



Prelim 2

Slight difference in two prelims; we'll curve as we did last time. Wait until all regrade requests have been answered.

Open up for regrade requests around noon.

No lunch meeting with me today. Instead, I am in my office to talk to people from 11:30 to 3:00.

Please don't use Piazza for questions about prelim. Talk to a consultant/TA/instructor, email instructors.

Concurrent Programs

A *thread* or *thread of execution* is a sequential stream of computational work.

Concurrency is about controlling access by multiple threads to shared resources.

Last time: Learned about

1. Race conditions
2. Deadlock
3. How to create a thread in Java.

Purpose of this lecture

Show you Java constructs for eliminating race conditions, allowing threads to access a data structure in a safe way but allowing as much concurrency as possible.

This requires

- (1) The locking of an object so that others cannot access it, called **synchronization**.
- (2) Use of two new Java methods: **wait()** and **notifyAll()**

As an example, throughout, we use a **bounded buffer**.

Look at JavaHyperText, entry Thread !!!!!!!

An Example: bounded buffer

finite capacity (e.g. 20 loaves)
implemented as a queue

Threads A: **produce** loaves of bread and put them in the queue

Threads B: **consume** loaves by taking them off the queue

An Example: bounded buffer

finite capacity (e.g. 20 loaves)

Separation of concerns:

1. How do you implement a queue in an array?
2. How do you implement a bounded buffer using the queue, which allows producers to add to it and consumers to take things from it, all in parallel?

Threads A: **produce** loaves of bread and put them in the queue

Threads B: **consume** loaves by taking them off the queue

ArrayQueue

Array b[0..5]

0 1 2 3 4 5 b.length
b

5	3	6	2	4	
---	---	---	---	---	--

put values 5 3 6 2 4 into queue

ArrayQueue

Array b[0..5]

0 1 2 3 4 5 b.length
b

5	3	6	2	4	
---	---	---	---	---	--

put values 5 3 6 2 4 into queue

get, get, get

ArrayQueue

Array b[0..5]

0 1 2 3 4 5 b.length
b

3	5		2	4	1
---	---	--	---	---	---

 Values wrap around!!

put values 5 3 6 2 4 into queue

get, get, get

put values 1 3 5

ArrayQueue

h
0 1 2 3 4 5 b.length
b



3	5		2	4	1
---	---	--	---	---	---

 Values wrap around!!



```
int[] b; // 0 <= h < b.length. The queue contains the  
int h; // n elements b[h], b[h+1], b[h+2], ...  
int n; // b[h+n-1] (all indices mod b.length)
```

```
/** Pre: there is space */      /** Pre: not empty */  
public void put(int v){        public int get(){  
    b[(h+n) % b.length]= v;    int v= b[h];  
    n= n+1;                    h= (h+1) % b.length;  
                                n= n-1;  
                                return v;  
}
```

Bounded Buffer

```
/** An instance maintains a bounded buffer of fixed size */  
class BoundedBuffer<E> {  
    ArrayQueue<E> aq;    
    /** Put v into the bounded buffer.*/  
    public void produce(E v) {  
        if (!aq.isFull()) { aq.put(v) };  
    }  
  
    /** Consume v from the bounded buffer.*/  
    public E consume() {  
        return aq.isEmpty() ? null : aq.get();  
    }  
}
```

Bounded Buffer

```
/** An instance maintains a bounded buffer of fixed size */  
class BoundedBuffer<E> {  
    ArrayQueue<E> aq;    
    /** Put v into the bounded buffer.*/  
    public void produce(E v) {  
        if (!aq.isFull()) { aq.put(v) };  
    }  
  
    Problems  
    1. Chef doesn't easily know whether bread was added.  
    2. Suppose race condition  
        (a) First chef finds it not full.  
        (b) Another chef butts in and adds a bread  
        (c) First chef tries to add and can't because  
            it's full. Need a way to prevent this
```

Bounded Buffer

13

```

/** An instance maintains a bounded buffer of fixed size */
class BoundedBuffer<E> {
    Queue<E> aq;

    /** Put v into the bounded buffer.*/
    public void produce(E v) {
        if (!aq.isFull()) { aq.put(v) };
    }
}
    
```

After finding `aq` not full, but before putting `v`, another chef might beat you to it and fill up buffer `aq`! **race condition**



Use of synchronized

14

synchronized (object) {
code

The object is the outhouse.
The code is the person,
waiting to get into the object.
If the key is on the door, the
code takes it, goes in, locks
the door, executes, opens the
door, comes out, and hangs
the key up.

Key is hanging the outhouse.

Anyone can grab the key, go
inside, and lock the door. They
have the key.

When they come out, they lock
the door and hang the key by
the front door. Anyone (only
one) person can then grab the
key, go inside, lock the door.

That's what **synchronized**
implements!

Synchronized block

15

```

/** An instance maintains a bounded buffer of fixed size */
class BoundedBuffer<E> {
    Queue<E> aq;

    /** Put v into the bounded buffer.*/
    public void produce(E v) {
        synchronized (aq) {
            if (!aq.isFull()) { aq.put(v) };
        }
    }
}
    
```

Synchronized blocks

```

public void produce(E v) {
    synchronized(this) {
        if (!aq.isFull()) { aq.put(v);
    }
}
    
```

You can synchronize (lock) any object, including this.

```

BB@10
BB@10 BB
aq_____
produce() {...} consume() {...}
    
```

Synchronized Methods

```

public void produce(E v) {
    synchronized(this) {
        if (!aq.isFull()) { aq.put(v); }
    }
}
    
```

You can synchronize (lock) any object, including **this**.

```

public synchronized void produce(E v) {
    if (!aq.isFull()) { aq.put(v); }
}
    
```

Or you can synchronize methods

This is the same as wrapping the entire method implementation
in a `synchronized(this)` block

Bounded buffer

18

```

/** An instance maintains a bounded buffer of fixed size */
class BoundedBuffer<E> {
    Queue<E> aq;

    /** Put v into the bounded buffer.*/
    public synchronized void produce(E v) {
        if (!aq.isFull()) { aq.put(v); }
    }
}
    
```

What happens if `aq` is full?

We want to wait until it becomes non-full —until there
is a place to put `v`.
Somebody has to buy a loaf of bread before we can put
more bread on the shelf.

Two lists for a synchronized object

For every synchronized object *obj*, Java maintains:

- locklist:** a list of threads that are waiting to obtain the lock on *obj*
- waitlist:** a list of threads that had the lock but executed `wait()`
 - e.g. because they couldn't proceed

Method `wait()` is defined in `Object`



Wait()

```
class BoundedBuffer<E> {
    ArrayQueue<E> aq;
    /** Put v into the bounded buffer.*/
    public synchronized void produce(E v) {
        while (aq.isFull()) {
            try { wait(); }
            catch (InterruptedException e) {}
        }
        aq.put(v);
        notifyAll();
    }
    ...
}
```

need while loop (not if statement) to prevent race conditions

puts thread on the wait list

threads can be interrupted if this happens just continue.



notify() and notifyAll()

- Methods `notify()` and `notifyAll()` are defined in `Object`
- `notifyAll()` moves all threads on the waitlist to the locklist
- `notify()` moves one thread from the waitlist to the locklist
 - Note: which thread is moved is arbitrary



notify() and notifyAll()

```
/** An instance maintains a bounded buffer of fixed size */
class BoundedBuffer<E> {
    ArrayQueue<E> aq;
    /** Put v into the bounded buffer.*/
    public synchronized void produce(E v) {
        while (aq.isFull()) {
            try { wait(); }
            catch (InterruptedException e){}
        }
        aq.put(v);
        notifyAll();
    }
    ...
}
```

WHY use of notify() may hang.

Two sets:

Work with a bounded buffer of length 1.

- Consumer W gets lock, wants White bread, finds buffer empty, and `wait()`s: is put in set 2.
 - Consumer R gets lock, wants Rye bread, finds buffer empty, `wait()`s: is put in set 2.
 - Producer gets lock, puts Rye in the buffer, does `notify()`, gives up lock.
 - The `notify()` causes one waiting thread to be moved from set 2 to set 1. Choose W.
 - No one has lock, so one Runnable thread, W, is given lock. W wants white, not rye, so `wait()`s: is put in set 2.
 - Producer gets lock, finds buffer full, `wait()`s: is put in set 2.
- All 3 threads are waiting in set 2. **Nothing more happens.**

1. lock: threads waiting to get lock.

2. wait: threads waiting to be notified

Should one use notify() or notifyAll()

But suppose there are two kinds of bread on the shelf—and one still picks the head of the queue, if it's the right kind of bread.



Using `notify()` can lead to a situation in which no one can make progress.

notifyAll() always works; you need to write documentation if you optimize by using notify()

Eclipse Example

25

Producer: produce random ints (rye or white bread)

Consumer 1: even ints (buy only rye bread)

Consumer 2: odd ints (buy only white bread)

Dropbox: 1-element bounded buffer (1 loaf of bread)

Locklist

Threads
wanting the
Dropbox



Waitlist

Threads who
had Dropbox
and waited



Dropbox: empty
or 1 integer (loaf of bread)



Word of warning with synchronized

26



BUT: You leave the back door open and tell your friends to go in whenever they want

Threads that don't synchronize can get in. Dangerous but useful to increase efficiency.

Key is hanging by front door.

Anyone can grab the key, go inside, and lock the door. They have the key.

When they come out, they lock the door and hang the key by the front door. Anyone (only one) person can then grab the key, go inside, lock the door.

That's what synchronized implements!

Using Concurrent Collections...

27

Java has a bunch of classes to make synchronization easier.

It has synchronized versions of some of the Collections classes

It has an Atomic counter.

From spec for HashSet

28

... this implementation is not synchronized. If multiple threads access a hash set concurrently, and at least one of the threads modifies the set, it must be synchronized externally. This is typically accomplished by synchronizing on some object that naturally encapsulates the set. If no such object exists, the set should be "wrapped" using method `Collections.synchronizedSet`. This is best done at creation time, to prevent accidental unsynchronized access to the set:

```
Set s = Collections.synchronizedSet(new HashSet(...));
```

Race Conditions

29

Thread 1

Thread 2

Initially, i = 0

tmp = load i; Load 0 from memory

Load 0 from memory tmp = load i;

tmp = tmp + 1; Store 1 to memory
store tmp to i;

Store 1 to memory tmp = tmp + 1;
store tmp to i;

time ↓

Finally, i = 1

Using Concurrent Collections...

30

```
import java.util.concurrent.atomic.*;

public class Counter {
    private static AtomicInteger counter;

    public Counter() {
        counter = new AtomicInteger(0);
    }

    public static int getCount() {
        return counter.getAndIncrement();
    }
}
```

Fancier forms of locking

Java. **synchronized** is the core mechanism

But. Java has a class Semaphore. It can be used to allow a limited number of threads (or kinds of threads) to work at the same time. Acquire the semaphore, release the semaphore

Semaphore: a kind of synchronized counter (invented by Dijkstra in 1962-63, THE multiprogramming system)

The Windows and Linux and Apple O/S have kernel locking features, like file locking

Python: acquire a **lock**, release the **lock**. Has semaphores

Summary

Use of multiple processes and multiple threads within each process can exploit concurrency

- may be real (multicore) or virtual (an illusion)

Be careful when using threads:

- synchronize shared memory to avoid race conditions
- avoid deadlock

Even with proper locking concurrent programs can have other problems such as "livelock"

Serious treatment of concurrency is a complex topic (covered in more detail in cs3410 and cs4410)

Nice tutorial at

<http://docs.oracle.com/javase/tutorial/essential/concurrency/index.html>