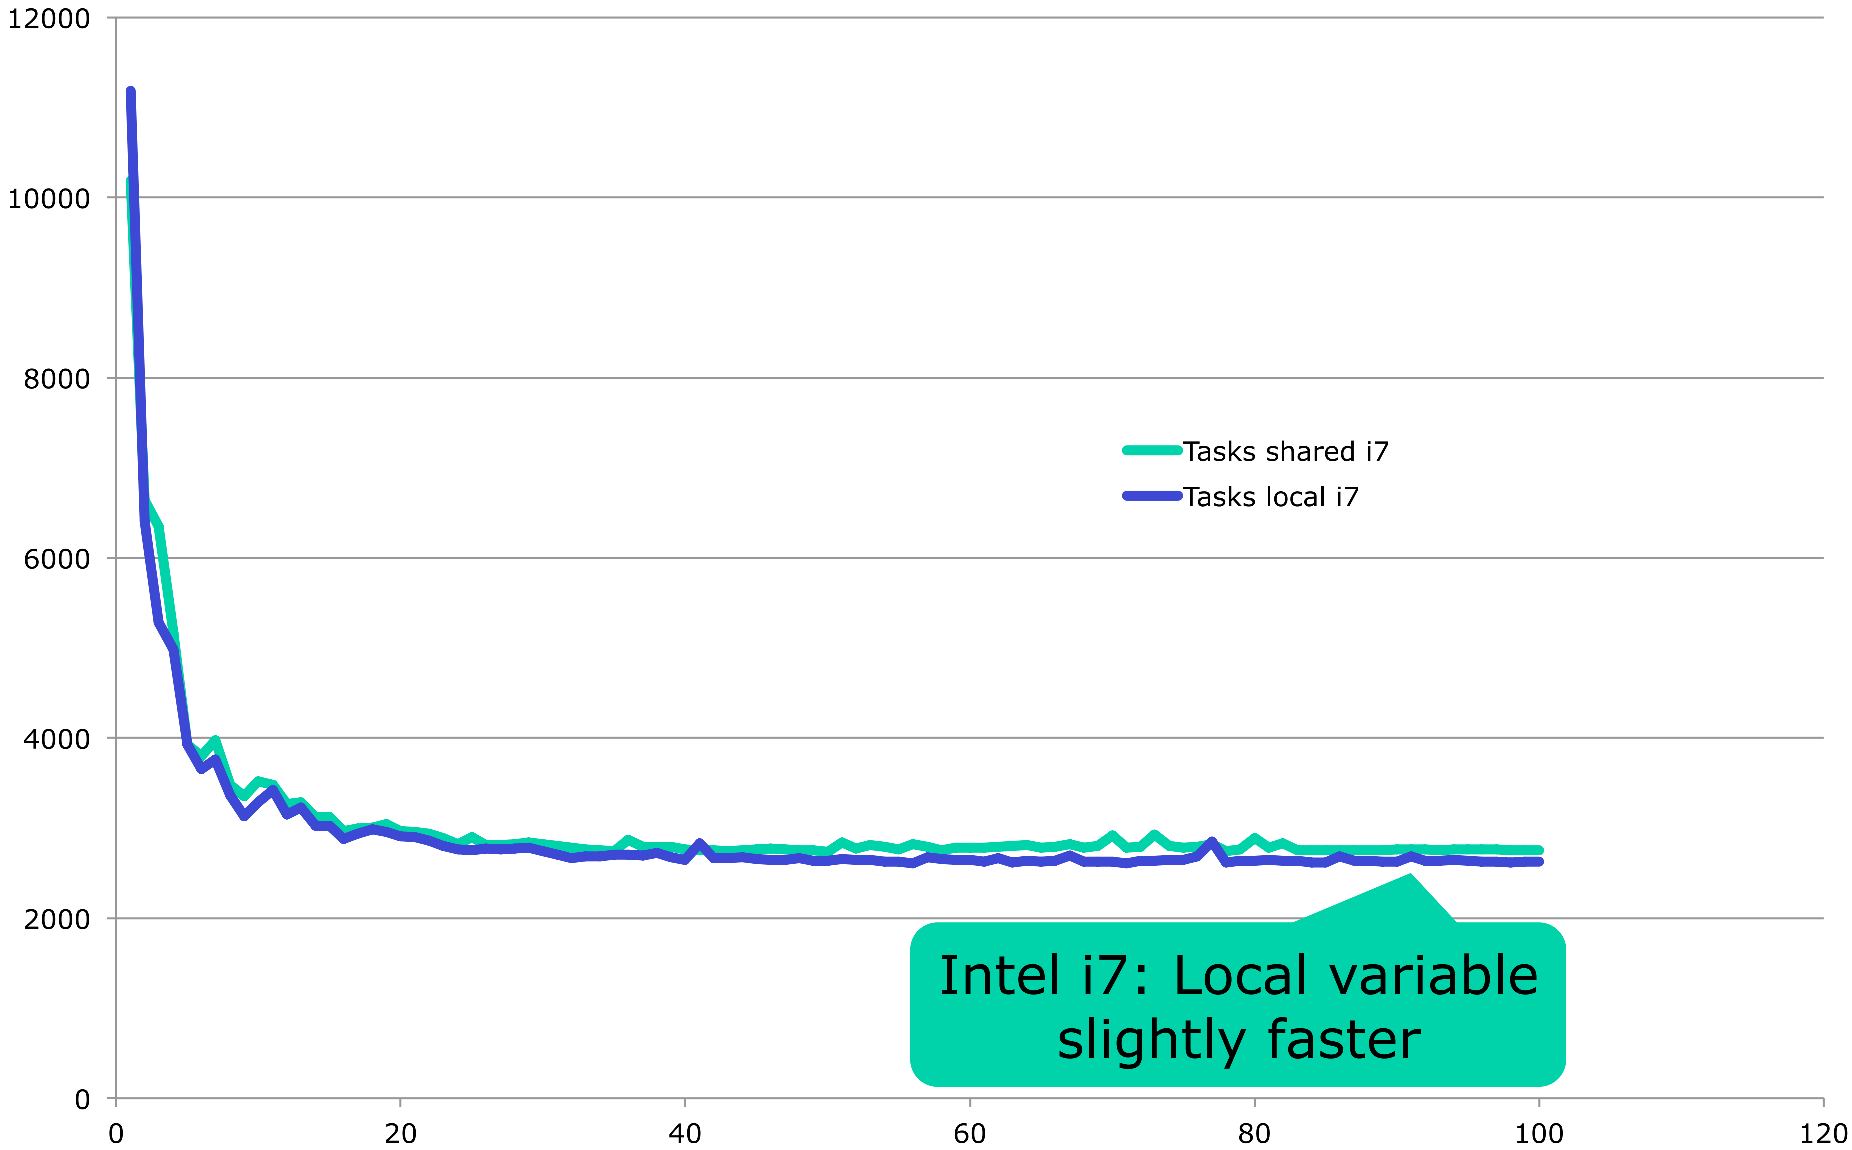

Computers differ a lot

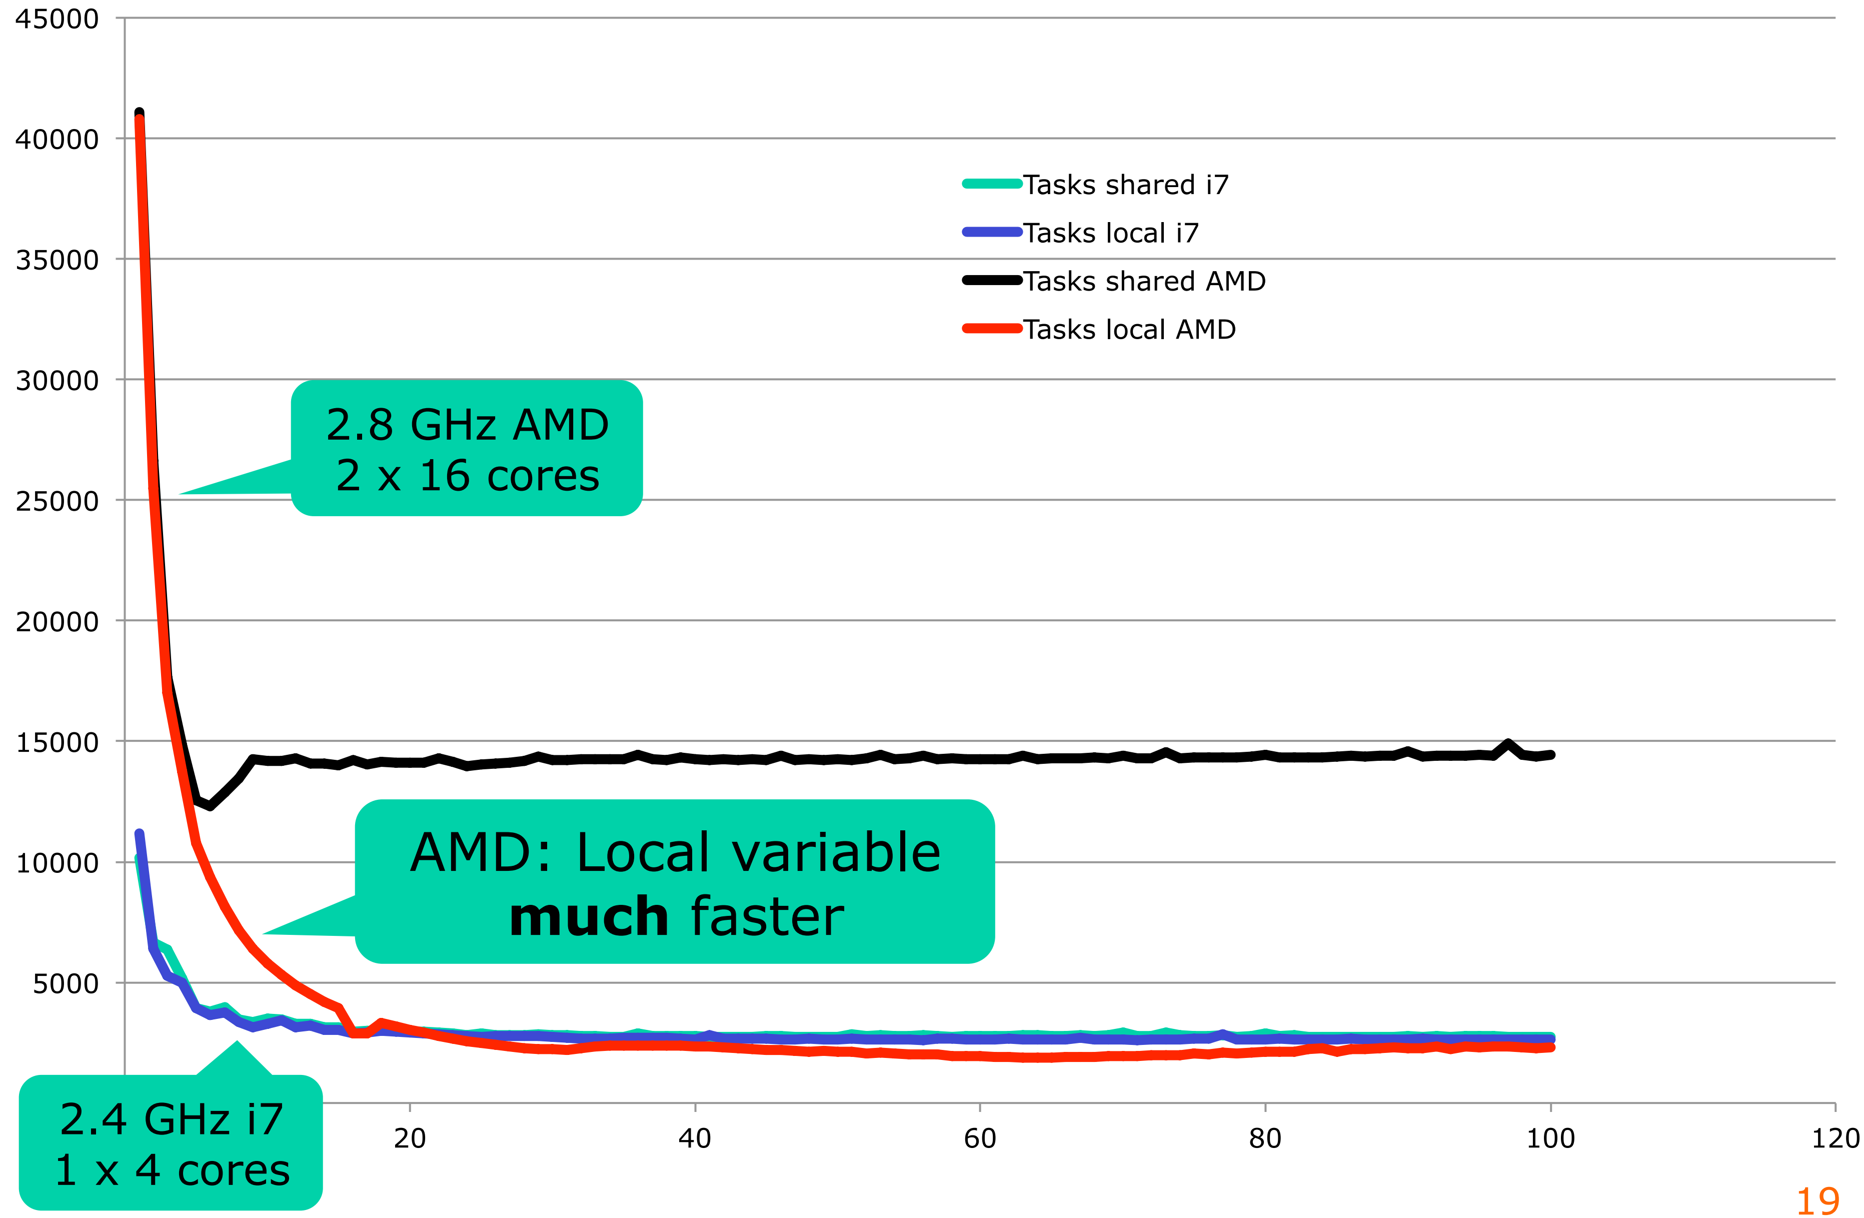

Plan for today

- Tasks and the Java executor framework
  - Executors, Runnables, Callables, Futures
- The states of a task
- Task creation overhead
- Using tasks to count prime numbers
- **Java versus the .NET Task Parallel Library**
- Producer-consumer pipelines
- Bounded queues, thread wait and notify
- The states of a thread

The .NET Task Parallel Library

- Since C#/.NET 4.0, 2010
- Easier to use and better language integration
  - **`async`** and **`await`** keywords in C#
  - .NET class library has more non-blocking methods
  - Java may get them in version 9 (2016)
- Namespace System.Threading.Tasks
- Class Task combines Runnable & Future
- Class Task<T> combines Callable<T> and Future<T>

- See *C# Precisely* chapters 22 and 23

Parallel prime counts in C#, shared

```
int perTask = range / taskCount;
LongCounter lc = new LongCounter();
Parallel.For(0, taskCount, t =>
  { int from = perTask * t,
      to = (t+1 == taskCount) ? range : perTask * (t+1);
    for (int i=from; i<to; i++)
      if (isPrime(i))
        lc.increment();
  });
return lc.get();
```

Create tasks, submit to standard executor, run

Create task t

TestCountPrimesTasks.cs

- Same concepts as in Java
  - much leaner notation
  - easier to use out of the box
- The tasks are executed on a thread pool
  - in an unknown order

Parallel prime counts in C#, local

```
long[] results = new long[taskCount];
Parallel.For(0, taskCount, t =>
  { int from = perTask * t,
    to = (t+1 == taskCount) ? range : perTask * (t+1);
    long count = 0;
    for (int i=from; i<to; i++)
      if (isPrime(i))
        count++;
    results[t] = count;
  });
return results.Sum();
```

TestCountPrimesTasks.cs

- Q: Why safe to write to **`results`** array?

Plan for today

- Tasks and the Java executor framework
  - Executors, Runnables, Callables, Futures
- The states of a task
- Task creation overhead
- Using tasks to count prime numbers
- Java versus the .NET Task Parallel Library
- **Producer-consumer pipelines**
- Bounded queues, thread wait and notify
- The states of a thread

Concurrent pipelines (Goetz §5.3)

- We parallelized prime counting by splitting the work into chunks:

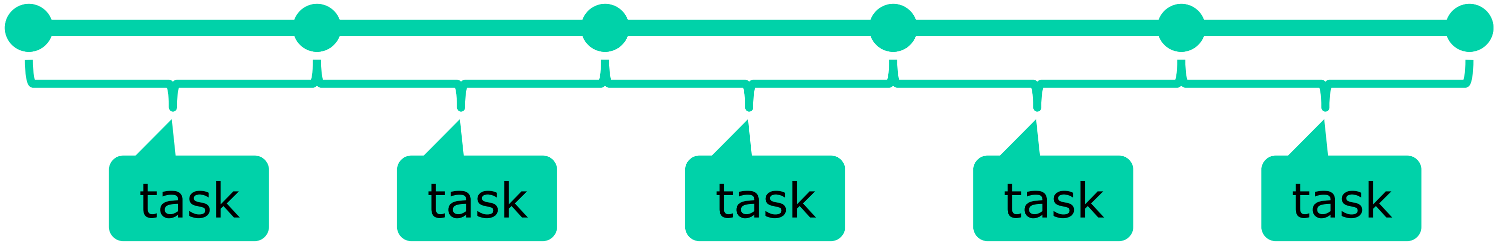

- A different way is to create a *pipeline*
- Example problem: Given long list of URLs,
  - For each URL,
  - download the webpage at that URL
  - scan the webpage for links **`<a href="link">`** …
  - for each link, print "**`url`** links to **`link`**"

Pipeline to produce URL, get webpage, scan for links, and print them

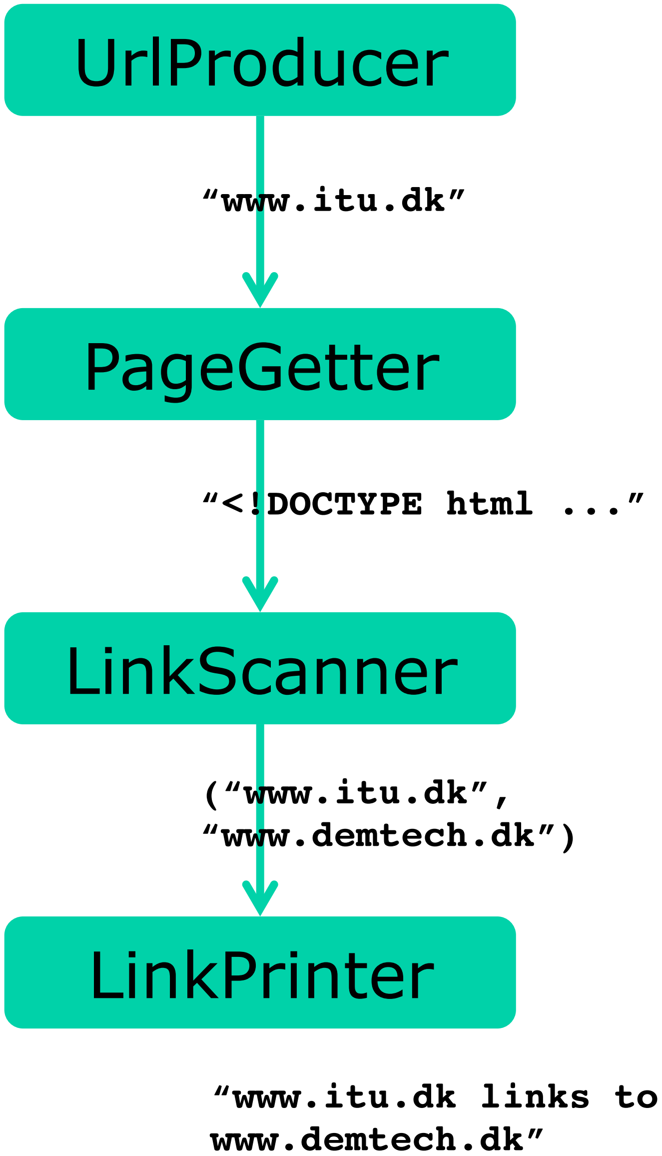

- There are four stages
- They can run in parallel
  - On four threads
  - Or as four tasks
- Each does a simple job
- Two stages communicate via a blocking queue
  - **`queue.put(item)`** sends data item to next stage; blocks until room for data
  - **`queue.take()`** gets data item from previous stage; blocks until data available

Sketch of a one-item queue

TestPipeline.java

```
interface BlockingQueue<T> {
  void put(T item);
  T take();
}
```

```
class OneItemQueue<T> implements BlockingQueue<T> {
  private T item;
  private boolean full = false;
  public void put(T item) {
    synchronized (this) {
      full = true;
      this.item = item;
    }
  }
  public T take() {
    synchronized (this) {
      full = false;
      return item;
    }
  }
}
```

Java monitor pattern, good

But: what if already full?

If queue full, we must wait for **another** thread to **`take()`** first

But: What if queue empty?

Other thread can **`take()`** only if we release lock first

Useless

Using wait() and notifyAll()

```
public void put(T item) {
  synchronized (this) {
    while (full) {
      try { this.wait(); }
      catch (InterruptedException exn) { }
    }
    full = true;
    this.item = item;
    this.notifyAll();
  }
}
```

If queue full, wait for notify from other thread

When non-full, save item, notify all waiting threads

TestPipeline.java

- **`this.wait()`**: release lock on **`this`**; do nothing until notified, then acquire lock and continue
  - Must hold lock in **`this`** before call
- **`this.notifyAll()`**: tell all threads **`wait()`**ing on **`this`** to wake up
  - Must hold lock on **`this`**, and keeps holding it

The take() method is similar

```
public T take() {
  synchronized (this) {
    while (!full) {
      try { this.wait(); }
      catch (InterruptedException exn) { }
    }
    full = false;
    this.notifyAll();
    return item;
  }
}
```

If queue empty, wait for notify from other thread

When non-empty, take item, notify all waiting threads

TestPipeline.java

- Only works if **all** methods locking on the queue are written correctly
- MUST do the wait() in a while loop; Q: Why?

Always use the wait loop idiom to invoke the wait method; never invoke it outside of a loop. The loop serves to test the condition before and after waiting.

Bloch p. 276

Java Thread states

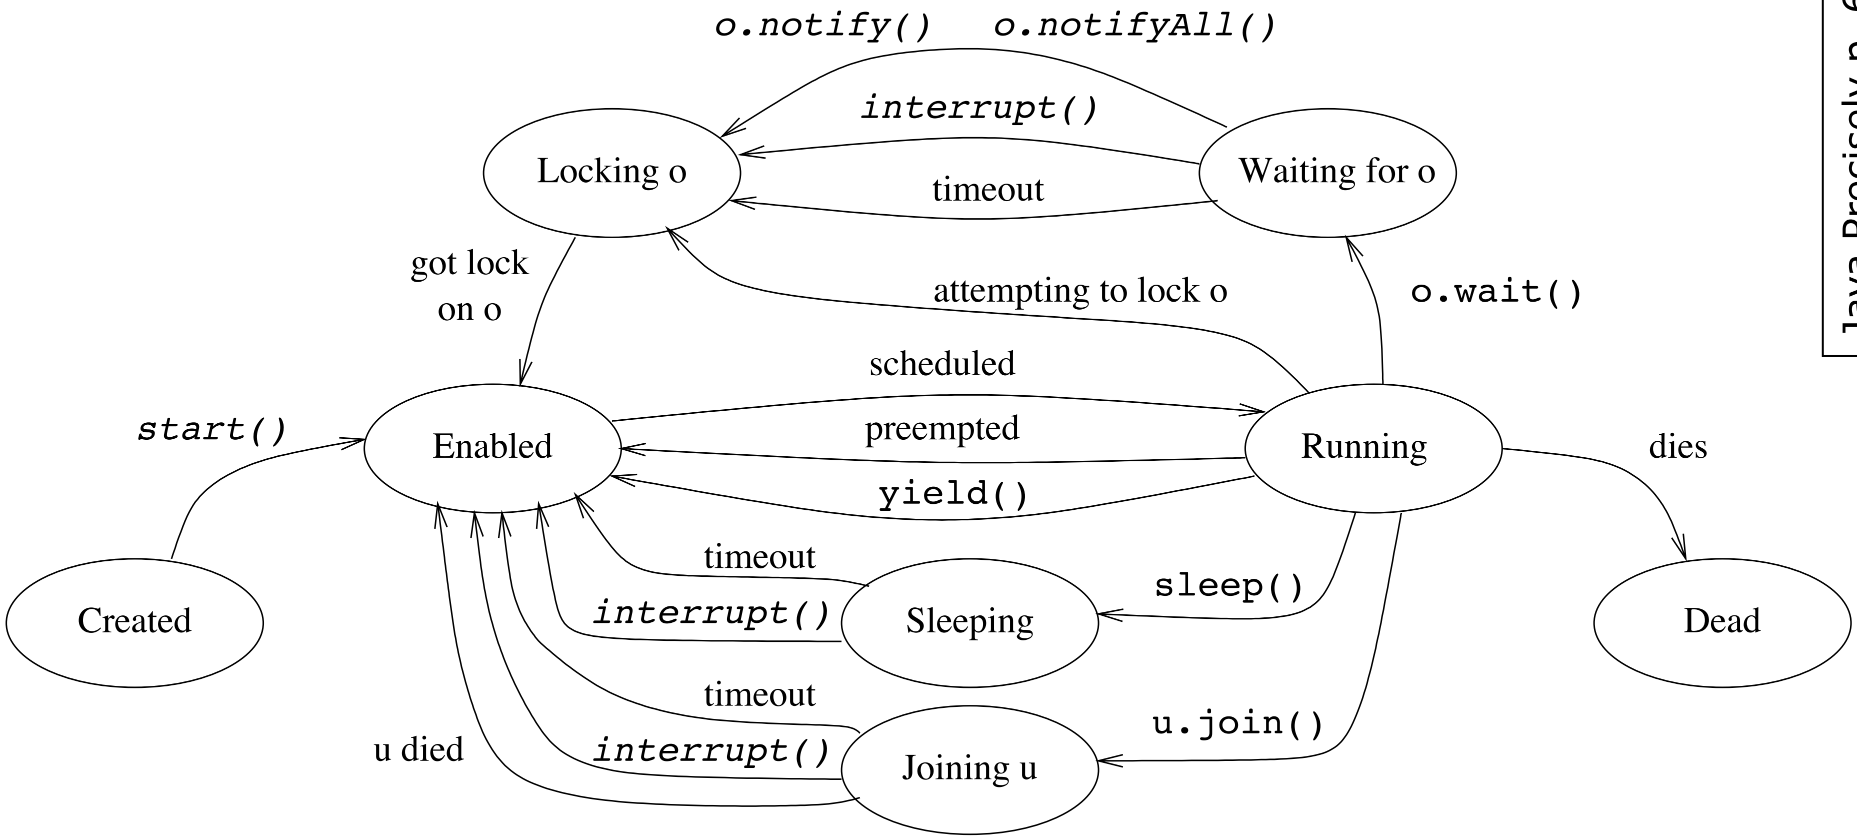

Java Precisely p. 67

- **o.wait()** is an action of the running thread itself
- ***o.notify()*** is an action by another thread, on the waiting one
- scheduled, preempted, ... are actions of the system

Producer-consumer pattern: Pipeline stages and connecting queues

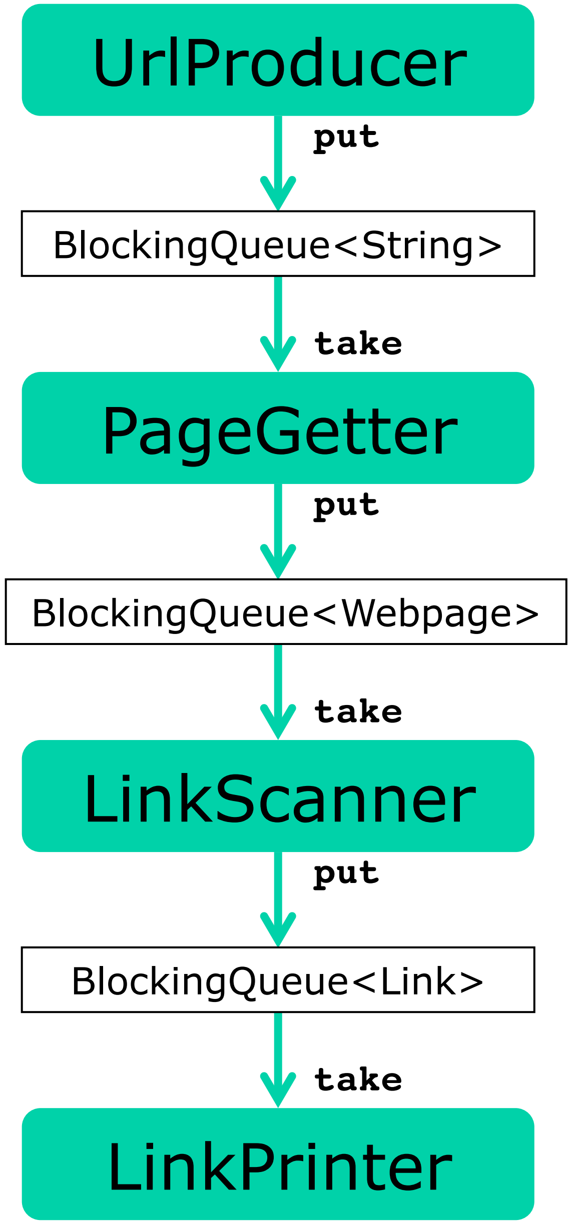

- The first stage is a producer only
- The middle stages are both consumers and producers
- The last stage is only a consumer

- A queue connects producer(s) to consumer(s) in a thread-safe way

How wait and notifyAll collaborate

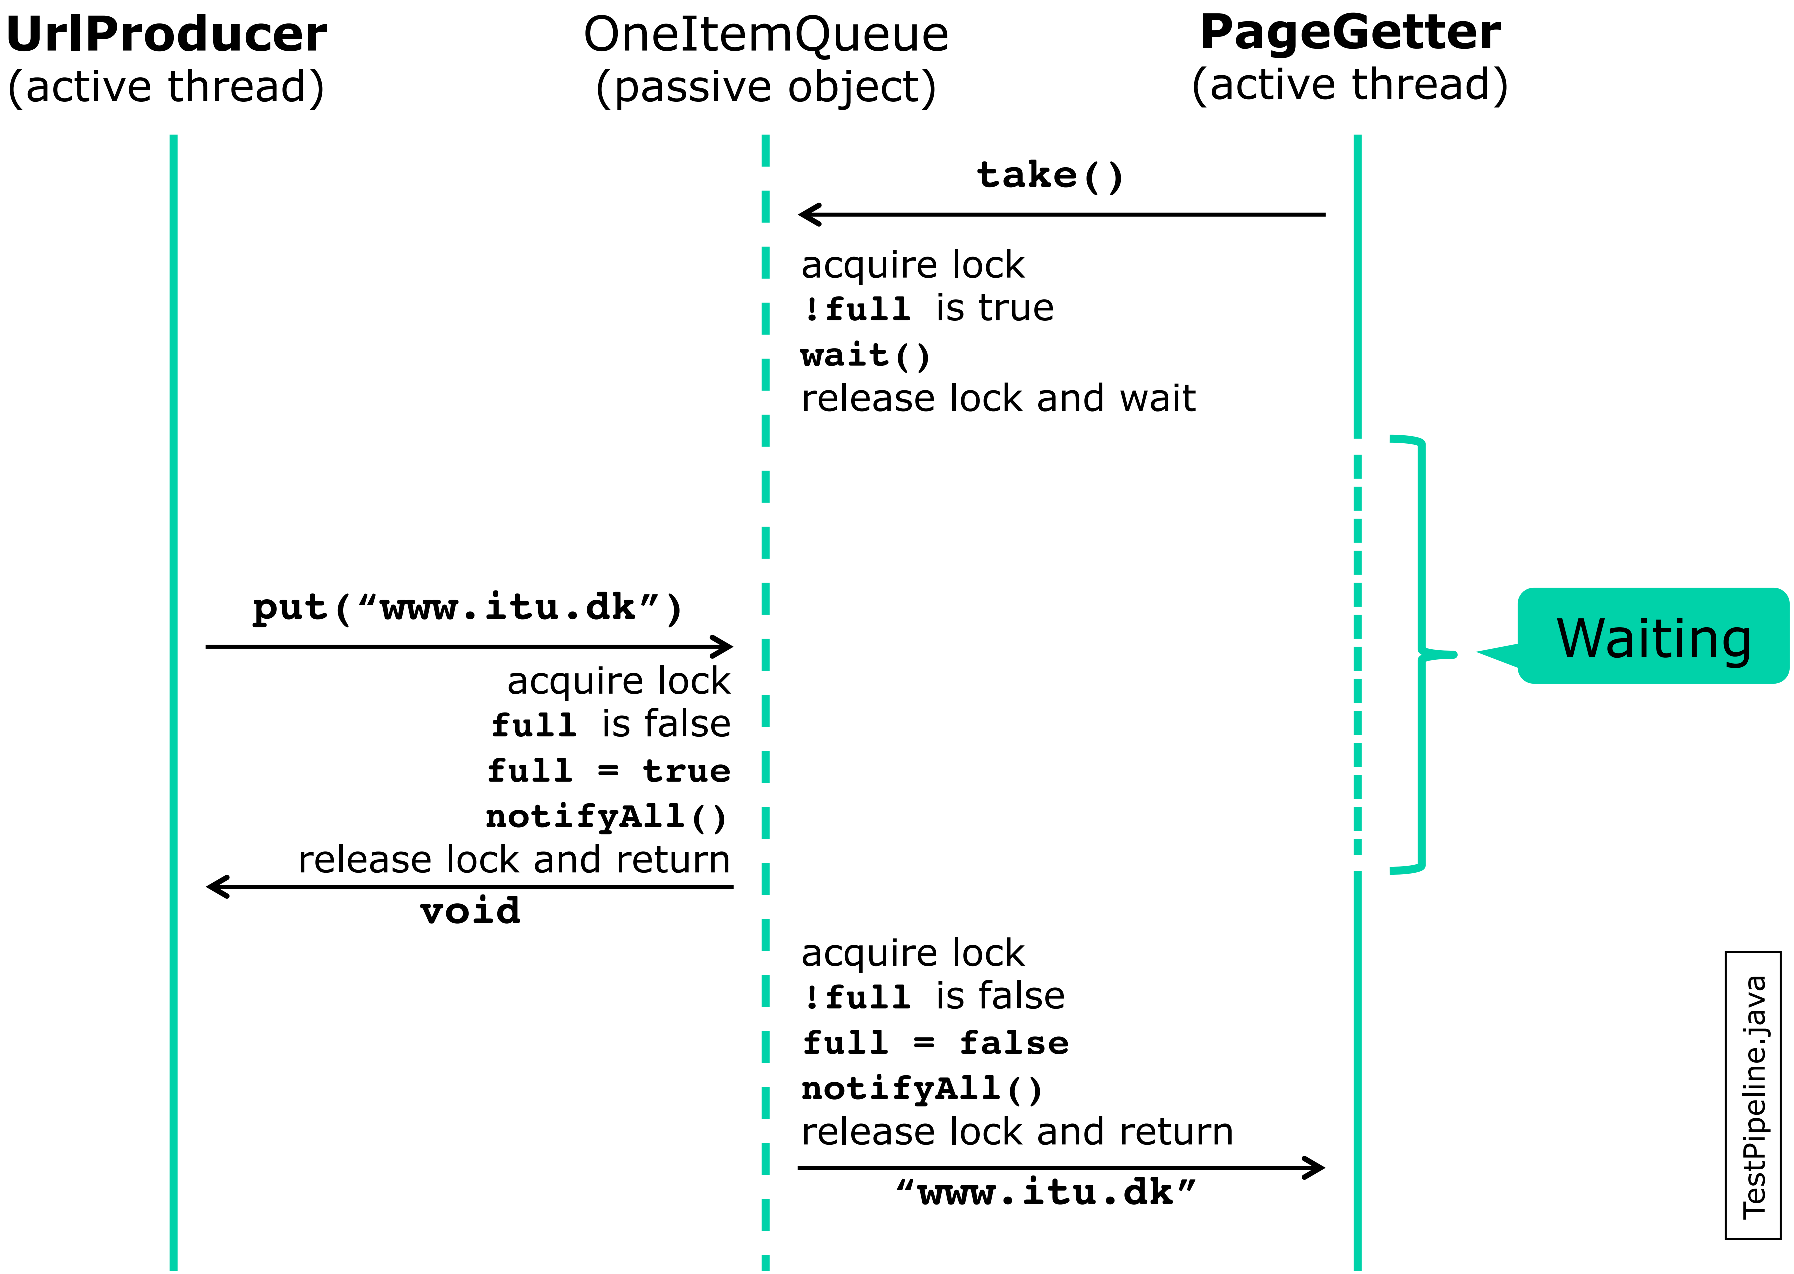

Stages 1 and 2

```
class UrlProducer implements Runnable {
  private final BlockingQueue<String> output;
  public UrlProducer(BlockingQueue<String> output) {
    this.output = output;
  }
  public void run() {
    for (int i=0; i<urls.length; i++)
      output.put(urls[i]);
  }
}
```

Produce URLs

```
class PageGetter implements Runnable {
  ...
  public void run() {
    while (true) {
      String url = input.take();
      try {
        String contents = getPage(url, 200);
        output.put(new Webpage(url, contents));
      } catch (IOException exn) { System.out.println(exn); }
    }
  }
```

Transform URL to webpage

TestPipeline.java

Stages 3 and 4

```
class LinkScanner implements Runnable {
  ...
  private final static Pattern urlPattern
    = Pattern.compile("a href=\"(\\p{Graph}*)\"");
  public void run() {
    while (true) {
      Webpage page = input.take();
      Matcher urlMatcher = urlPattern.matcher(page.contents);
      while (urlMatcher.find()) {
        String link = urlMatcher.group(1);
        output.put(new Link(page.url, link));
} } } }
```

Transform web page to link stream

```
class LinkPrinter implements Runnable {
  ...
  public void run() {
    while (true) {
      Link p = input.take();
      System.out.printf("%s links to %s%n", p.from, p.to);
} } }
```

Consume links and print them

Putting stages and queues together

```
final BlockingQueue<String> urls = new OneItemQueue<String>();
final BlockingQueue<Webpage> pages = new OneItemQueue<Webpage>();
final BlockingQueue<Link> refPairs = new OneItemQueue<Link>();
Thread t1 = new Thread(new UrlProducer(urls));
Thread t2 = new Thread(new PageGetter(urls, pages));
Thread t3 = new Thread(new LinkScanner(pages, refPairs));
Thread t4 = new Thread(new LinkPrinter(refPairs));
t1.start(); t2.start(); t3.start(); t4.start();
```

TestPipeline.java

- Each stage does *one* job
  - Simple to implement and easy to modify
  - Separation of concerns, simple control flow
- Easy to add new stages
  - For instance, discard duplicate links
- Can achieve high throughput
  - May run multiple copies of a slow stage

"Prefer concurrency utilities to wait and notify"

Bloch item 69

- It's instructive to use **`wait`** and **`notify`**
- … but easy to do it wrong
- Package java.util.concurrent has
  - BlockingQueue<T> interface
  - ArrayBlockingQueue<T> class and much more
- Better use those in practice

- Next week: same pipeline with Java 8 streams
  - Simpler, and *very* easy to parallelize
  - (But of course still needs to be thread-safe)

This week

- Reading
  - Goetz et al chapters 5.3, 6 and 8
  - Bloch items 68, 69

- Exercises week 5
  - Show that you can use tasks and the executor framework, and modify a concurrent pipeline

- Read before next week's lecture
  - Goetz chapter 10
  - Bloch item 67