

CS 10:


# Problem solving via Object Oriented Programming

Synchronization

# Main goals


- Implement mechanisms to handle threads that are accessing shared resources at the same time

# Agenda

- 
1. Threads and interleaving execution
  2. Producer/consumer
  3. Deadlock, starvation

# Threads are a way for multiple processes to run concurrently


## Threads



```
main() {  
MyThread t = new MyThread();  
  
//start thread at run method, main  
thread keeps running  
t.start()  
  
//halt main until thread finishes  
t.join()  
}
```

# Concurrent threads can access the same resources; this can cause problems

## Concurrency



# Let's make it interesting, what is the final value of total?

## Incrementer.java

```
7 public class Incrementer extends Thread {
8     private static int total = 0; // a variable shared by all incrementers
9     private static final int times = 1000000; // how many times to increment total, in each thread
10
11     /**
12      * Increments total the specified number of times
13      */
14     public void run() {
15         for (int i = 0; i < times; i++) {
16             total++;
17         }
18     }
19
20     public static void main(String [] args) throws Exception {
21         Incrementer inc1 = new Incrementer();
22         Incrementer inc2 = new Incrementer();
23
24         // Fire off threads and wait for them to complete
25         inc1.start();
26         inc2.start();
27         inc1.join();
28         inc2.join();
29
30         System.out.println("total at end = " + total);
31     }
32 }
```

One million (not a trick)

What will *total* be at end?  
Top three guesses?

# Threads can be interrupted at any point, this can cause unexpected behavior

## Incrementer.java

```
7 public class Incrementer extends Thread {
8     private static int total = 0;                // a variable shared by all incrementers
9     private static final int times = 1000000;    // how many times to increment total, in each thread
10
11     /**
12      * Increments total the specified number of times
13      */
14     public void run() {
15         for (int i = 0; i < times; i++) {
16             total++;
17         }
18     }
19
20     public static void main(String [] args) throws Exception {
21         Incrementer inc1 = new Incrementer();
22         Incrementer inc2 = new Incrementer();
23
24         // Fire off threads and wait for them to complete
25         inc1.start();
26         inc2.start();
27         inc1.join();
28         inc2.join();
29
30         System.out.println("total at end = " + total);
31     }
32 }
```

# IncrementerInterleaving.java demonstrates interruptions (sometimes)

## IncrementerInterleaving.java

```
6 public class IncrementerInterleaving extends Thread {
7     private static int total = 0;           // a variable shared by all incrementers
8     private static final int times = 5;    // how many times to increment total, in each thread
9     private String name;                   // for display purposes
10
11     public IncrementerInterleaving(String name) {
12         this.name = name;
13     }
14
15     /**
16      * Increments total the specified number of times
17      */
18     public void run() {
19         for (int i = 0; i < times; i++) {
20             int temp = total;
21             System.out.println(name + " gets " + temp);
22             temp = temp + 1;
23             total = temp;
24             System.out.println(name + " puts " + temp);
25         }
26     }
27
28     public static void main(String [] args) throws Exception {
29         IncrementerInterleaving inc1 = new IncrementerInterleaving("one");
30         IncrementerInterleaving inc2 = new IncrementerInterleaving("two");
31
32         // Fire off threads and wait for them to complete
33         inc1.start();
34         inc2.start();
35         inc1.join();
36         inc2.join();
37
38         System.out.println("total at end = " + total);
39     }
40 }
41
```



# DEMO: IncrementerInterleaving.java

- Run several times
- Interrupted execution causes tricky bugs
- Sometimes it works as expected
- Sometimes it doesn't...

# Java provides the keyword `synchronized` to make some operations “atomic”

## IncrementerTotal.java

```
public class IncrementerTotal {  
    private int total = 0;  
    public synchronized void inc() {  
        total++;  
    }  
}
```

- `synchronized` keyword in front of `inc` method means only one thread can be running this code at a time
- If multiple threads try to run `synchronized` code, one thread runs, all others block until first one finishes
- Once first thread finishes, OS selects another thread to run
- `synchronized` makes this code “atomic” (e.g., as if it were one instruction)
- This `synchronized` approach is called a “mutex” (or monitor), acts like a “lock” on static `total` variable

# IncrementerSync.java uses atomic operations to ensure desired behavior

## IncrementerSync.java

```
8 public class IncrementerSync extends Thread {
9     private static IncrementerTotal total = new IncrementerTotal(); // a variable shared by all incrementers
10    private static final int times = 1000000; // how many times to increment total, in each thread
11
12    /**
13     * Increments total the specified number of times
14     */
15    public void run() {
16        for (int i = 0; i < times; i++) {
17            total.inc();
18        }
19    }
20
21    public static void main(String [] args) throws Exception {
22        IncrementerSync inc1 = new IncrementerSync();
23        IncrementerSync inc2 = new IncrementerSync();
24
25        // Fire off threads and wait for them to complete
26        inc1.start();
27        inc2.start();
28        inc1.join();
29        inc2.join();
30
31        System.out.println("total at end = " + total.total);
32    }
33 }
```

```
public class IncrementerTotal {
    private int total = 0;
    public synchronized void inc() {
        total++;
    }
}
```

# Agenda

1. Interleaving execution

 2. Producer/consumer

3. Deadlock, starvation

# Producers tell Consumers when ready, Consumers tell Producers when done

**Big idea: keep Producers and Consumers in sync**

## **Producer:**

- Tell Consumer when item is ready (`notify` or `notifyAll`)
- Block until woken up by Consumer that item handled (`wait`)
- Tell Consumer when next item is ready (`notify` or `notifyAll`)
- There can be multiple Producers

## **Consumer:**

- Block until woken up by Producer that item ready (`wait`)
- Process item and tell Producer when done (`notify` or `notifyAll`)
- Block until woken up by Producer (`wait`)
- There can be multiple Consumers

# Producers and Consumers synchronized with `wait`, `notify` or `notifyAll`

## `wait()`

- Pauses and removes Thread from synchronized method
- Tells Operating System to put this Thread into a list of Threads waiting to resume execution
- `wait()` allows another Thread to enter synchronized method

## `notify()`

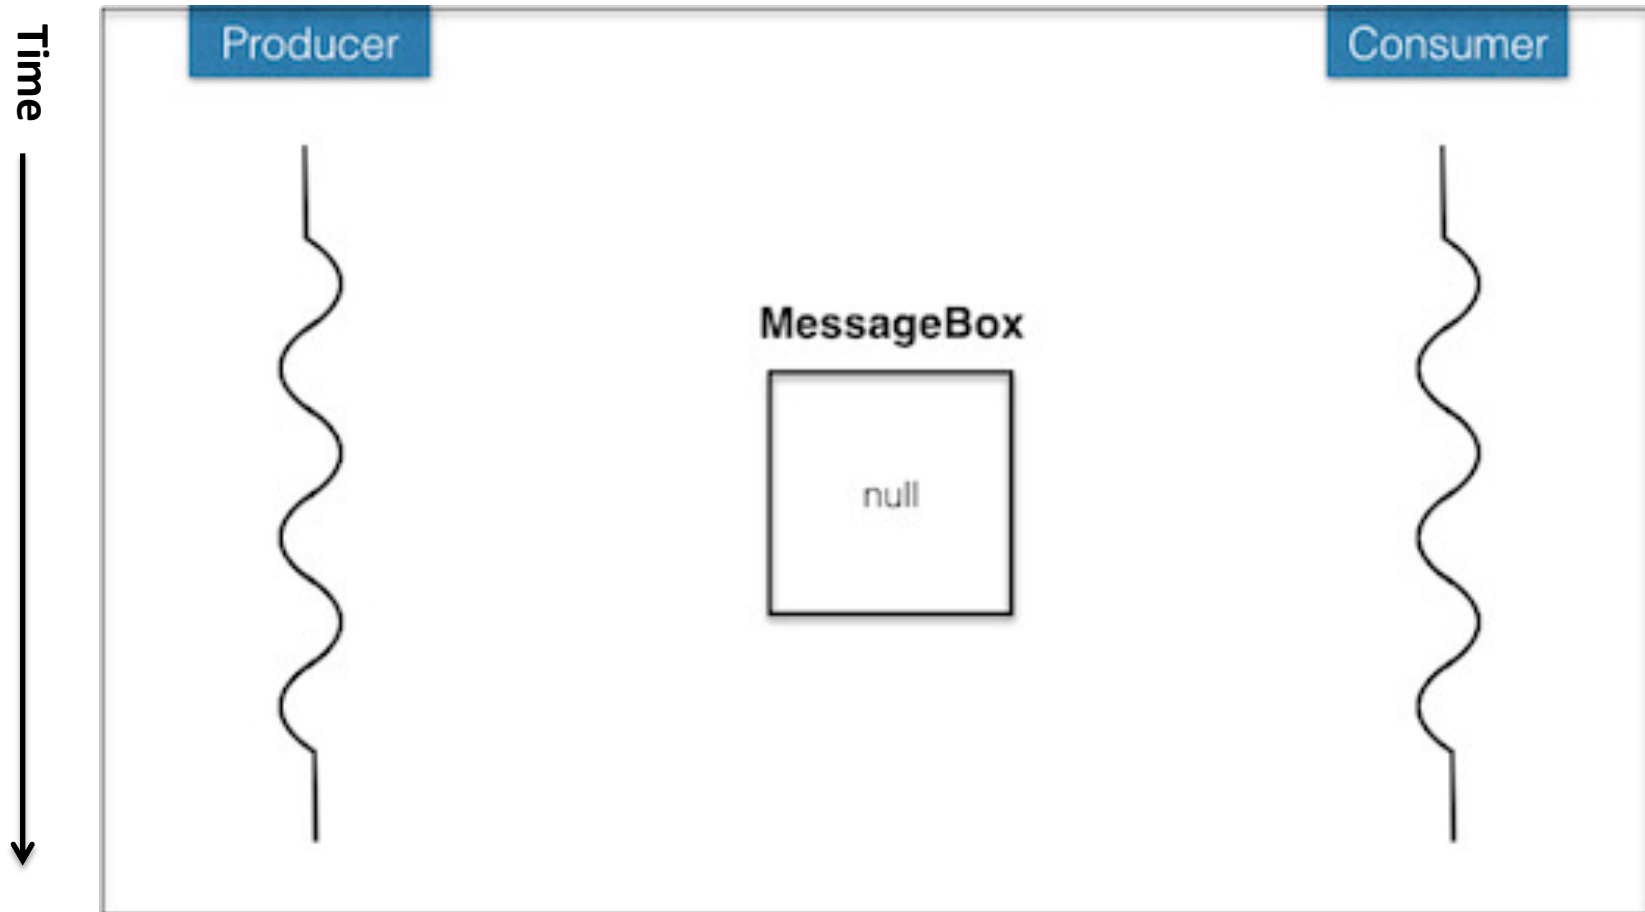
- Tells Operating System to pick a waiting Thread and let it run again (not a FIFO queue, OS decides – take CS58 for more)
- Thread should check that conditions are met for it to continue

## `notifyAll()`

- Wake up all waiting Threads
- Each Thread should check that conditions are met for it to continue

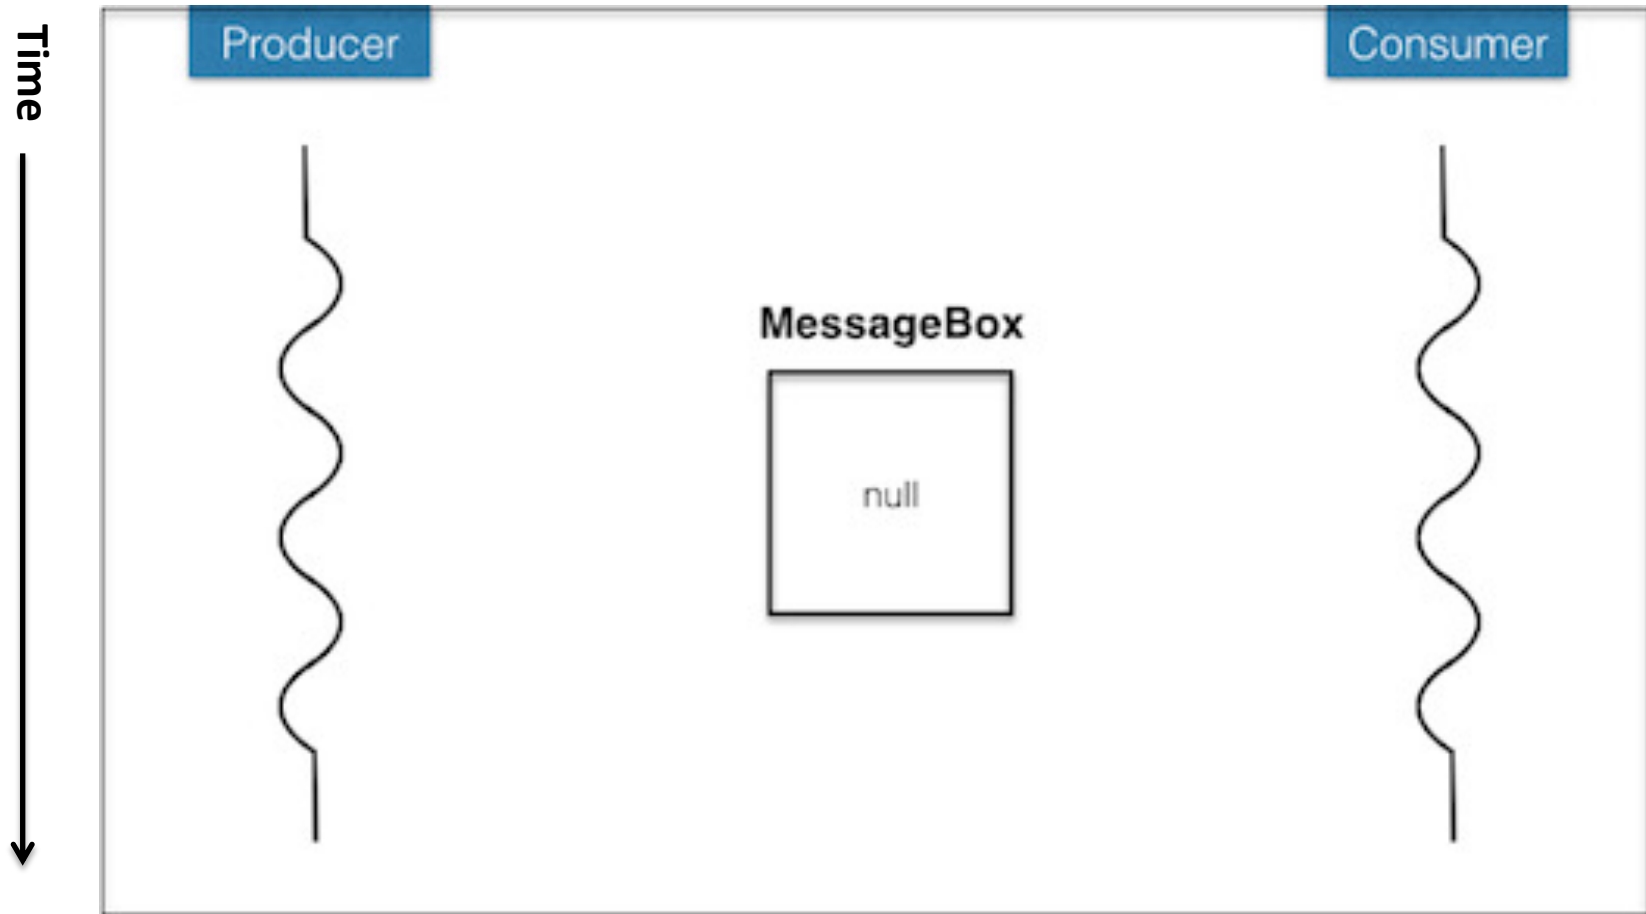
# Scenario: Producers produce messages for Consumers, need to keep in sync

## Example



# We can use a semaphore to keep Producers and Consumers in sync

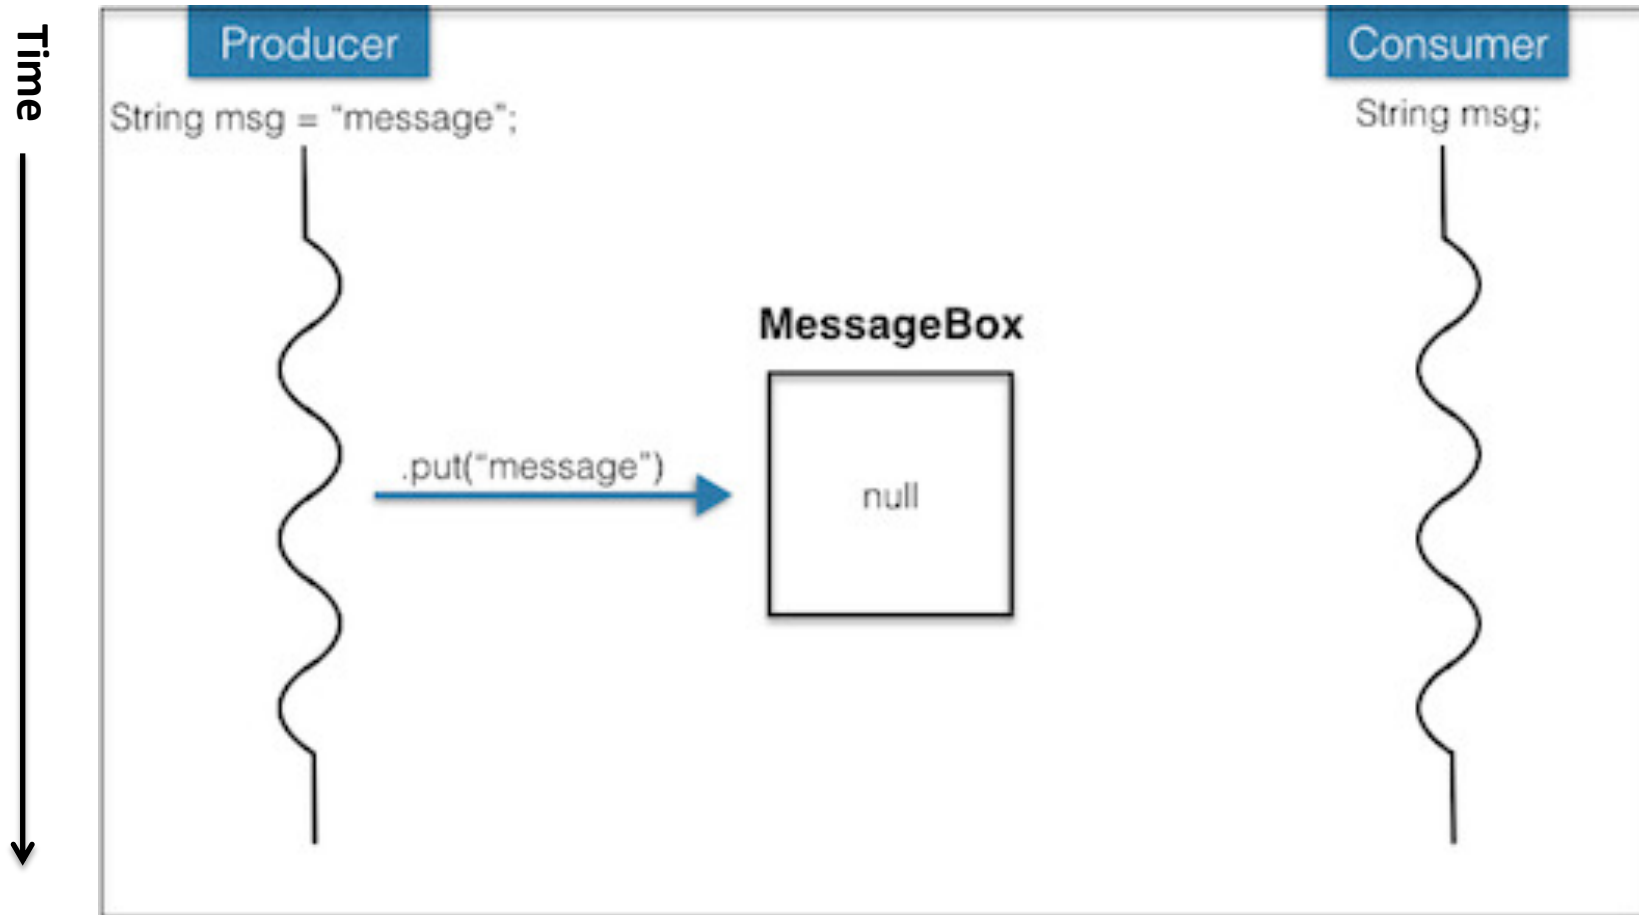
## Example





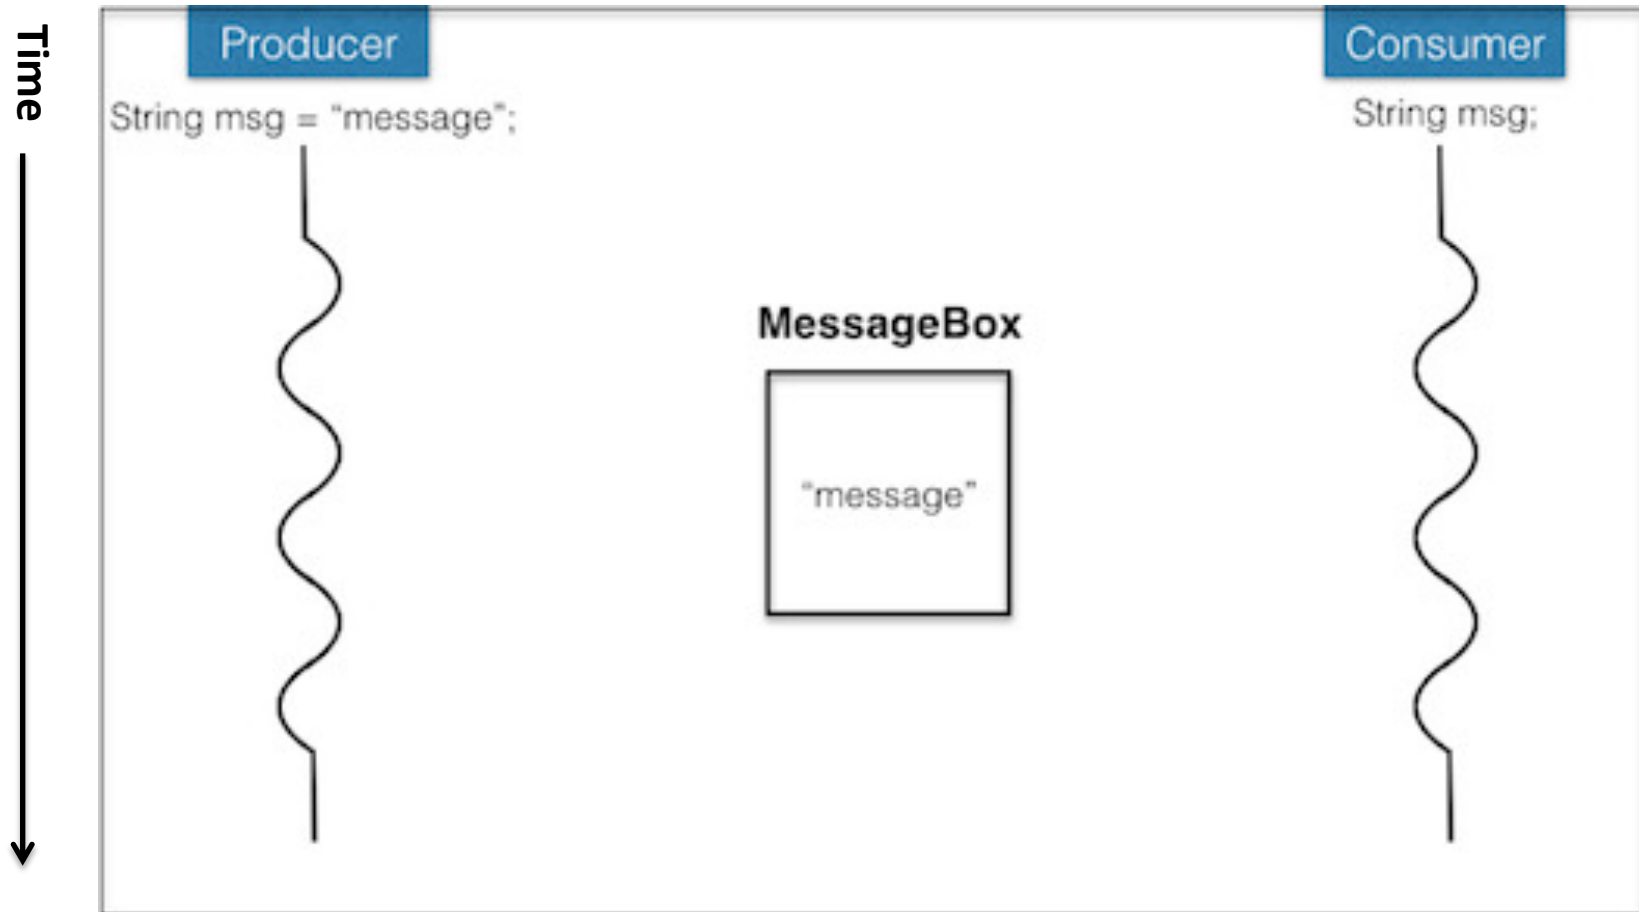
# Producer passing messages to Consumer using semaphore

## Example



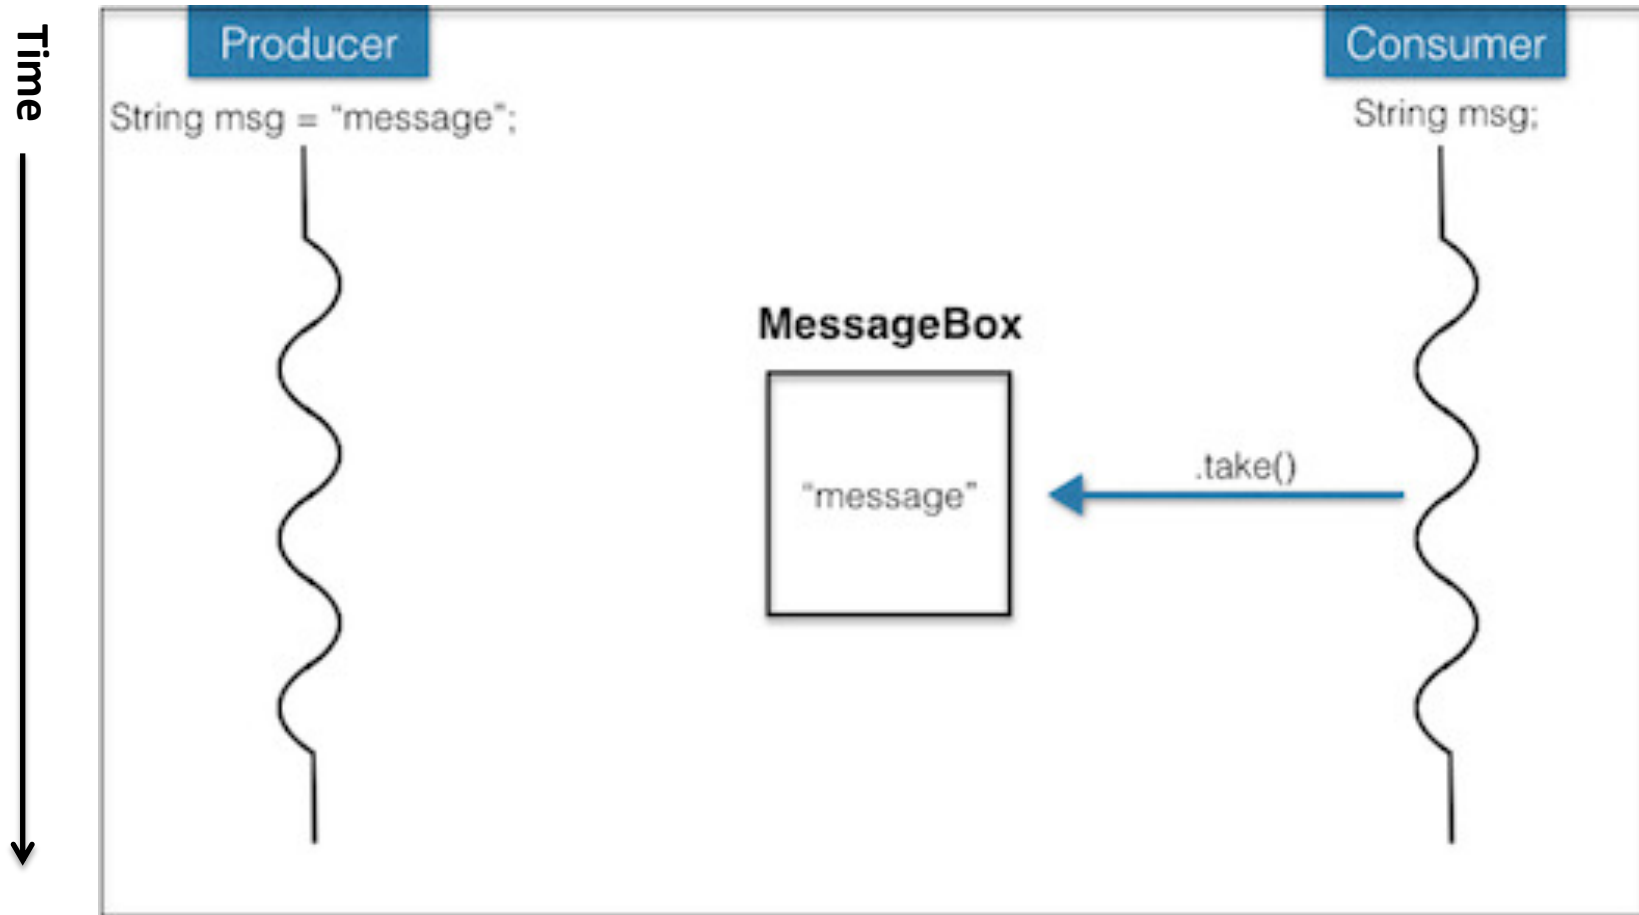
# Producer passing messages to Consumer using semaphore

## Example



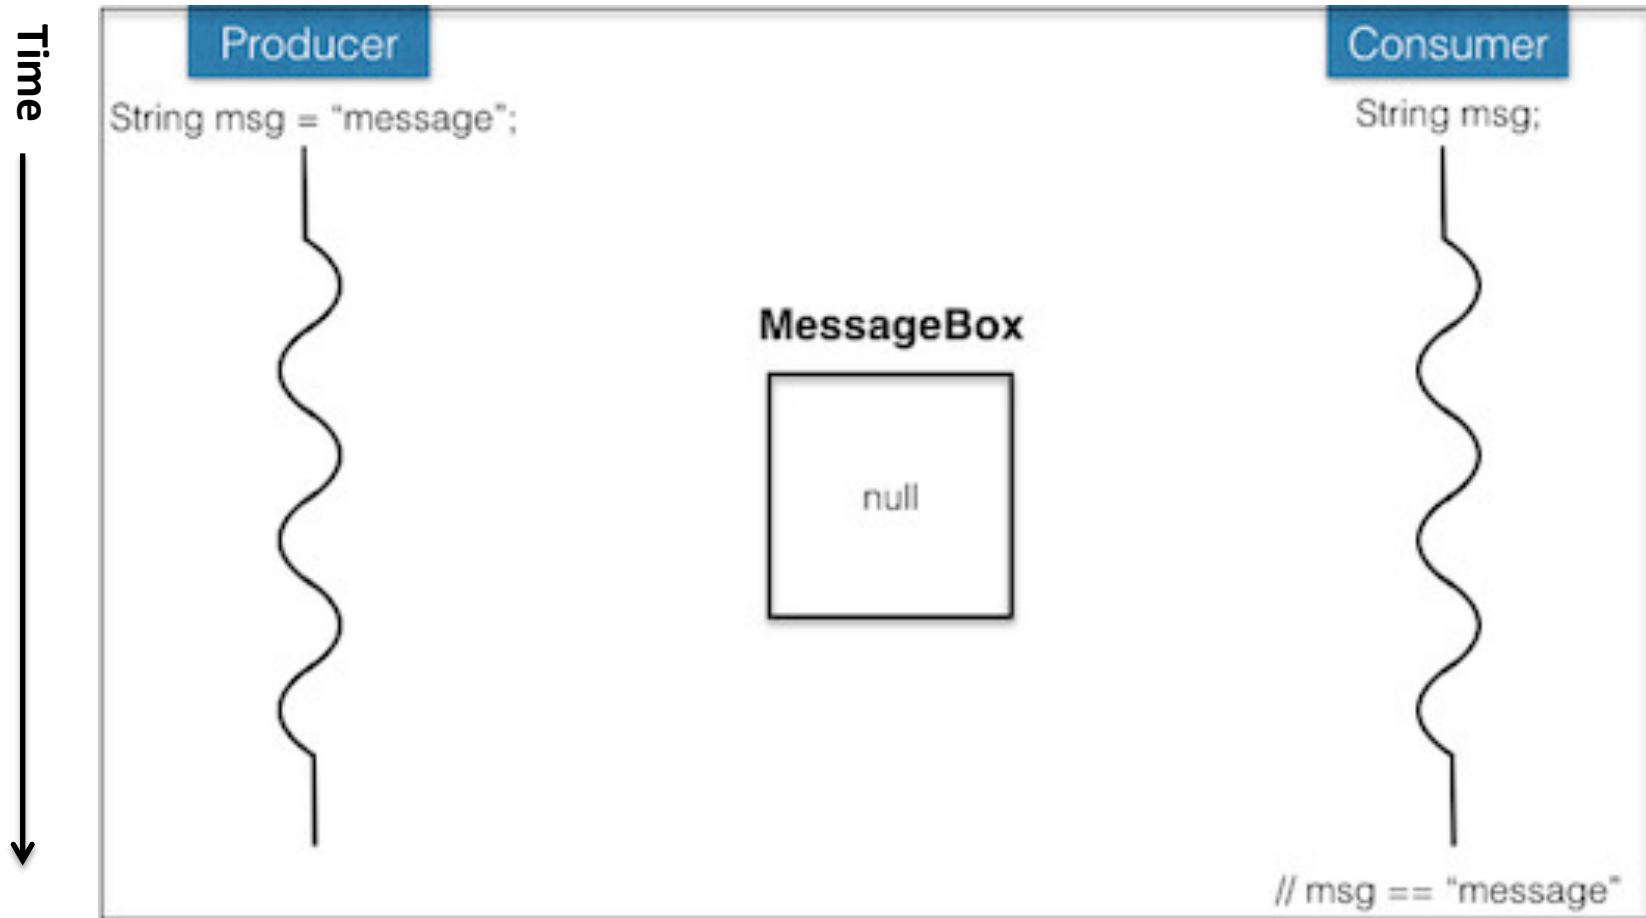
# Producer passing messages to Consumer using semaphore

## Example



# Producer passing messages to Consumer using semaphore


## Example



# MessageBox.java implements a semaphore that holds one String


## MessageBox.java

```
7 public class MessageBox {
8     private String message = null;
9
10    /**
11     * Put m as message once it's okay to do so (current message has been taken)
12     */
13    public synchronized void put(String m) throws InterruptedException {
14        //check to see if message is not null, might have been woken by put() notifyAll
15        while (message != null) {
16            wait();
17        }
18        message = m;
19        notifyAll(); //wakes producers AND consumers
20    }
21
22    /**
23     * Takes message once it's there, leaving empty message
24     */
25    public synchronized String take() throws InterruptedException {
26        //check to see if message is null, might have been woken by take() notifyAll
27        while (message == null) {
28            wait();
29        }
30        String m = message;
31        message = null;
32        notifyAll(); //wakes producers AND consumers
33        return m;
34    }
35 }
```



# MessageBox.java implements a semaphore that holds one String


## MessageBox.java



```
7 public class MessageBox {
8     private String message = null;
9
10    /**
11     * Put m as message once it's okay to do so (current message has been taken)
12     */
13    public synchronized void put(String m) throws InterruptedException {
14        //check to see if message is not null, might have been woken by put() notifyAll
15        while (message != null) {
16            wait();
17        }
18        message = m;
19        notifyAll(); //wakes producers AND consumers
20    }
21
22    /**
23     * Takes message once it's there, leaving empty message
24     */
25    public synchronized String take() throws InterruptedException {
26        //check to see if message is null, might have been woken by take() notifyAll
27        while (message == null) {
28            wait();
29        }
30        String m = message;
31        message = null;
32        notifyAll(); //wakes producers AND consumers
33        return m;
34    }
35 }
```

# Producers use MessageBox to pass messages to Consumers


## Producer.java



```
6 public class Producer extends Thread {
7     private MessageBox box;
8     private int numberToSend;
9
10    public Producer(MessageBox box, int numberToSend) {
11        this.box = box;
12        this.numberToSend = numberToSend;
13    }
14
15    /**
16     * Wait for a while then puts a message
17     * Puts "EOF" when # messages have been put
18     */
19    public void run() {
20        try {
21            for (int i = 0; i < numberToSend; i++) {
22                sleep((int)(Math.random()*5000)); //sleep for random time up to 5 seconds
23                box.put("message #" + i); //put a new message in MessageBox
24            }
25            box.put("EOF"); //EOF means end of file
26        }
27        catch (InterruptedException e) {
28            System.err.println(e);
29        }
30    }
31 }
```

# Consumers retrieve messages from the MessageBox

## Consumer.java




```
6 public class Consumer extends Thread {
7     private MessageBox box;
8
9     public Consumer(MessageBox box) {
10         this.box = box;
11     }
12
13     /**
14      * Takes messages from the box and prints them, until receiving EOF
15      */
16     public void run() {
17         try {
18             String message;
19             while (!(message = box.take()).equals("EOF")) {
20                 System.out.println(message);
21             }
22         }
23         catch (InterruptedException e) {
24             System.err.println(e);
25         }
26     }
27 }
```



# ProducerConsumer uses all three components to pass messages

## ProducerConsumer.java

```
8 public class ProducerConsumer {
9     public static final int numMessages = 5; // how many messages to send from produc
10    private Producer producer;
11    private Consumer consumer;
12
13    public ProducerConsumer() {
14        MessageBox box = new MessageBox();
15        producer = new Producer(box, numMessages);
16        consumer = new Consumer(box);
17    }
18
19    /**
20     * Just starts the producer and consumer running
21     */
22    public void communicate() {
23        producer.start();
24        consumer.start();
25    }
26
27    public static void main(String[] args) {
28        new ProducerConsumer().communicate();
29        System.out.println("Peace out! (threads are still running but I'm done)");
30    }
31 }
```



Problems Javadoc Declaration Console Debug Expressions Error Log Console Call Hierarchy  
<terminated> ProducerConsumer [Java Application] /Library/Java/JavaVirtualMachines/jdk1.8.0\_112.jdk/Contents/Home/bin/java (Feb 22, 2018, 11:55:46 AM)

Peace out! (threads are still running but I'm done)

message #0  
message #1  
message #2  
message #3  
message #4

**main() ends, but Producers and Consumers run to completion  
(daemon not set to true)**

# Agenda

1. Interleaving execution

2. Producer/consumer

 3. Deadlock, starvation

# Synchronization can lead to two problems: deadlock and starvation

## Deadlock


- Objects lock resources
- Execution cannot proceed because object needs a resource another locked
- Object A locks resource 1
- Object B locks resource 2
- A needs resource 2 to proceed but B has it locked
- B needs resources 1 to proceed but A has it locked
- A and B are deadlocked

## Starvation

- Thread never gets resource it needs
- Thread A needs resource 1 to complete
- Other threads always take resource 1 before A can get it
- We say A is *starved*

# Dinning Philosophers explains deadlock and starvation

## Dining Philosophers




### Problem set up

- Five philosophers ( $P_0$ - $P_4$ ) sit at a table to eat spaghetti
- There are forks between each of them (five total forks)
- Each philosopher needs two forks to eat
- After acquiring two forks, philosopher eats, then puts both forks down
- Another philosopher can then pick up and use fork previously put down (gross!)

# Dinning Philosophers explains deadlock and starvation

## Dining Philosophers



### Naïve approach

- Each philosopher picks up fork on left
- Then picks up fork on right
- Deadlock occurs if all philosophers get left fork, none get right fork

# For deadlock to occur four conditions must be met

## Deadlock conditions

### 1. Mutual exclusion

- At least one resource class must have non-sharable access. That is:
  - Either one process is using a resource (and others wait), or
  - Resource is free

### 2. Hold and wait

- At least one process is holding a resource instance, while also waiting to be granted another resource instance. (e.g., Each philosopher is holding on to their left fork, while waiting to pick up their right fork)

### 3. No preemption

- Resources cannot be pre-empted; a resource can be released only voluntarily by the process holding it (e.g., can't force philosophers to drop their forks.)

### 4. Circular wait

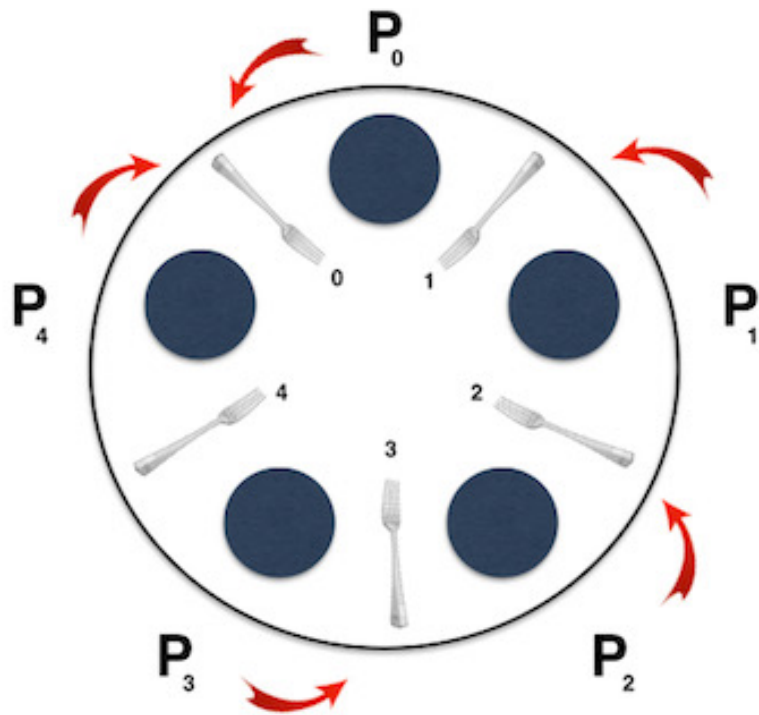
- There must exist a circular chain of at least two processes, each of whom is waiting for a resource held by another one. (e.g., each `Philosopher[i]` is waiting for `Philosopher[(i+1) mod 5]` to drop its fork.)

# Three ways to ensure deadlock does not occur

1. Ensure circular wait cannot occur by numbering Forks and reaching for smallest numbered Fork first
2. Prevent circular wait by making one of the philosophers wait until at least one other philosopher is finished
3. Prevent hold and wait by making Fork acquisition an atomic operation (e.g., must get both Forks in one step)

# We can break the deadlock by ensuring the “circular wait” does not occur

## Dinning Philosophers



Could also force one of the Philosophers to wait at first

## Eliminate circular wait

- Number each fork in circular fashion
- Make each philosopher pick up lowest numbered fork first
- All pick up right fork, except  $P_4$  who tries to pick up left fork 0
- Either  $P_0$  or  $P_4$  get fork 0
- If  $P_0$  gets it,  $P_4$  waits for fork 0 before picking up fork 4, so  $P_3$  eats
- $P_3$  eventually releases both forks and  $P_2$  eats
- Others eat after  $P_2$
- Cannot deadlock



# Fork.java models forks in the Dining Philosophers problem

## Fork.java

```
6 public class Fork {
7     private boolean available = true;
8
9     public synchronized void acquire() throws InterruptedException {
10         while (!available) {
11             wait();
12         }
13         available = false;
14     }
15
16     public synchronized void release() {
17         available = true;
18         notifyAll();
19     }
20 }
21
```

# Philosophers try to eat by getting both the left and right Forks

## Philosopher.java

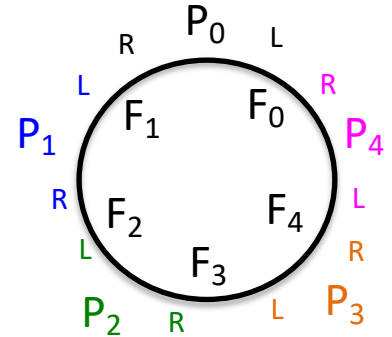
```
6 public class Philosopher extends Thread {
7     private int num;           // for message printout
8     private Fork left, right;  // the resources
9
10    public Philosopher(int num, Fork left, Fork right) {
11        this.num = num;
12        this.left = left;
13        this.right = right;
14    }
15
16    /**
17     * Waits a bit -- 1 to 5 seconds
18     */
19    private void randPause() throws InterruptedException {
20        sleep(1000 + (int)(Math.random()*4000));
21    }
22
23    /**
24     * Start the rounds of resource acquisition
25     */
26    public void run () {
27        for (int meal = 0; meal < 3; meal++) {
28            eat();
29            System.out.println(num + " finished meal " + meal);
30        }
31        System.out.println(num + " all done");
32    }
33
34    /**
35     * One round
36     */
37    public void eat() {
38        try {
39            System.out.println(num + " contemplating the universe, working up an appetite");
40            randPause();
41            System.out.println(num + " hungry; going for left fork");
42            left.acquire();
43            System.out.println(num + " got left fork");
44            randPause();
45            System.out.println(num + " going for right fork");
46            right.acquire();
47            System.out.println(num + " got right fork; chowing down");
48            randPause();
49            System.out.println(num + " finished eating; dropping forks");
50            right.release();
51            left.release();
52        }
53        catch (InterruptedException e) {
54            System.err.println(e);

```

# DiningPhilosophers.java uses five Philosophers and five Forks

## DiningPhilosopher.java

```
8 public class DiningPhilosophers {
9     private ArrayList<Philosopher> philosophers;
10
11     /**
12      * Creates the forks and philosophers
13      */
14     public DiningPhilosophers() {
15         ArrayList<Fork> forks = new ArrayList<Fork>();
16         for (int fork = 0; fork < 5; fork++) {
17             forks.add(new Fork());
18         }
19
20         philosophers = new ArrayList<Philosopher>();
21         for (int phil = 0; phil < 5; phil++) {
22             philosophers.add(new Philosopher(phil, forks.get(phil), forks.get((phil+1)%5)));
23         }
24     }
25
26     /**
27      * Gets each philosopher started at the table
28      */
29     public void dine() {
30         for (Philosopher phil : philosophers) {
31             phil.start();
32         }
33     }
34
35     public static void main(String[] args) {
36         new DiningPhilosophers().dine();
37     }
38 }
```




# DEMO: DiningPhilosophers.java

- Run several times
- Sometimes deadlocks
- Try adjusting pause time to longer to make it less likely to deadlock

# Another approach is to prevent “hold and wait” by picking up both forks atomically

## Dinning Philosophers



### Eliminate hold and wait

- Make picking up both forks an atomic operation
- Forks no longer control their destiny as in prior code
- Now we lock both with a mutex
- Could lead to starvation if one philosopher always picks up before another
- In this case starvation will eventually end because the philosophers only eat a limited number of meals

# Prevent deadlocks by making getting both Forks an atomic operation

## MonitoredDiningPhilosopher.java

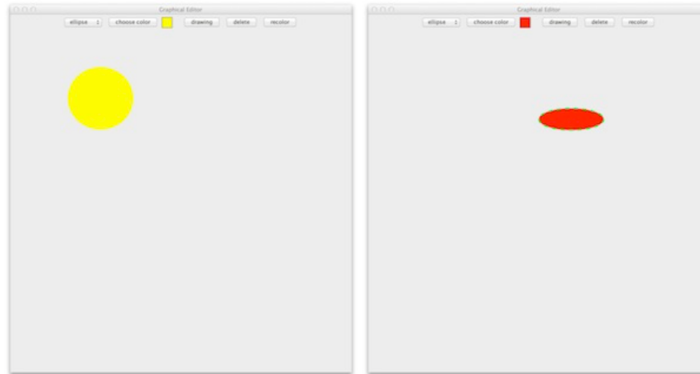
```
9 public class MonitoredDiningPhilosophers {
10     private ArrayList<MonitoredPhilosopher> philosophers;
11
12     /**
13      * Creates the forks and philosophers
14      */
15     public MonitoredDiningPhilosophers() {
16         ArrayList<MonitoredFork> forks = new ArrayList<MonitoredFork>();
17         for (int fork = 0; fork < 5; fork++) {
18             forks.add(new MonitoredFork());
19         }
20
21         philosophers = new ArrayList<MonitoredPhilosopher>();
22         for (int phil = 0; phil < 5; phil++) {
23             philosophers.add(new MonitoredPhilosopher(this, phil, forks.get(phil), forks.get((phil+1)%5)));
24         }
25     }
26
27     /**
28      * Gets each philosopher started at the table
29      */
30     public void dine() {
31         for (MonitoredPhilosopher phil : philosophers) {
32             phil.start();
33         }
34     }
35
36     /**
37      * Simultaneously acquires both resources
38      */
39     public synchronized void acquire(MonitoredFork left, MonitoredFork right) throws InterruptedException {
40         while (!left.available || !right.available) {
41             wait();
42         }
43         left.available = false;
44         right.available = false;
45     }
46
47     /**
48      * Releases both resources
49      */
50     public synchronized void release(MonitoredFork left, MonitoredFork right) {
51         left.available = true;
52         right.available = true;
53         notifyAll();
54     }
55
56     public static void main(String[] args) {
57         new MonitoredDiningPhilosophers().dine();
58     }
59 }
```

# SA-10

## Exercises

This short assignment will build up a core piece of a graphical editor that we'll expand upon in PS-6. The next problem set will flesh it out and make it concurrent (like a shared canvas, with multiple people drawing on it). Now we'll just get some basic user interface machinery in place.

This core part only supports drawing and modifying a single ellipse. When the "draw" radio button is selected, an ellipse is drawn by pressing the mouse button for one corner of the ellipse (well, the bounding box around it) and dragging to the other corner. When the "move" radio button is selected, the ellipse can be moved by clicking on it and dragging. It can likewise be deleted or recolored by first selecting the appropriate radio button and then clicking on the shape.



Most of the GUI is given in this scaffold: [EditorOne.java](#); there are places for you to plug in some code to make it all work. The ellipse itself is handled by a separate class, [Ellipse.java](#) implementing an interface, [Shape.java](#), which will have a number of other implementations in the problem set. The shape stuff is a bit wedded to the Java AWT Graphics machinery (i.e., combining the state and the presentation), in a manner analogous to AWT's own Geometry classes. But those classes have both more and less than we need here, and it's more fun to do our own anyway.

While the task is really just to translate comments to Java, do make sure you understand how it all fits together. I've put some print statements to help. Try printing the current ellipse too – I provided a convenient `toString`.

A few notes (largely echoed in the comments):

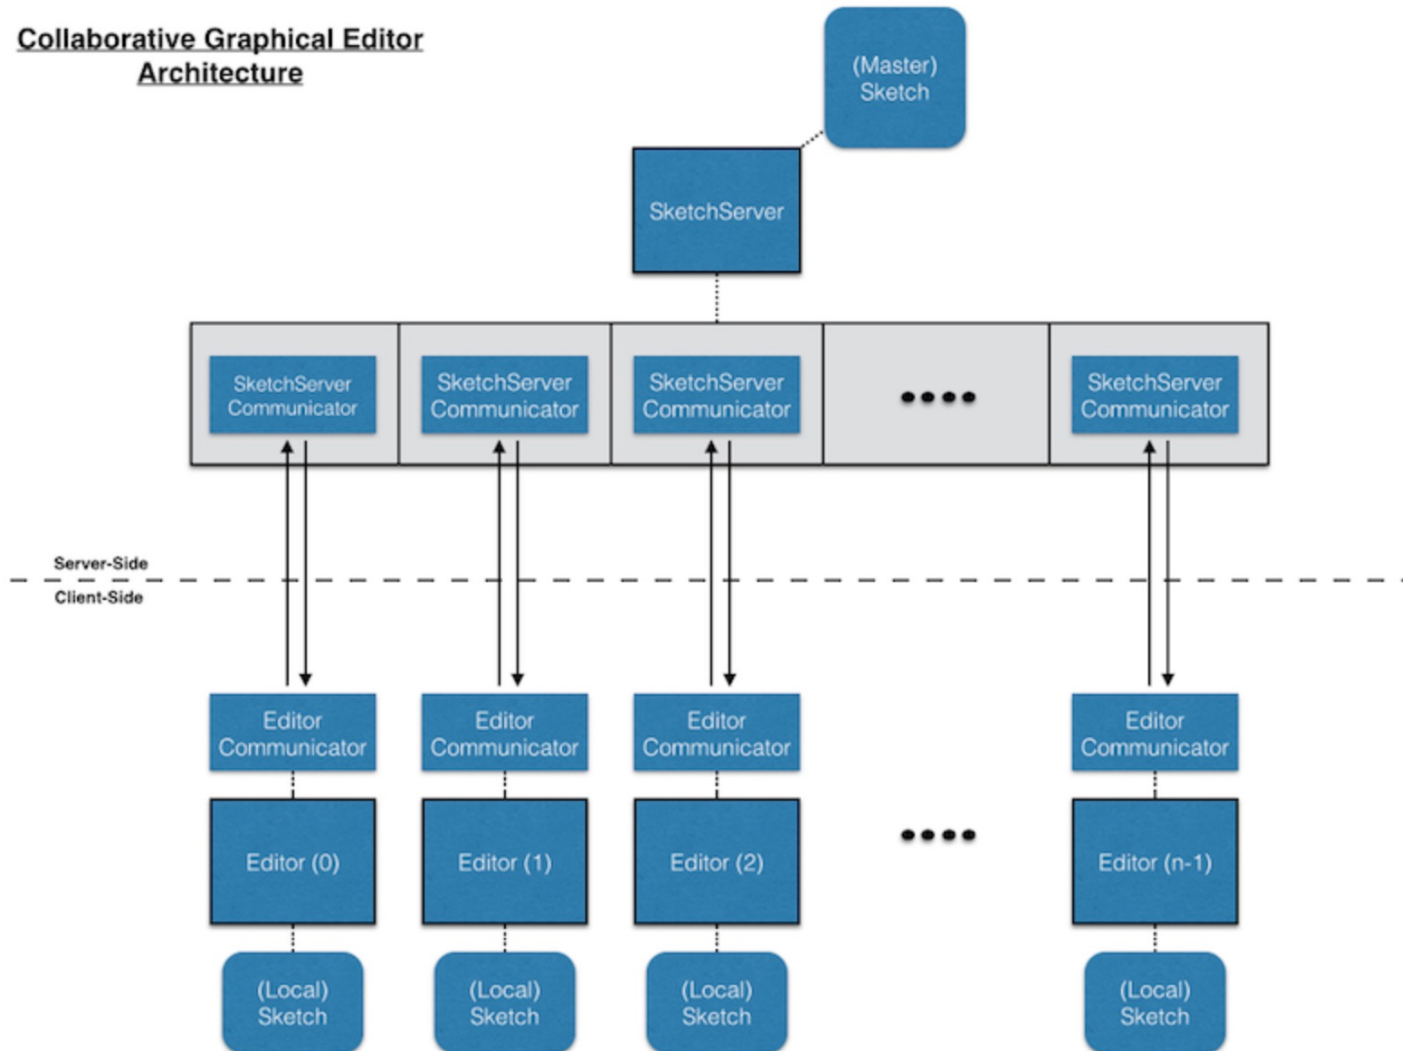
- The GUI elements and canvas largely follow the style of the Flickr search tool. The `JColorChooser` illustrates how much can be packaged up in a widget, with a callback to tell us what color was clicked on.
- The `mode` variable indicates which of the radio buttons has been selected. An `enum` is just a nice and safe way to have a bunch of related constant values. Various other actions (e.g., does a drag expand the ellipse or move it?) depend on the setting of this variable. For example, you can say `if (mode == Mode.DRAW) ...`.
- The `shape` variable holds either the one and only ellipse drawn, or null.
- The `drawFrom` variable should indicate where the mouse was first pressed to begin a new ellipse. The `moveFrom` variable likewise should indicate where it last was during dragging.
- Recall that the `repaint` method can be invoked to cause a refresh of the display after things have changed (recoloring, etc.) You'll need to sprinkle it around, but think about when you really want to do that.

## Submission Instructions

Turn in your completed Java code and a snapshot of a most beautiful ellipse.

# PS-6

## Collaborative Graphical Editor Architecture





# Summary

- Unexpected behavior when working with threads as they can be interrupted at any point
  - Use of synchronized to make the operation atomic
- Producers/consumers paradigm with notifying and waiting for synchronization
- Synchronization can lead to deadlock and starvation
  - Ensure to make atomic operations
  - Limit resource use


# Additional Resources

Concurrency and shared resources

# **ANNOTATED SLIDES**

# Threads are a way for multiple processes to run concurrently

## Threads



```
main() {  
MyThread t = new MyThread();  
  
//start thread at run method, main  
thread keeps running  
t.start()  
  
//halt main until thread finishes  
t.join()  
}
```

Assume `MyThread` is a class that extends `Thread`


`MyThread` must implement a `run` method

Execution begins by calling `start` on a `MyThread` object, `run` method then executes

Can call `join` to halt main program until thread finishes

# Concurrent threads can access the same resources; this can cause problems

## Concurrency



- Remember a static variable is a Class variable, there is only one
- Every Object of MyThread Class references the same static variable
- Threads can be interrupted at any time by the Operating System and another Thread may run
- When each Thread tries to increment `total`, it gets a current copy of `total`, adds 1, then stores it back in memory
- What can go wrong?

# Threads can be interrupted at any point, this can cause unexpected behavior

## Incrementer.java

```
7 public class Incrementer extends Thread {
8     private static int total = 0; // a variable shared by all incrementers
9     private static final int times = 1000000; // how many times to increment total, in each thread
10
11     /**
12      * Increments total the specified number of times
13      */
14     public void run() {
15         for (int i = 0; i < times; i++) {
16             total++;
17         }
18     }
19
20     public static void main(String [] args) throws Exception {
21         Incrementer inc1 = new Incrementer();
22         Incrementer inc2 = new Incrementer();
23
24         // Fire off threads and wait for them to complete
25         inc1.start();
26         inc2.start();
27         inc1.join();
28         inc2.join();
29
30         System.out.println("total at end = " + total);
31     }
32 }
```

***total* is static so it is a Class variable (one *total* for all Incrementer Objects)**

**Increment *total* one million times:**

- Get value of *total* from memory
- Add one to *total*
- Write *total* back to memory

**Two Incrementer Objects that extend Thread (so must implement *run()* method)**

- *start()* begins Thread running and calls *run()* method
- *main()* continues running after *inc1.start()*, so *inc2* starts immediately after *inc1* (*main()* does not block and wait for *inc1* to finish)

- *inc1.join()* causes *main()* to block until *inc1.run()* finishes
- *inc2.join()* causes *main()* to block until *inc2.run()* finishes

# Threads can be interrupted at any point, this can cause unexpected behavior

## Incrementer.java

```
7 public class Incrementer extends Thread {
8     private static int total = 0; // a variable shared by all incrementers
9     private static final int times = 1000000; // how many times to increment total, in each thread
10
11     /**
12      * Increments total the specified number of times
13      */
14     public void run() {
15         for (int i = 0; i < times; i++) {
16             total++;
17         }
18     }
19
20     public static void main(String [] args) throws Exception {
21         Incrementer inc1 = new Incrementer();
22         Incrementer inc2 = new Incrementer();
23
24         // Fire off threads and wait for them to complete
25         inc1.start();
26         inc2.start();
27         inc1.join();
28         inc2.join();
29
30         System.out.println("total at end = " + total);
31     }
32 }
```

### Increment *total* one million times:

- Get value of *total* from memory
- Add one to *total*
- Write *total* back to memory

### Operating System might interrupt a Thread at any point:

- *inc1* reads value of *total* from memory (say it's 10)
- *inc1* gets interrupted and *inc2* begins running
- *inc2* reads value of *total* (10), increments and writes back (*total*=11)

### Say *inc2* runs for 5 iterations (*total*=15)

- *inc2* interrupted and *inc1* resumes running
- *inc1* increments *total* to 11 and writes it back
- *total* now 11 not 16 as expected

# IncrementerInterleaving.java demonstrates interruptions (sometimes)

## IncrementerInterleaving.java

```
6 public class IncrementerInterleaving extends Thread {
7     private static int total = 0; // a variable shared by all incrementers
8     private static final int times = 5; // how many times to increment total, in each thread
9     private String name; // for display purposes
10
11     public IncrementerInterleaving(String name) {
12         this.name = name;
13     }
14
15     /**
16      * Increments total the specified number of times
17      */
18     public void run() {
19         for (int i = 0; i < times; i++) {
20             int temp = total;
21             System.out.println(name + " gets " + temp);
22             temp = temp + 1;
23             total = temp;
24             System.out.println(name + " puts " + temp);
25         }
26     }
27
28     public static void main(String [] args) throws Exception {
29         IncrementerInterleaving inc1 = new IncrementerInterleaving("one");
30         IncrementerInterleaving inc2 = new IncrementerInterleaving("two");
31
32         // Fire off threads and wait for them to complete
33         inc1.start();
34         inc2.start();
35         inc1.join();
36         inc2.join();
37
38         System.out.println("total at end = " + total);
39     }
40 }
41
```

**total** static as before  
Will loop 5 times in *run()* method  
Each Thread gets a name for clarity

- Printing to console is *sloowwww*
- Gives more time for OS to interrupt
- Console output shows when read and write **total**
- Might expect total to be 10 (5 from inc1 and 5 from inc2)
- Sometimes total is 10
- Most of the time it is not
- Bugs caused by multiple threads can be devilishly tricky to find



# Java provides the keyword `synchronized` to make some operations “atomic”

## IncrementerTotal.java

```
public class IncrementerTotal {  
    private int total = 0;  
    public synchronized void inc() {  
        total++;  
    }  
}
```

- **IncrementerTotal Class keeps a *total* instance variable**
- **Value of *total* incremented via *inc()* method**
- ***inc()* method is synchronized so only one Thread at a time can be inside *inc()***
- **IncrementerTotal Class used on next slide**

- `synchronized` keyword in front of `inc` method means only one thread can be running this code at a time
- If multiple threads try to run `synchronized` code, one thread runs, all others block until first one finishes
- Once first thread finishes, OS selects another thread to run
- `synchronized` makes this code “atomic” (e.g., as if it were one instruction)
- This synchronized approach is called a “mutex” (or monitor), acts like a “lock” on static `total` variable

# IncrementerSync.java uses atomic operations to ensure desired behavior

## IncrementerSync.java

```
8 public class IncrementerSync extends Thread {
9     private static IncrementerTotal total = new IncrementerTotal(); // a variable shared by all incrementers
10    private static final int times = 1000000; // how many times to increment total, in each thread
11
12    /**
13     * Increments total the specified number of times
14     */
15    public void run() {
16        for (int i = 0; i < times; i++) {
17            total.inc();
18        }
19    }
20
21    public static void main(String [] args) throws Exception {
22        IncrementerSync inc1 = new IncrementerSync();
23        IncrementerSync inc2 = new IncrementerSync();
24
25        // Fire off threads and wait for them to complete
26        inc1.start();
27        inc2.start();
28        inc1.join();
29        inc2.join();
30
31        System.out.println("total at end = " + total.total);
32    }
33 }
```

**total now an IncrementerTotal Object**  
**total.inc() is synchronized**

- **Synchronized total.inc() ensures only one Thread inside inc() at a time**
- **inc() runs to completion before another Thread allowed in**

```
public class IncrementerTotal {
    private int total = 0;
    public synchronized void inc() {
        total++;
    }
}
```

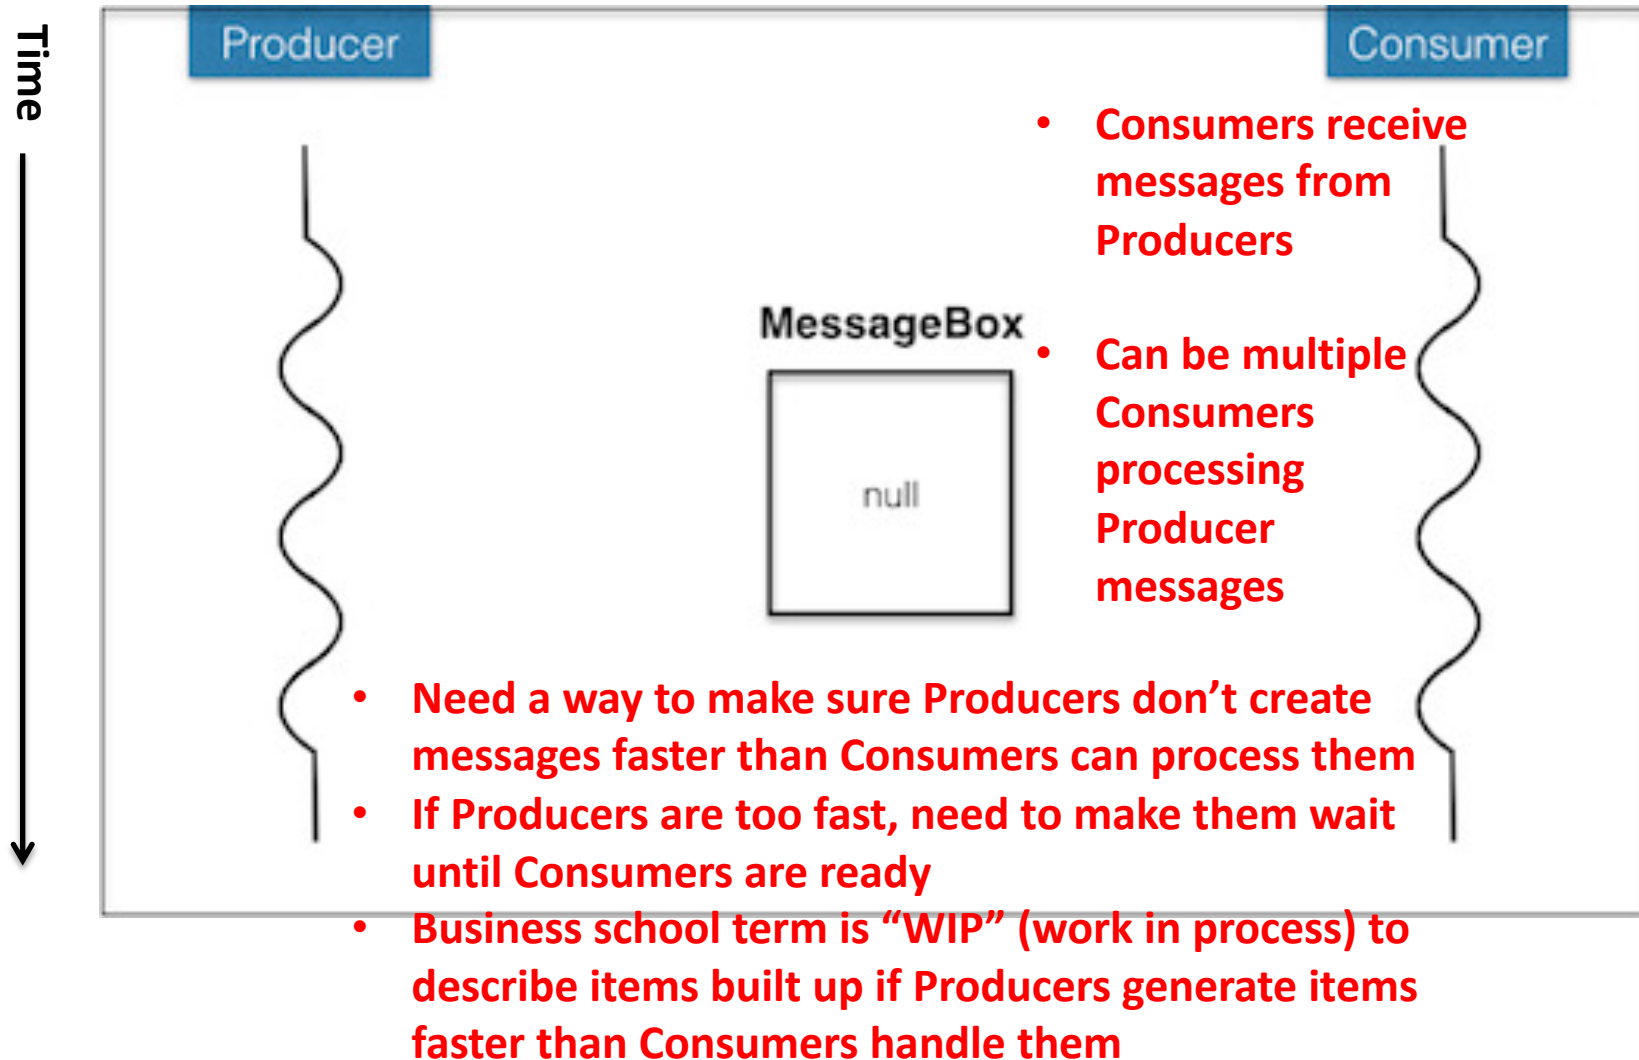
**total.total now always 2 million**

Producers/consumers

# **ANNOTATED SLIDES**

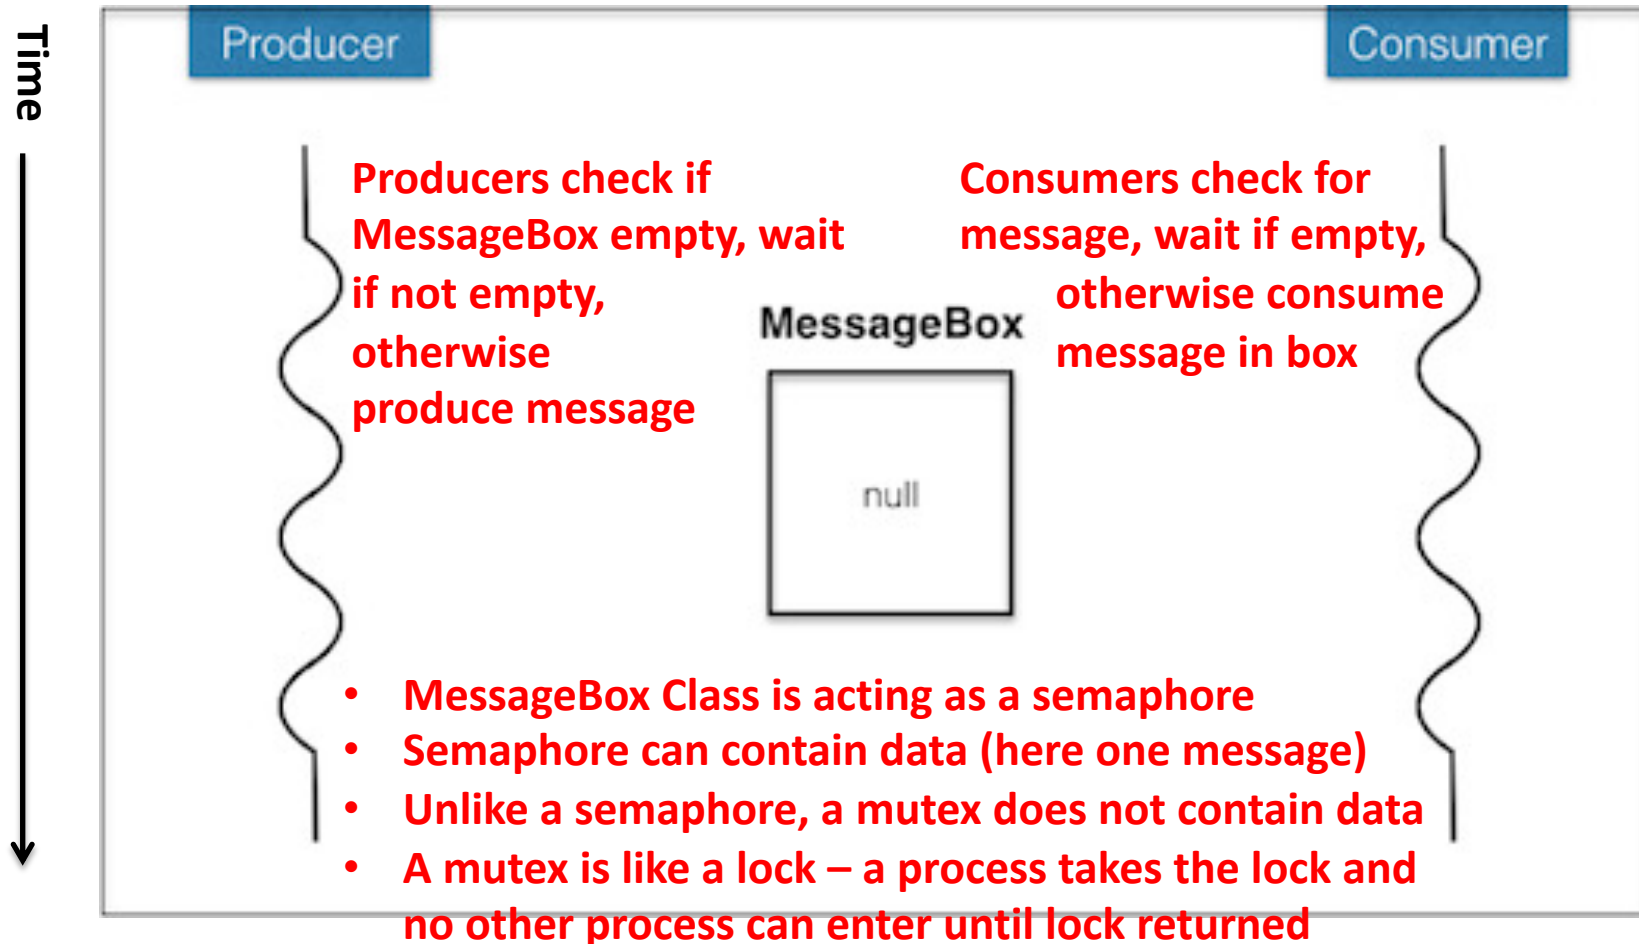
# Scenario: Producers produce messages for Consumers, need to keep in sync

## Example



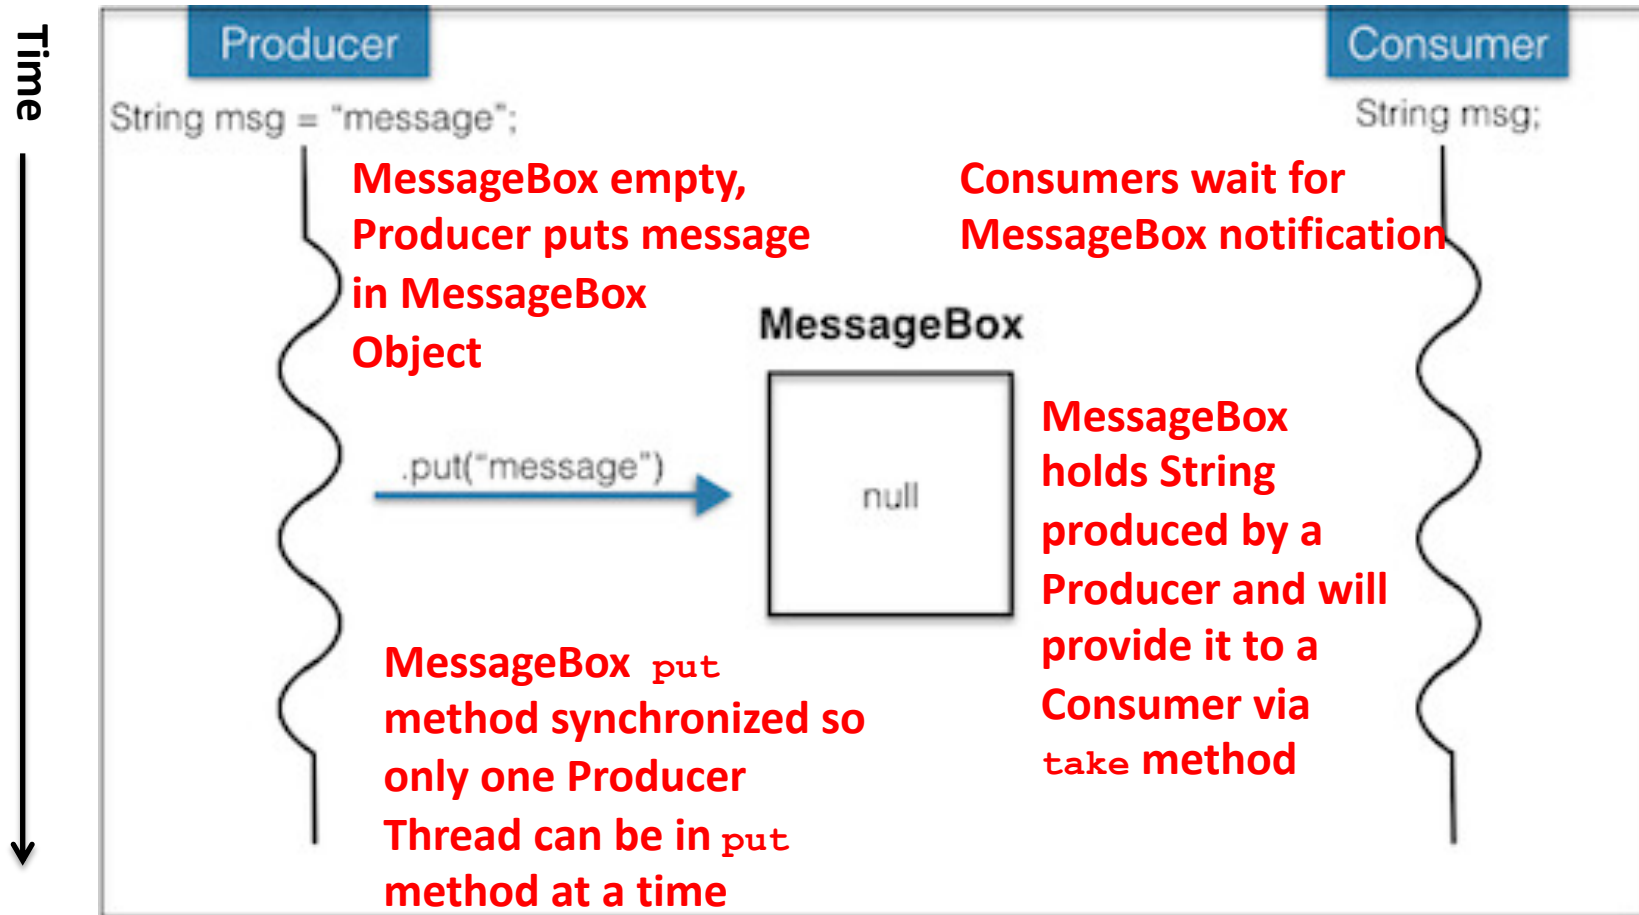
# We can use a semaphore to keep Producers and Consumers in sync

## Example



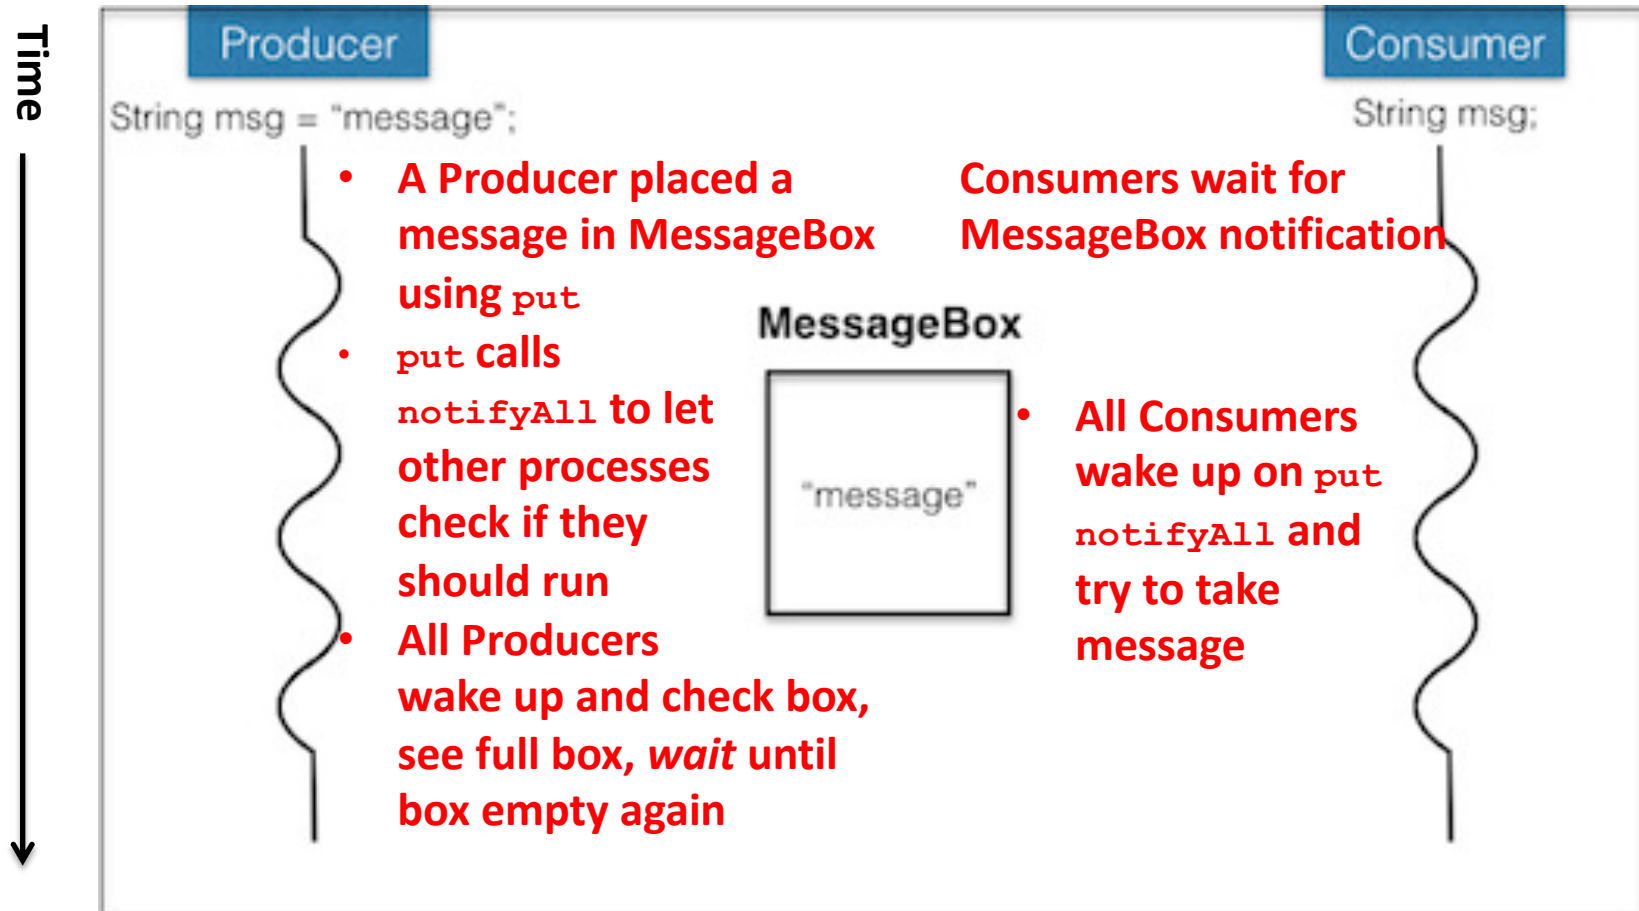
# Producer passing messages to Consumer using semaphore

## Example



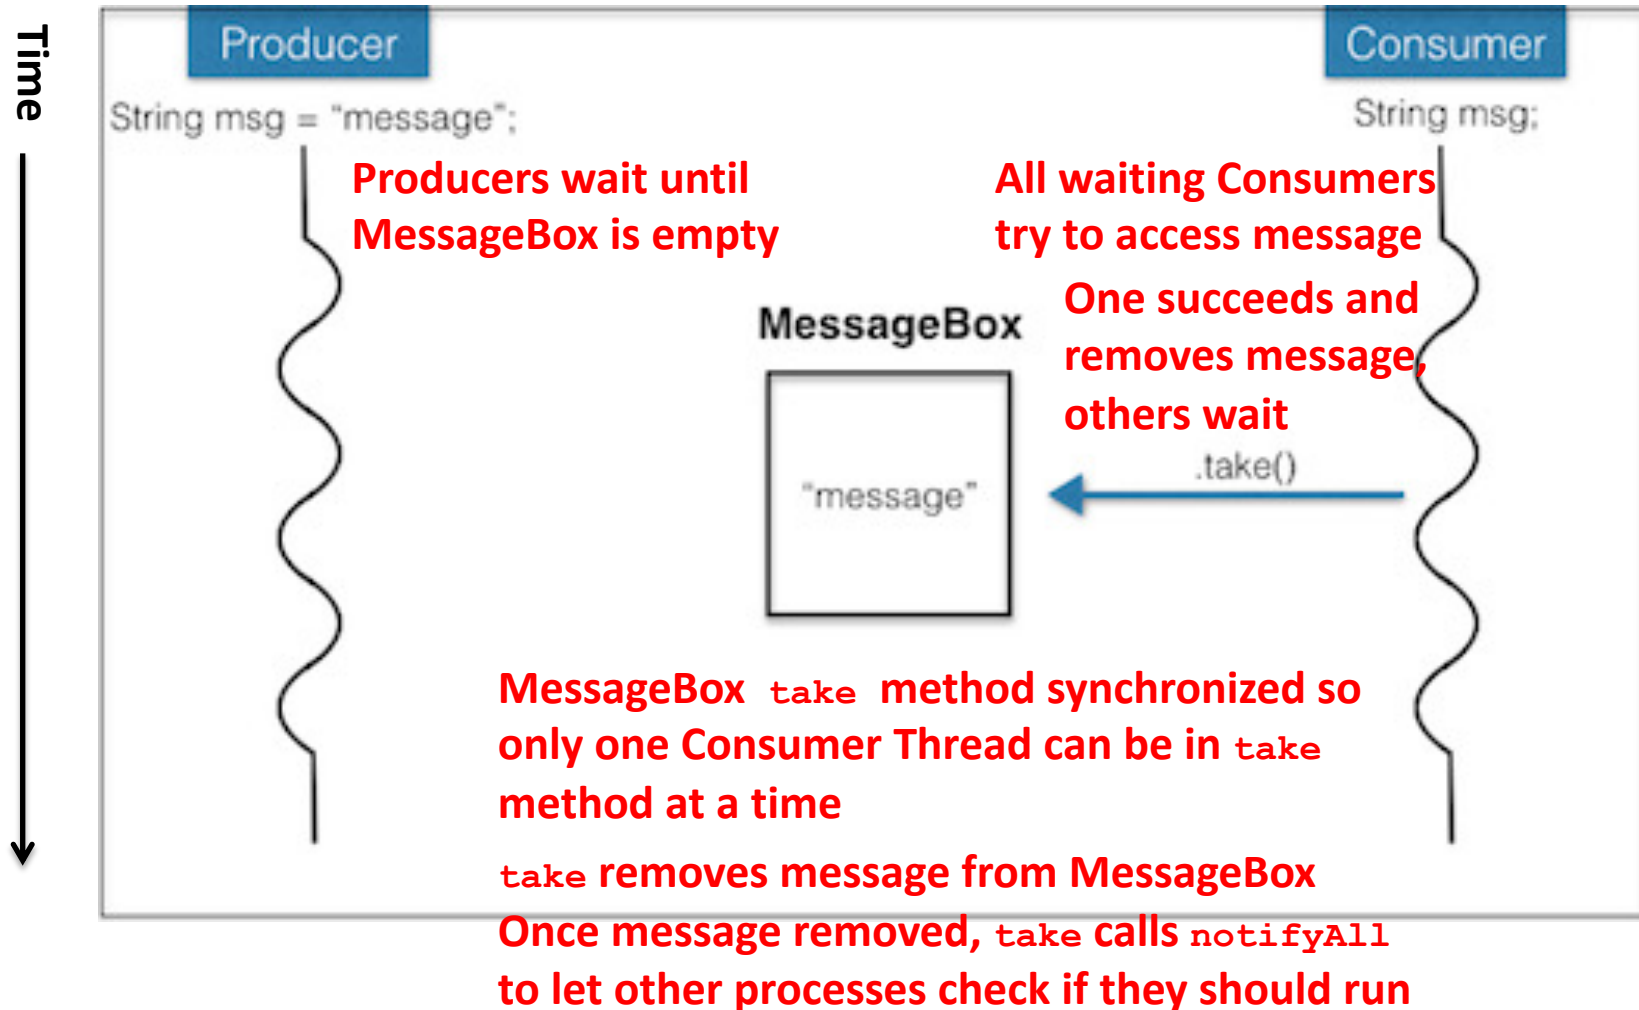
# Producer passing messages to Consumer using semaphore

## Example



# Producer passing messages to Consumer using semaphore

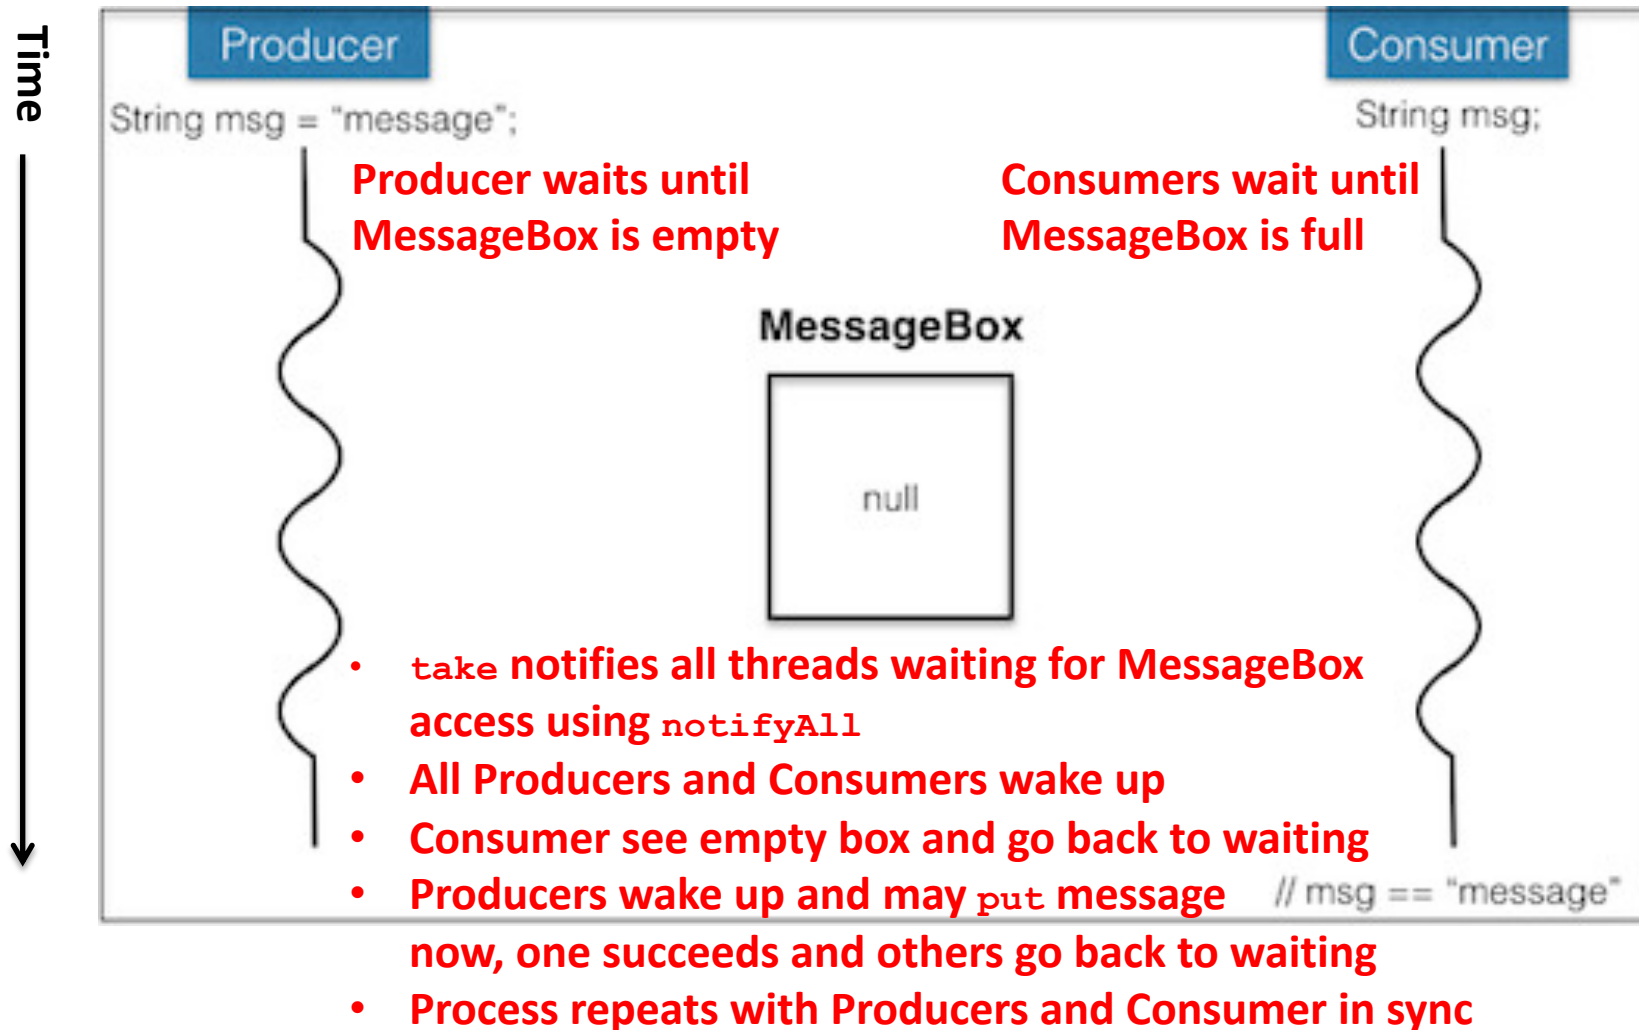
## Example






# Producer passing messages to Consumer using semaphore

## Example



# MessageBox.java implements a semaphore that holds one String



## MessageBox.java

```
7 public class MessageBox {
8     private String message = null;
9     /**
10     * Put m as message once it's okay to do so (current message has been taken)
11     */
12     public synchronized void put(String m) throws InterruptedException {
13         //check to see if message is not null, might have been woken by put() notifyAll
14         while (message != null) {
15             wait();
16         }
17         message = m;
18         notifyAll(); //wakes producers AND consumers
19     }
20     /**
21     * Takes message once it's there, leaving empty message
22     */
23     public synchronized String take() throws InterruptedException {
24         //check to see if message is null, might have been woken by take() notifyAll
25         while (message == null) {
26             wait();
27         }
28         String m = message;
29         message = null;
30         notifyAll(); //wakes producers AND consumers
31         return m;
32     }
33 }
34 }
35 }
```

**MessageBox holds one String called *message***

**Producers will fill *message* using *put()* method**

**Consumer will process message using *take()* method**


**Synchronized *put()* makes sure only one Producer at a time can store *message***

**↑ Notify all Threads (Producers and Consumers) to check MessageBox**

- Wait until MessageBox is empty
- If woken up (resume running at *wait*), make sure to check if MessageBox is empty

- It could be the case that many Producers were woken up and another Producer already filled the MessageBox
- An if statement wouldn't suffice, need a while to go back to sleep if box filled

# MessageBox.java implements a semaphore that holds one String



## MessageBox.java

```
7 public class MessageBox {
8     private String message = null;
9
10    /**
11     * Put m as message once it's okay to do so (current message has been taken)
12     */
13    public synchronized void put(String m) throws InterruptedException {
14        //check to see if message is not null, might have been woken by put() notifyAll
15        while (message != null) {
16            wait();
17        }
18        message = m;
19        notifyAll(); //wakes producers AND consumers
20    }
21
22    /**
23     * Takes message once it's there, leaving empty message
24     */
25    public synchronized String take() throws InterruptedException {
26        //check to see if message is null, might have been woken by take() notifyAll
27        while (message == null) {
28            wait();
29        }
30        String m = message;
31        message = null;
32        notifyAll(); //wakes producers AND consumers
33        return m;
34    }
35 }
```

**Synchronized ensures only one Consumer can take *message***

**If woken up, check *message*:**

- If empty, go back to waiting (another Consumer already took it)
- If not, return *message* and set to null


**MessageBox now empty, notify all Threads to wake up and check MessageBox**

# Producers use MessageBox to pass messages to Consumers

## Producer.java

```
6 public class Producer extends Thread {
7     private MessageBox box;
8     private int numberToSend;
9
10    public Producer(MessageBox box, int numberToSend) {
11        this.box = box;
12        this.numberToSend = numberToSend;
13    }
14
15    /**
16     * Wait for a while then puts a message
17     * Puts "EOF" when # messages have been put
18     */
19    public void run() {
20        try {
21            for (int i = 0; i < numberToSend; i++) {
22                sleep((int)(Math.random()*5000)); //sleep for random time up to 5 seconds
23                box.put("message #" + i); //put a new message in MessageBox
24            }
25            box.put("EOF"); //EOF means end of file
26        }
27        catch (InterruptedException e) {
28            System.err.println(e);
29        }
30    }
31 }
```

**MessageBox as parameter**  
**If multiple Producers, all would get the *same* MessageBox**



- When Thread starts, try to put a *message* in the MessageBox using *put()* after random interval
- *put()* will cause this Producer to *wait()* if there is already a message
- That will remove this Thread from *put()* and add it to a list of Threads waiting to run

**Send EOF when all messages sent**


- When *notifyAll()* received, this Thread will wake up and resume running in *put()* method of MessageBox
- If MessageBox is empty it will store its message and return here

# Consumers retrieve messages from the MessageBox

## Consumer.java


```
6 public class Consumer extends Thread {
7     private MessageBox box;
8
9     public Consumer(MessageBox box) {
10        this.box = box;
11    }
12
13    /**
14     * Takes messages from the box and prints them, until receiving EOF
15     */
16    public void run() {
17        try {
18            String message;
19            while (!(message = box.take()).equals("EOF")) {
20                System.out.println(message);
21            }
22        }
23        catch (InterruptedException e) {
24            System.err.println(e);
25        }
26    }
27 }
```

Store same MessageBox that Producers use



Take *message* from MessageBox  
If no message, *take()* will cause this Thread to wait  
If this Thread retrieves message, check for EOF and exit

# ProducerConsumer uses all three components to pass messages



Create a MessageBox, a Producer, and a Consumer

Pass the same MessageBox Object to both the Producer and the Consumer (here 1 producer and 1 consumer)

Producer *run()* will wait a random period, then put a *message* in MessageBox, then wait until MessageBox empty  
Consumer will wake up on *notifyAll()* from MessageBox and *take()* message

*take()* issues *notifyAll()* after taking *message*, waking Producer to *put()* next *message*

*main()* thread will complete after starting both Producer and Consumer Objects

After creating ProducerConsumer Object, call *communicate()*

*main()* ends, but Producers and Consumers run to completion (daemon not set to true)

## ProducerConsumer.java

```
8 public class ProducerConsumer {
9     public static final int numMessages = 5; // how many messages to send from produc
10    private Producer producer;
11    private Consumer consumer;
12
13    public ProducerConsumer() {
14        MessageBox box = new MessageBox();
15        producer = new Producer(box, numMessages);
16        consumer = new Consumer(box);
17    }
18
19    /**
20     * Just starts the producer and consumer running
21     */
22    public void communicate() {
23        producer.start();
24        consumer.start();
25    }
26
27    public static void main(String[] args) {
28        new ProducerConsumer().communicate();
29        System.out.println("Peace out! (threads are still running but I'm done)");
30    }
31 }
```


```
<terminated> ProducerConsumer [Java Application] /Library/Java/JavaVirtualMachines/jdk1.8.0_112.jdk/Contents/Home/bin/java (Feb 22, 2018, 11:55:46 AM)
Peace out! (threads are still running but I'm done)
message #0
message #1
message #2
message #3
message #4
```

Deadlock and starvation

# **ANNOTATED SLIDES**

# Dinning Philosophers explains deadlock and starvation

## Dining Philosophers



### Problem set up

- Five philosophers ( $P_0$ - $P_4$ ) sit at a table to eat spaghetti
- There are forks between each of them (five total forks)
- Each philosopher needs two forks to eat
- After acquiring two forks, philosopher eats, then puts both forks down
- Another philosopher can then pick up and use fork previously put down (gross!)



# Fork.java models forks in the Dining Philosophers problem

## Fork.java

```
6 public class Fork {
7     private boolean available = true;
8
9     public synchronized void acquire() throws InterruptedException {
10         while (!available) {
11             wait();
12         }
13         available = false;
14     }
15
16     public synchronized void release() {
17         available = true;
18         notifyAll();
19     }
20 }
```

***available*** tracks if this Fork Object is being used

**Synchronized *acquire()*** causes wait if Fork is not available  
If acquire Fork, set ***available*** false

- ***release()*** makes Fork available to others
- Use ***notifyAll()*** to tell Philosophers a Fork is free

# Philosophers try to eat by getting both the left and right Forks

## Philosopher.java

```
6 public class Philosopher extends Thread {
7     private int num;           // for message printout
8     private Fork left, right;  // the resources
9
10    public Philosopher(int num, Fork left, Fork right) {
11        this.num = num;
12        this.left = left;
13        this.right = right;
14    }
15
16    /**
17     * Waits a bit -- 1 to 5 seconds
18     */
19    private void randPause() throws InterruptedException {
20        sleep(1000 + (int)(Math.random()*4000));
21    }
22
23    /**
24     * Start the rounds of resource acquisition
25     */
26    public void run () {
27        for (int meal = 0; meal < 3; meal++) {
28            eat();
29            System.out.println(num + " finished meal " + meal);
30        }
31        System.out.println(num + " all done");
32    }
33
34    /**
35     * One round
36     */
37    public void eat() {
38        try {
39            System.out.println(num + " contemplating the universe, working up an appetite");
40            randPause();
41            System.out.println(num + " hungry; going for left fork");
42            left.acquire();
43            System.out.println(num + " got left fork");
44            randPause();
45            System.out.println(num + " going for right fork");
46            right.acquire();
47            System.out.println(num + " got right fork; chowing down");
48            randPause();
49            System.out.println(num + " finished eating; dropping forks");
50            right.release();
51            left.release();
52        }
53        catch (InterruptedException e) {
54            System.err.println(e);
55        }
56    }
57 }
```

Philosopher runs on a Thread and is passed left and right Fork (also passed a philosopher number)

Philosophers try to eat three meals

- *eat()* tries to *acquire()* the left and right fork (after universe contemplation of course)
- Always tries to get Fork on left first (could be a problem if Forks not numbered properly)
- *acquire()* will cause a wait if Fork not available
- Once philosopher has both Forks, he can eat
- Philosopher releases both Forks after eating

# DiningPhilosophers.java uses five Philosophers and five Forks

## DiningPhilosopher.java

```
8 public class DiningPhilosophers {
9     private ArrayList<Philosopher> philosophers;
10
11     /**
12      * Creates the forks and philosophers
13      */
14     public DiningPhilosophers() {
15         ArrayList<Fork> forks = new ArrayList<Fork>();
16         for (int fork = 0; fork < 5; fork++) {
17             forks.add(new Fork());
18         }
19
20         philosophers = new ArrayList<Philosopher>();
21         for (int phil = 0; phil < 5; phil++) {
22             philosophers.add(new Philosopher(phil, forks.get(phil), forks.get((phil+1)%5)));
23         }
24     }
25
26     /**
27      * Gets each philosopher started at the table
28      */
29     public void dine() {
30         for (Philosopher phil : philosophers) {
31             phil.start();
32         }
33     }
34
35     public static void main(String[] args) {
36         new DiningPhilosophers().dine();
37     }
38 }
```

Will hold multiple Philosophers in ArrayList

Set up five Fork Objects in ArrayList

Create five Philosophers and pass the left and right Fork Objects


$P_0$  left =  $F_0$ , right =  $F_1$

$P_4$  left =  $F_4$ , right =  $F_0$

Could deadlock!

Reverse Forks for  $P_4$  and won't deadlock

Start each Philosopher dining (calls *run()* on previous slide)



# Prevent deadlocks by making getting both Forks an atomic operation

## MonitoredDiningPhilosopher.java

```
9 public class MonitoredDiningPhilosophers {
10     private ArrayList<MonitoredPhilosopher> philosophers;
11
12     /**
13      * Creates the forks and philosophers
14      */
15     public MonitoredDiningPhilosophers() {
16         ArrayList<MonitoredFork> forks = new ArrayList<MonitoredFork>();
17         for (int fork = 0; fork < 5; fork++) {
18             forks.add(new MonitoredFork());
19         }
20
21         philosophers = new ArrayList<MonitoredPhilosopher>();
22         for (int phil = 0; phil < 5; phil++) {
23             philosophers.add(new MonitoredPhilosopher(this, phil, forks.get(phil), forks.get((phil+1)%5)));
24         }
25     }
26
27     /**
28      * Gets each philosopher started at the table
29      */
30     public void dine() {
31         for (MonitoredPhilosopher phil : philosophers) {
32             phil.start();
33         }
34     }
35
36     /**
37      * Simultaneously acquires both resources
38      */
39     public synchronized void acquire(MonitoredFork left, MonitoredFork right) throws InterruptedException {
40         while (!left.available || !right.available) {
41             wait();
42         }
43         left.available = false;
44         right.available = false;
45     }
46
47     /**
48      * Releases both resources
49      */
50     public synchronized void release(MonitoredFork left, MonitoredFork right) {
51         left.available = true;
52         right.available = true;
53         notifyAll();
54     }
55
56     public static void main(String[] args) {
57         new MonitoredDiningPhilosophers().dine();
58     }
59 }
```

- Move *acquire()* and *release()* to main program, not controlled by individual Forks now
- Synchronized only allows one Philosopher in *acquire()* at a time, wait if left and right Forks not available
- Pick up both Forks while here
- *release()* also synchronized
- Drop both Forks while here
- *notifyAll()* when Forks are available