


THREADS AND THREAD SYNCHRONIZATION


Instructor: Prasun Dewan (FB 150, dewan@unc.edu)




DISTRIBUTION VS. CONCURRENCY PROGRAM



Distribution, no fine-grained concurrency




Concurrency, not distribution



Distribution and fine-grained concurrency (typical)



DEMO: HALLOWEEN SIMULATION



WAITING SINGLE-THREAD SOLUTION

```
loop
```

```
    wait for local user input
```

```
    process user input
```

```
    wait for remote user input
```

```
    process remote user input
```

```
end
```

Should not impose order on user input

Not a chess game!



POLLING SINGLE-THREAD SOLUTION

```
loop
    if (local user input received)
        process user input
    if (remote user input received)
        process remote user input
        sleep (interval)
    end
```

Busy waiting

Wastes computer resources

Reduces response time based on
sleep time



MULTIPLE THREADS

```
loop  
    wait for local user input  
    process user input  
end
```

```
loop  
    wait for remote user input  
    process remote user input  
end
```

Much cleaner code

One reason threads were added to
the OS Kernel



THREADS AND SYNCYRONIZATION

Programming: Abstraction use

Systems: Abstraction design and implementation

Theory: Models and algorithms

Some repetition for those who have seen
threads and synchronization before




TOPICS

- Threads
- Command objects
- Bounded Buffer
- Generics
- Interrupts
- Critical sections
- Busy waiting
- Semaphores
- Monitors
 - Entry procedures
 - Conditions
 - Hints and Absolutes
 - Java and General Hints
- Path Expressions




PROGRAM VS. PROCESS VS. THREAD



Program

```
public class AP2PAPAliceTOTSimulation {  
    public static String ID = "9100";  
    public static String NAME = "Alice";  
    public static int USER_NUMBER = 0;  
    public static void main (String[] args) {  
        Tracer.showInfo(true);  
        AP2PTOTSessionsClientCreator.createP2PDelayed.  
    }  
}
```

Execution instance



Process is execution instance of program, associated with program and memory

Thread is an independent activity, within a process, associated with a process and a stack




THREAD AS AN ACTIVE AGENT WITH DATA STRUCTURES

The screenshot shows a Java development environment with the following components:

- Left Panel:** Shows the project structure with files: ABoundedBuffer.java, AProducer.java, and ABoundedBufferMain.java.
- Middle Panel:** Displays the Java Virtual Machine (JVM) monitor. It lists two suspended threads:
 - Thread [Thread-0] (Suspended (breakpoint at line 16 in ABoundedBufferMain))
 - Thread [Thread-1] (Suspended (breakpoint at line 16 in ABoundedBufferMain))Each thread has a stack trace showing calls to ABoundedBuffer.put(ElementType), AProducer.run(), and Thread.run().
- Right Panel:** Shows the variable values for the current context. A table lists:

Name	Value
this	ABoundedBuffer<ElementType>
buffer	Object[10] (id=30)
[0]	"Hello" (id=29)
[1]	null
[2]	null
[3]	null
- Bottom Right:** An "Outline" view showing class members: nextIn, nextOut, put(ElementType), and netIn.


THREAD AS A DATA STRUCTURE TO IMPLEMENTATION



To the thread implementation thread is a data structure

Thread implementation can be operating system, programming language, library

THREAD AS AN INDEPENDENT ACTIVITY




Whether current, ready, non ready captured by thread status

What triggers context switching?




RESCHEDULING



Context switching occurs as a result of time quantum expiring or some other higher priority thread becoming ready

With multiple CPUs, multiple current threads, and maybe even multiple ready queues

THREAD AS A DATA STRUCTURE



Where does a thread start executing

WHAT DOES A THREAD EXECUTE?

Executes some identifiable portion of a program

Executes a method call asynchronously

Thread creator does not wait for method to terminate

PROCEDURE CALL VS. THREAD CREATION

```
buffer.put ("hello");
System.out.println("Put complete");
```


```
Thread putThread = async buffer.put ("hello");
System.out.println("Put started");
```

Not real Java code, Java syntax used to illustrate design choices

Why thread object?

What thread operations?

THREAD AS AN OBJECT



A thread is an object representing an independent execution of code and can be started., suspended, resumed, interrupted while sleeping, given lower/higher priority ...

PROCEDURE CALL VS. THREAD CREATION

```
Thread putThread = async buffer.put ("hello");  
System.out.println("Put started");
```

Without language syntax?

PROCEDURE CALL VS. THREAD CREATION

```
Thread putThread = (new Thread(buffer, put, "ca va")).start();  
System.out.println("Put started");
```

Object
address

Method
address

Method
parameters

Thread constructor must take varying number of arguments

Varying number of parameters to
constructor /method not allowed
in type safe languages

Method parameters not allowed in
most OO languages

COMMAND OBJECT

Command Object
=Embedded Operation +
Parameters

execute ()

constructor (targetObject,
param1, param2, ...);

Provides an execute operation
defined by a well-defined interface
to execute some procedure

Constructor takes target object
and parameters of operation as
arguments

The execute operation takes no
arguments

A command represents a
procedure call



JAVA RUNNABLE COMMAND OBJECT

Runnable Implementation

run()

constructor (targetObject,
param1, param2, ...);

```
package java.lang;  
public interface Runnable {  
    public void run();  
}
```

Command object defined by the Runnable interface



EXAMPLE RUNNABLE IMPLEMENTATION

```
public class AProducer implements Runnable {  
    BoundedBuffer<String> boundedBuffer;  
    String element;  
    public AProducer(  
        BoundedBuffer<String> aBoundedBuffer,  
        String anElement) {  
        boundedBuffer = aBoundedBuffer;  
        element = anElement;  
    }  
    @Override  
    public void run() {  
        boundedBuffer.put(element);  
    }  
}
```

The diagram illustrates the flow of type parameters in the code. It shows two light blue boxes connected by arrows. One arrow points from the 'constructor (targetObject, param1, param2, ...)' box to the 'run()' box. Another arrow points from the 'run()' box to the 'Type parameter assigned as in List' box, which contains the text 'Type parameter assigned as in List'. The 'run()' box also has a direct connection to the 'Type parameter assigned as in List' box.

run()

constructor (targetObject, param1, param2, ...);

Type parameter assigned as in List




THREAD CREATION

```
public class ABoundedBufferMain {  
    public static void main(String[] args) {  
        BoundedBuffer<String> greetings =  
            new ABoundedBuffer();  
        Runnable producer1 =  
            new AProducer<String>(greetings, "Hello") ;  
        Runnable producer2 =  
            new AProducer<String>(greetings, "Ca Va") ;  
        (new Thread(producer1)).start();  
        (new Thread(producer2)).start();  
    }  
}
```



DELEGATION VS. INHERITANCE BASED THREAD/COMMAND OBJECT



GENERIC ELABORATION

```
public class ABoundedBufferMain {  
    public static void main(String[] args) {  
        BoundedBuffer<String> greetings =  
            new ABoundedBuffer();  
        Runnable producer1 =  
            new AProducer<String>(greetings, "Hello");  
        Runnable producer2 =  
            new AProducer<String>(greetings, "Ca Va");  
        (new Thread(producer1)).start();  
        (new Thread(producer2)).start();  
    }  
}
```

BoundedBuffer elaborated

How defined?



GENERIC BOUNDED BUFFER INTERFACE

```
public interface BoundedBuffer<ElementType> {  
    void put(ElementType element);  
    ElementType get();  
}
```

Type parameter declared

Generic BoundedBuffer Interface



JAVA GENERIC TYPES (TYPE PARAMETERS)

A scope name (class, interface, method) can be succeeded by a series of type parameters within angle brackets <A, B, C, ...> that have the same value in all places

```
public interface I<T> {  
    public void addElement (T t);  
    public T elementAt (int index);  
    public int size();  
}
```

Create an **elaboration** of generic by giving actual value to type parameter

```
I<String> stringI;  
I<Point> pointI;
```

A method, class, interface with type parameter is called a **generic**.

A single implementation is shared by all of its elaborations

Assigning values to type parameters is a compile time activity for type checking only



GENERIC BOUNDED BUFFER CLASS

```
public class ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    // data structures
    ...
    void put(ElementType element) {
        ...
    }
    ElementType get() {
        ...
    }
}
```

Can we give them different names?

Multiple occurrences of a type variable assigned (unified to) the same value

Why repeat type parameter in class and interface?

If a class is elaborated with String then so is the interface



GENERIC BOUNDED BUFFER CLASS

```
public class ABoundedBuffer
    implements BoundedBuffer<String>{
    // data structures
    ...
    void put(String element) {
        ...
    }
    ElementType get() {
        ...
    }
}
```

Unparameterized class elaborating parameter of interface

May want non generic class implementing generic interface



RENAMING AND MULTIPLE TYPE VARIABLES

```
public interface I1 <T>{  
    ...  
}
```

```
public interface I2 <T>{  
    ...  
}
```

```
public class C<T1, T2> implements I1 <T1>, I2<T2>{  
    ...  
}
```

Parameterized class may
implement multiple parameterized
interfaces

GENERIC BOUNDED BUFFER INTERFACE

```
public interface BoundedBuffer<ElementType> {  
    void put(ElementType element);  
    ElementType get();  
}
```

Implementation?



A BOUNDED BUFFER (DATA STRUCTURES)

```
public class ABounddBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int size = 0;
    int nextIn = 0;
    int nextOut = 0;
```

Convention: Class A<T> implements interface <T>
containing public methods of A



BOUNDED BUFFER (METHODS)

```
public void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
}  
  
public ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
}  
}  
}  
} Requires polling to wait for  
non empty and empty slots
```

Thread safe?

Will it work if multiple
threads access it



EXAMPLE THREADS

```
public class ABoundedBufferMain {  
    public static void main(String[] args) {  
        BoundedBuffer<String> greetings =  
            new ABoundedBuffer();  
        Runnable producer1 =  
            new AProducer<String>(greetings, "Hello") ;  
        Runnable producer2 =  
            new AProducer<String>(greetings, "Ca Va") ;  
        (new Thread(producer1)).start() ;  
        (new Thread(producer2)).start() ;  
    }  
}
```

“Hello” and “Ca Va” threads created



“HELLO” THREAD

Screenshot of a Java debugger interface showing the state of threads and memory.

Threads:

- A Bounded Buffer Main [Java Application]
- synchronization.ABoundedBufferMain at localhost:55905
- Thread [Thread-0] (Suspended (breakpoint at line 13 in ABoundedBuffer))
 - ABoundedBuffer<ElementType>.put(ElementType) line: 13
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available
- Thread [DestroyJavaVM] (Running)
- Thread [Thread-1] (Suspended (breakpoint at line 13 in ABoundedBuffer))
 - ABoundedBuffer<ElementType>.put(ElementType) line: 13
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available
- C:\Program Files\Java\jre6\bin\javaw.exe (Aug 26, 2011 9:07:52 AM)

Memory View:

Name	Value
this	ABoundedBuffer<ElementType> (id=20)
buffer	Object[4] (id=32)
[0]	null
[1]	null
[2]	null
[3]	null
nextIn	0
nextOut	0
size	0
element	"Hello" (id=28)

Code View:

```
int nextOut = 0;

public void put(ElementType element) {
    if (size >= MAX_SIZE)
        return;
    buffer[nextIn] = element;
    nextIn++;
    size++;
}
```

Annotations:

- “Hello” Thread Scheduled
- “Hello” Thread Starts Executing put()



“HELLO” THREAD PROCEEDS

```
J A BoundedBufferMain [Java Application]
  synchronization.ABoundedBufferMain at localhost:55905
    Thread [Thread-0] (Suspended (breakpoint at line 16 in ABoundedBuffer))
      ABoundedBuffer<ElementType>.put(ElementType) line: 16
      AProducer<ElementType>.run() line: 12
      Thread.run() line: not available
    Thread [DestroyJavaVM] (Running)
    Thread [Thread-1] (Suspended (breakpoint at line 13 in ABoundedBuffer))
      ABoundedBuffer<ElementType>.put(ElementType) line: 13
      AProducer<ElementType>.run() line: 12
      Thread.run() line: not available
C:\Program Files\Java\jre6\bin\javaw.exe (Aug 26, 2011 9:07:52 AM)
```

Name	Value
this	ABoundedBuffer<ElementType> (id=20)
buffer	Object[4] (id=32)
[0]	"Hello" (id=28)
[1]	null
[2]	null
[3]	null
nextIn	0
nextOut	0
size	0
element	"Hello" (id=28)

Hello

```
APopulatedQuestGorge  ObjectAdapter.java  AnAvatarQueue.java  AMovableGeneralizedQ  ABoundedBuffer.java  ABoun
int nextOut = 0;

public void put(ElementType element) {
    if (size >= MAX_SIZE)
        return;
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
}
```

“Hello” Thread Adds Item

nextIn not incremented before rescheduling happens



“CA VA” THREAD

```
J A BoundedBufferMain [Java Application]
  synchronization.ABoundedBufferMain at localhost:55905
    Thread [Thread-0] (Suspended (breakpoint at line 13 in ABoundedBuffer))
      ABoundedBuffer<ElementType>.put(ElementType) line: 13
      AProducer<ElementType>.run() line: 12
      Thread.run() line: not available
    Thread [DestroyJavaVM] (Running)
    Thread [Thread-1] (Suspended (breakpoint at line 13 in ABoundedBuffer))
      ABoundedBuffer<ElementType>.put(ElementType) line: 13
      AProducer<ElementType>.run() line: 12
      Thread.run() line: not available
C:\Program Files\Java\jre6\bin\javaw.exe (Aug 26, 2011 9:07:52 AM)
```

Name	Value
>this	ABoundedBuffer<ElementType> (id=20)
buffer	Object[4] (id=32)
[0]	null
[1]	null
[2]	null
[3]	null
nextIn	0
nextOut	0
size	0
element	"Ca Va" (id=34)
Ca Va	

```
int nextOut = 0;

public void put(ElementType element) {
    if (size >= MAX_SIZE)
        return;
    buffer[nextIn] = element;
    nextIn++;
    size++
```

“Ca Va” Thread Scheduled



CONTEXT SWITCH TO “CA VA” THREAD

A Bounded Buffer Main [Java Application]

- synchronization.ABoundedBufferMain at localhost:55905
- Thread [Thread-0] (Suspended (breakpoint at line 16 in ABoundedBuffer))
 - ABoundedBuffer<ElementType>.put(ElementType) line: 16
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available
- Thread [DestroyJavaVM] (Running)
- Thread [Thread-1] (Suspended (breakpoint at line 16 in ABoundedBuffer))
 - ABoundedBuffer<ElementType>.put(ElementType) line: 16
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available

C:\Program Files\Java\jre6\bin\javaw.exe (Aug 26, 2011 9:07:52 AM)

Name	Value
this	ABoundedBuffer<ElementType> (id=20)
buffer	Object[4] (id=32)
[0]	"Ca Va" (id=34)
[1]	null
[2]	null
[3]	null
nextIn	0
nextOut	0
size	0
element	"Ca Va" (id=34)

```
int nextOut = 0;

public void put(ElementType element) {
    if (size >= MAX_SIZE)
        return;
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
}
```

Previous greeting overwritten

Non deterministic program

Program behavior depends on when rescheduling occurs



CRITICAL SECTION



Critical Section

Code that can be executed only if certain conditions are met by other threads

Other threads should not be executing the same critical section

Other threads should not be executing other critical sections that access the same data



CRITICAL SECTION (REVIEW)



Critical Section

Code that can be executed only if certain conditions are met by other threads

Other threads should not be executing the same critical section

Other threads should not be executing other critical sections that access the same data



CRITICAL SECTION

```
lock.lock()
```

Critical Section

```
lock.unlock()
```

```
public interface Lock {  
    void lock();  
    void unlock();  
}
```

CRITICAL SECTION (REVIEW)

```
lock.lock()
```

Critical Section

```
lock.unlock()
```

```
public interface Lock {  
    void lock();  
    void unlock();  
}
```



USE OF LOCK

```
public class ALockingBoundedBuffer<ElementType>
    extends ABoundedBuffer<ElementType>{
    Lock lock;
    public ALockingBoundedBuffer(Lock aLock) {
        lock = aLock;
    }
    public void put(ElementType element) {
        lock.lock();
        super.put(element);
        lock.unlock();
    }
    public ElementType get() {
        lock.lock();
        ElementType retVal = super.get();
        lock.unlock();
        return retVal;
    }
}
```

How to implement lock?



DISABLE CONTEXT SWITCHING

```
public class AContextSwitchingLock implements Lock{  
    int originalInterrupts;  
    public void lock() {  
        originalInterrupts = System.disableInterrupts();  
    }  
    public void unlock() {  
        System.restoreInterrupts(originalInterrupts);  
    }  
}
```

Context switching occurs as a result of a blocking call or an interrupt

Timer interrupt can cause time quantum to expire

Input/disk/network interrupt can cause higher priority thread to becoming eligible

Disable interrupts and make no blocking call in critical section



PROS AND CONS

Coarse grained: stops all threads, even those that do not access same or any critical section

Insecure: a thread can cheat and hog processor time


Simple

Needed to implement critical sections in operating system

Need critical sections to implement critical sections!



FINE-GRAINED CRITICAL SECTION



BOOLEAN OBJECT CAPTURING LOCK STATUS

```
public class ABooleanObject implements BooleanObject{  
    boolean value;  
    public boolean get() {  
        return value;  
    }  
    public void set(boolean newVal) {  
        value = newVal;  
    }  
}
```

boolean is not an object

Boolean is not an mutable



BUSY WAITING LOCK

```
public class APollingLock implements Lock{  
    BooleanObject booleanObject = new ABooleanObject();  
    public void lock() {  
        while (booleanObject.get()) {  
            ; // do nothing  
        }  
        booleanObject.set(true);  
    }  
    public void unlock() {  
        booleanObject.set(false);  
    }  
}
```

Test of lock

Set of lock

Does this work?

Test and set not done atomically

Context switch can occur between
test and set



THREAD 1 GETS LOCK

ABusyWaitingBoundedBufferMain [Java Application]

- synchronization.ABusyWaitingBoundedBufferMain at localhost:59103
 - Thread [Thread-0] (Suspended (breakpoint at line 8 in APollingLock))
 - APollingLock.lock() line: 8
 - ABusyWaitingBoundedBuffer<ElementType>.put(ElementType) line: 8
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available
 - Thread [DestroyJavaVM] (Running)
 - Thread [Thread-1] (Suspended (breakpoint at line 8 in APollingLock))
 - APollingLock.lock() line: 8
 - ABusyWaitingBoundedBuffer<ElementType>.put(ElementType) line: 8
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available

C:\Program Files\Java\jre6\bin\javaw.exe (Aug 27, 2011 10:10:27 AM)

Name	Value
this	APollingLock (id=28)
booleanObject	ABooleanObject (id=32)
value	false

ABoundedBuffer.java AboundedBufferMain.j APollingLock.java TestAndSetHardware.j Foo.java ATestAn

```
package synchronization;
public class APollingLock implements Lock{
    BooleanObject booleanObject = new ABooleanObject();
    public void lock() {
        while (booleanObject.get()) {
            ;
        }
        booleanObject.set(true);
    }
}
```

Thread 1 gets lock and context switches



THREAD 2 ALSO GETS LOCK

The screenshot shows a Java application running in a debugger. The left pane displays the stack traces for two suspended threads:

- Thread [Thread-0] (Suspended (breakpoint at line 8 in APollingLock))
 - APollingLock.lock() line: 8
 - ABusyWaitingBoundedBuffer<ElementType>.put(ElementType) line: 8
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available
- Thread [DestroyJavaVM] (Running)
- Thread [Thread-1] (Suspended (breakpoint at line 8 in APollingLock))
 - APollingLock.lock() line: 8
 - ABusyWaitingBoundedBuffer<ElementType>.put(ElementType) line: 8
 - AProducer<ElementType>.run() line: 12
 - Thread.run() line: not available

The right pane shows the local variable table for the current thread:

Name	Value
this	APollingLock (id=28)
booleanObject	ABooleanObject (id=32)
value	false

Below the debugger, the code editor shows the APollingLock.java file:

```
package synchronization;
public class APollingLock implements Lock{
    BooleanObject booleanObject = new ABooleanObject();
    public void lock() {
        while (booleanObject.get()) {
            ;
        }
        booleanObject.set(true);
    }
}
```

A blue callout box points to the line `booleanObject.set(true);` with the text "Thread 2 gets lock".

BUSY WAITING LOCK

```
public class APollingLock implements Lock{  
    BooleanObject booleanObject = new ABooleanObject();  
  
    public void lock() {  
        while (booleanObject.get()) {  
            ; // busy waiting  
        }  
        booleanObject.set(true);  
    }  
  
    public void unlock() {  
        booleanObject.set(false);  
    }  
}
```

Test of lock

Set of lock

Test and set not done atomically

Fix how?

Context switch can occur between
test and set

TEST AND SET HARDWARE INSTRUCTION

```
public class TestAndSetHardware {  
    public static atomic boolean  
        testAndSet(BooleanObject aBooleanObject) {  
            boolean retVal = aBooleanObject.get(); // test  
            aBooleanObject.set(true); // set  
            return retVal; // return tested value  
        }  
}
```

Function with side effects!

At the end of instruction, boolean object (memory slot) is always true

returns false if object was false at start of instruction

A TEST AND SET LOCK

```
public class ATestAndSetLock implements Lock{  
    BooleanObject booleanObject;  
    public void lock() {  
        while (TestAndSetHardware.testAndSet(booleanObject)) {  
            ;  
        }  
    }  
    public void unlock() {  
        booleanObject.set(false);  
    }  
}
```

Also called spin locks, as tester is spinning

BUSY WAITING PROS AND CONS

Busy waiting: polling wastes computer resources

Higher priority polling thread would not allow lower priority
thread in critical section to release lock

Needed to implement critical sections in multiprocessor operating
system

Disabling interrupts on one processor does not disable on others

Fine-grained approach with out disadvantages?

SEMAPHORE

```
public class ASemaphore implements Semaphore{  
    Queue<Thread> queue = new AQueue<Thread>();  
    int count;  
    public ASemaphore(int initialCount) {  
        count = initialCount;  
    }  
    public void semWait() {  
        int originalInterrupts = System.disableInterrupts();  
        count--;  
        if (count < 0) {  
            Thread currentThread = Thread.currentThread();  
            currentThread.setStatus(Thread.WAIT);  
            queue.add(currentThread);  
        }  
        System.restoreInterrupts(originalInterrupts);  
    }  
}
```

Count and queue manipulation must be done atomically

Disable and restore interrupts to ensure atomicity



SEMAPHORE (CONTD)

```
public void semSignal() {  
    int originalInterrupts = System.disableInterrupts();  
    count++;  
    if (count <= 0) {  
        Thread nextThread = queue.remove();  
        nextThread.setStatus(Thread.READY);  
        Thread.reschedAndRestoreInterrupts(  
            originalInterrupts);  
    } else {  
        System.restoreInterrupts(originalInterrupts);  
    }  
}
```

If count is positive, indicates how many semWait() can be done without blocking

If count is negative, its absolute value indicates queue size

SEMAPHORE-BASED LOCK

```
public class ASemaphoreLock implements Lock {  
    // let first thread go  
    Semaphore semaphore = new ASemaphore(1);  
    public void lock() {  
        semaphore.semWait();  
    }  
    public void unlock() {  
        semaphore.semSignal();  
    }  
}
```

SEMAPHORE PROS AND CONS

Fine-grained locking that works and is efficient

Like other schemes discussed so far, requires explicit lock and unlock

Programmer overhead in creating and using lock

Programmer may forget to lock

Worse, may forget to unlock

Program reader must look at code to find critical sections

MANUAL USE OF LOCK

```
public class ALockingBoundedBuffer<ElementType>
    extends ABoundedBuffer<ElementType>{
    Lock lock;
    public ALockingBoundedBuffer(Lock aLock) {
        lock = aLock;
    }
    public void put(ElementType element) {
        lock.lock();
        super.put(element);
        lock.unlock();
    }
    public ElementType get() {
        lock.lock();
        ElementType retVal = super.get();
        lock.unlock();
        return retVal;
    }
}
```

Create declarative rather than procedural solution

DECLARATIVE VS. PROCEDURAL PRIMITIVE

```
public class ALockingBoundedBuffer<ElementType>
    extends ABoundedBuffer<ElementType>{
    Lock lock;
    public ALockingBoundedBuffer(Lock aLock) {
        lock = aLock;
    }
    public void put(ElementType element) {
        lock.lock();
        super.put(element);
        lock.unlock();
    }
    public ElementType get() {
        lock.lock();
        ElementType retVal = s
        lock.unlock();
        return retVal;
    }
}
```

Declarative: Directive to system by giving declaration specifying *what*

Procedural: Directive to system by executing statements specifying *how*

A la Regular expressions vs. Finite State Automata



HIGH-LEVEL SUPPORT




Critical Section

Other threads should not be executing other critical sections that access the same data

Connect lock to shared data and methods that manipulate them



CLASS WITH ALL ATOMIC METHODS



When declaring a class, provide a label (e.g. synchronized) that says that on an instance a method cannot be executed if the same or another method is executing on that instance

Such a class is a monitor class, and its instance a monitor


Not all classes have critical sections, hence labeling can prevent unnecessary code and data structures in such non monitor classes

LABELING A CLASS AS A MONITOR

```
public synchronized ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int size = 0;
    int nextIn = 0;
    int nextOut = 0;
```

Some methods of the class may not be critical sections

ATOMIC VS NON ATOMIC METHODS



When declaring a method, provide a label (e.g. synchronized). A single synchronized method can execute at one time on an instance

LABELING METHODS AS ATOMIC

```
public synchronized void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
}  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    A lock associated with class instance is automatically obtained on  
    entry to the marked procedure and released on its exit  
    ElementType result = buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    Such a method is classically called an entry procedure  
    return result;  
}
```

Java calls it a **synchronized method**, we will call it an **atomic method**

Entry Queue

ENTRY QUEUE

```
public synchronized void put(
    ElementType element) {
    if (size >= MAX_SIZE) {
        return;
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    nonEmpty.condSignal();
}
public synchronized ElementType get() {
    if (size == 0) {
        return null;
    }
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    nonFull.condSignal()
    return retVal;
}
```



GET CALLED

Entry Queue

Cons. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
→ public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```



GET PROCEEDS AND PUT CALLED

Entry Queue

Prod. 1

Cons. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```



PRODUCER IN ENTRY QUEUE

Entry Queue

Prod. 1

Cons. 1

```
blic synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```



PUT CALLED AGAIN

Entry Queue

Prod. 1

Prod. 2

Cons. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

PRODUCER IN ENTRY QUEUE

Entry Queue

Prod. 2

Prod. 1

Cons. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

CONSUMER ABOUT TO RETURN

Entry Queue

Prod. 2

Prod. 1

Cons. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

CONSUMER LEAVES

Entry Queue

Prod. 2

Prod. 1

Monitor unlocked

```
public synchronized void put(
    ElementType element) {
    if (size >= MAX_SIZE) {
        return;
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    nonEmpty.condSignal();
}

public synchronized ElementType get() {
    if (size == 0) {
        return null;
    }
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    nonFull.condSignal()
    return retVal;
}
```

FIFO SERVICE OF ENTRY QUEUE

Entry Queue

Prod. 2

Prod. 1

```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

MONITOR PROS AND CONS

Higher-level declarative solution

No programmer overhead in creating and using lock

No danger of forgetting to unlock

Can simply look at method headers to determine critical sections

Can forget to label method

Does not provide (general) thread synchronization

SYNCHRONIZATION VS. MUTUAL EXCLUSION

Thread general synchronization (Thread coordination): a thread waiting for another thread to do something

Thread mutual exclusion: : a thread waiting for another thread to leave a critical section

Threads are “competing” rather than “cooperating”

Thread synchronization: general synchronization other than mutual exclusion

Threads are “cooperating” rather than “competing”

THREAD SYNCHRONIZATION NEED

```
public synchronized void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        return;  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return null;  
    }  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
}
```

Semaphores? How many?

Should wait for non full buffer

Should wait for non empty buffer

Monitors do not support general thread synchronization

THREAD SYNCHRONIZATION WITH SEMAPHORES

```
public synchronized void put(ElementType element) {  
    nonFull.semWait();  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.semSignal();  
}  
  
public synchronized ElementType get() {  
    nonEmpty.semWait();  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.semSignal();  
    return retVal;  
}  
}
```

Initial counts of semaphores?

BOUNDED BUFFER (DATA STRUCTURES)

```
public class ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int size = 0;
    int nextIn = 0;
    int nextOut = 0;
    Semaphore nonFull = new ASemaphore(MAX_SIZE);
    Semaphore nonEmpty = new ASemaphore(0);
```

should be able to deposit MAX_SIZE buffers without waiting

General synchronization solution:
Monitors + Semaphores?

SEMAPHORES TOO HEAVYWEIGHT

```
public synchronized void put(ElementType element) {
    nonFull.semWait();
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    nonEmpty.semSignal();
}

public synchronized ElementType get() {
    nonEmpty.semWait();
    ElementType retVal = (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    nonFull.semSignal();
    return retVal;
}
}
```

Semaphore takes steps to ensure atomic execution of semaphore operations

But method has already taken steps to ensure atomicity

Unnecessary expensive system calls to disable context switching



CONDITION (SYNCHRONIZER)

```
public class AConditionSynchronizer
    implements ConditionSynchronizer{
Queue<Thread> queue = new AQueue<Thread>();
public void condWait() {
    Thread currentThread = Thread.currentThread();
    currentThread.setStatus(Thread.WAIT);
    queue.add(currentThread);
}
public void condSignal() {
    Thread nextThread = queue.remove();
    nextThread.setStatus(Thread.READY);
    Thread.reschedAndRestoreInterrupts(
        originalInterrupts);
}
```

Non atomic operations – no disabling of context switching

int count is only one kind of condition for synchronization

No integer count

Synchronization condition checked by caller



POLICY AND MECHANISM TIED IN

```
public synchronized void put(ElementType element) {  
    nonFull.semWait();  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.semSignal();  
}  
  
public synchronized ElementType get() {  
    nonEmpty.semWait();  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.semSignal();  
    return retVal;  
}  
}
```



POLICY AND MECHANISM ORTHOGONAL

```
public synchronized void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return nonEmpty.condWait();  
    }  
    ElementType retVal = (ElementType, buffer[nextOut]);  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Must check programmer-defined boolean expression before waiting

Policy for synchronizing and mechanism for blocking are orthogonal

Could we put check outside semaphore also



LIGHTWEIGHT SEMAPHORE

```
public class ASemaphore implements Semaphore{  
    Queue<Thread> queue = new AQueue<Thread>();  
    public void semWait() {  
        int originalInterrupts = System.disableInterrupts();  
        Thread currentThread = Thread.currentThread();  
        currentThread.setStatus(Thread.WAIT);  
        queue.add(currentThread);  
        System.restoreInterrupts(originalInterrupts);  
    }  
    public void semSignal() {  
        int originalInterrupts = System.disableInterrupts();  
        Thread nextThread = queue.remove();  
        nextThread.setStatus(Thread.READY);  
        Thread.reschedAndRestoreInterrupts(  
            originalInterrupts);  
    }  
}
```

Like regular semaphore, except
no count

Like condition, except there is
atomicity implementation



LIGHTER WEIGHT SEMAPHORE USE

```
public void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.semWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.semSignal();  
}  
  
public ElementType get() {  
    if (size == 0) {  
        return nonEmpty.semWait();  
    }  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.semSignal();  
    return retVal;  
}
```

Test of lock

Enqueuing

Test and enqueueing are not atomic



SYNCHRONIZER OPERATIONS IN ATOMIC METHODS

```
public synchronized void put(ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        return nonEmpty.condWait();  
    }  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Blocking/unblocking operations should be done atomically

Conditions for blocking/unblocking should also be checked atomically with blocking/unblocking operations

Force blocking/unblocking operations to be done only in atomic methods, which can then check and set arbitrary conditions atomically



BOUNDED BUFFER WITH CONDITION (SYNCHRONIZERS): DATA STRUCTURES

```
public class ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int size = 0;
    int nextIn = 0;
    int nextOut = 0;
    ConditionSynchronizer nonFull =
        new AConditionSynchronizer();
    ConditionSynchronizer nonEmpty =
        new AConditionSynchronizer();
```



QUEUES WITH CONDITIONS

Entry Queue

nonEmpty Condition Queue

nonFull Condition Queue

```
public synchronized void put(
    ElementType element) {
    if (size >= MAX_SIZE) {
        nonFull.condWait();
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    nonEmpty.condSignal();
}

public synchronized ElementType get() {
    if (size == 0) {
        nonEmpty.condWait();
    }
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    nonFull.condSignal()
    return retVal;
}
```



GET CALLED

Cons. 1

Entry Queue

nonEmpty Condition Queue

nonFull Condition Queue



```
public synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```



GET PROCEEDS AND PUT CALLED

Entry Queue

nonEmpty Condition Queue

nonFull Condition Queue

Prod. 1

Cons. 1

```
blic synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```



PRODUCER IN ENTRY QUEUE

Entry Queue

Prod. 1

nonEmpty Condition Queue


nonFull Condition Queue

```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Cons. 1



GET CALLED AGAIN



CONSUMER WAITS FOR CONDITION

Entry Queue

Cons. 2

Prod. 1

nonEmpty Condition Queue


nonFull Condition Queue

```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```

Cons. 1



CONSUMER UNLOCKS MONITOR



PRODUCER ENTERS

Entry Queue

Cons. 2

nonEmpty Condition Queue

Cons. 1


nonFull Condition Queue

Prod. 1

```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```



PRODUCER SIGNALS, CONSUMER DEQUEUED



```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal = buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Producer not exited
Where does consumer go next?



SIGNALLED THREAD IN MONITOR?

Entry Queue

Cons. 2

nonEmpty Condition Queue

nonFull Condition Queue



```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Prod. 1

Cons. 1

Violates Monitor Definition



SIGNALLED THREAD IN ENTRY QUEUE

Entry Queue

Cons. 1

Cons. 2

nonEmpty Condition Queue

nonFull Condition Queue

```
blic synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```


Cons. 1 will get the buffer first

Starvation possible

Prod. 1



SIGNALLED THREAD IN URGENT QUEUE



```
blic synchronized void put(  
    ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
    get();  
}  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType result;  
    (ElementType) buffer[nextOut];  
    nextOut = Urgent Q serviced before Entry Q;  
    size--;  
    n  
    r  
}  
Signaled thread in urgent Q  
Signaler may violate condition or previously  
signaled thread may violate it
```

Prod. 1

Signaled thread in urgent Q


Urgent Q serviced before Entry Q

Condition for Entering May Not Exist Later

Signaler may violate condition or previously signaled thread may violate it

SIGNALLED THREAD IN MONITOR AND SIGNALING

T Entry Queue URGENT QUEUE




```
public synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```


Cons. 1

Signaling thread must restore monitor invariant before signaling

SIGNALLED THREAD ABOUT TO LEAVE



URGENT QUEUE SERVICED FIRST



blic synchronized void put(ElementType element) {
 if (size >= MAX_SIZE) {
 nonFull.condWait();
 }
 buffer[nextIn] = element;
 nextIn = (nextIn + 1) % MAX_SIZE;
 size++;
 nonEmpty.condSignal();
}
public synchronized ElementType get() {
 if (size == 0) {
 nonEmpty.condWait();
 }
 ElementTy
 (ElementType) nextOut = buffer[nextIn];
 size--;
 nonFull.c
 return retVal;
}

Prod. 1

Signaled process enters only to leave!

Unnecessary context switch

MAKE SIGNAL A RETURN, LET SIGNALLED THREAD EXIT

Entry Queue

Cons. 2

nonEmpty Condition Queue

nonFull Condition Queue

```
public synchronized void put( ElementType element) {  
    if (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condSignal();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal = (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Prod. 1

X

A single signal allowed

(ElementType) buffer[nextOut];

nextOut = No statement allowed after signal
size--;

nonFull.condSignal()
return retVal; X

Make signal a hint



HINT VS. ABSOLUTE

Absolute: Application is provided information with absolute guarantee

A la non probabilistic algorithm

The server is on machine M a la the person is at address A

The condition for which you were waiting is true

Hint: Application is provided information that is likely to be true but can be checked

A la probabilistic algorithm

The server is probably on machine M, and if the message bounces, check with some central authority at some cost

The condition for which you are waiting is probably true, but must check



HINT VS. ABSOLUTE


Programmer must have code to check truth

If check fails, operations are more expensive

Usual case is efficient, unusual case is possible but not efficient

Servers do not move usually, condition does not usually get violated

SIGNAL IS A HINT



Make usual case efficient and other cases possible

HINT VS. ABSOLUTE PROGRAMMING

```
if (!okToProceed) {  
    condSynchronizer.condWait();  
}
```


```
while (!okToProceed) {  
    condSynchronizer.condWait();  
}
```

Polling?

Not polling as each iteration has a wait

Allows broadcast: all waiting threads
are unblocked


BROADCAST



Prod. 1

```
blic synchronized void put( ElementType element) {  
    while (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condBroadcastSignal();  
}  
public synchronized ElementType get() {  
    while (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```

BROADCAST



```
blic synchronized void put( ElementType element) {  
    while (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condBroadcastSignal();  
}  
  
public synchronized ElementType get() {  
    while (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (Ele  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal();  
    return retVal;  
}
```

Prod. 1

Put all threads on condition Q of condition synchronizer in urgent Q

FIRST SIGNALLED THREAD ENTERS

Entry Queue

nonEmpty Condition Queue

nonFull Condition Queue

Urgent Queue

Cons. 2




```
blic synchronized void put( ElementType element) {  
    while (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condBroadcastSignal();  
}  
  
public synchronized ElementType get() {  
    while (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```

Cons. 1



FIRST SIGNALLED THREAD FINISHES



blic synchronized void put(ElementType element) {
while (size >= MAX_SIZE) {
nonFull.condWait();
}
buffer[nextIn] = element;
nextIn = (nextIn + 1) % MAX_SIZE;
size++;
nonEmpty.condBroadcastSignal();
}
public synchronized ElementType get() {
while (size == 0) {
nonEmpty.condWait();
}
ElementType retVal =
(ElementType) buffer[nextOut];
nextOut = (nextOut + 1) % MAX_SIZE;
size--;
nonFull.condSignal()
return retVal;
}

SECOND SIGNALLED THREAD ENTERS

Entry Queue

nonEmpty Condition Queue

nonFull Condition Queue


Urgent Queue



```
blic synchronized void put( ElementType element) {  
    while (size >= MAX_SIZE) {  
        nonFull.condWait();  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    nonEmpty.condBroadcastSignal();  
}  
  
public synchronized ElementType get() {  
    while (size == 0) {  
        nonEmpty.condWait();  
    }  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    nonFull.condSignal()  
    return retVal;  
}
```

Cons. 2

SECOND SIGNALLED RE-ENTERS CONDITION Q



```
public synchronized void put(
    ElementType element) {
    while (size >= MAX_SIZE) {
        nonFull.condWait();
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    nonEmpty.condBroadcastSignal();
}

public synchronized ElementType get() {
    while (size == 0) {
        nonEmpty.condWait();
    }
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    nonFull.condSignal()
    return retVal;
}
```

HINT PROS AND CONS

Extra check for each signaled thread

Signaled thread may in unusual cases wait again

Signaling thread does not have to restore invariant

Usual case avoids context switch (faster than extra check) or restricted signal (on return)

Allows broadcast semantics

JAVA H

Entry Queue

Single Condition Queue

Urgent Queue

```
public synchronized void put(
    ElementType element) {
    while (size >= MAX_SIZE) {
        try {
            this.wait();
        } catch (Exception e) {}
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    this.notify();
}
```

Single hint-based condition synchronizer and Q) {
in each monitor, named by object (this)

```
try {
    this.wait();
} catch (Exception e) {}
}
ElementType retVal =
(ElementType) buffer[nextOut];
nextOut = (nextOut + 1) % MAX_SIZE;
size--;
this.notify();
return retVal;
```

Both producer and consumer waiting on same
condition synchronizer but different conditions



JAVA HINT SEMANTICS FOR BOUNDED BUFFER

Entry Queue

Single Condition Queue

Prod. 1 X Cons. 2

Urgent Queue

```
public synchronized void put(  
    ElementType element) {  
    while (size >= MAX_SIZE) {  
        try {  
            this.wait();  
        } catch (Exception e) {}  
    }  
    nextIn = (nextIn + 1) % MAX_SIZE;  
    size++;  
    this.notify(); this.notifyAll();  
}  
  
public synchronized ElementType get() {  
    if (size == 0) {  
        this.wait();  
    }  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    size--;  
    this.notify(); this.notifyAll();  
    return retVal;  
}
```

notifyAll for Bounded Buffer?

Buffer cannot be full and empty at the same time

Bounded Buffer is driving problem

Multiple conditions share a synchronizer, hence modified name



JAVA VS. GENERAL HINT SEMANTICS

Every object has an entry and condition queue

Probably why primitive values are not objects as in Smalltalk

May need broadcast semantics even when one thread eligible

Difficult to imagine cases when that happens

No need for special class labeling as monitor

No need to declare condition synchronizer

ALWAYS NONEMPTY, NONFULL BUFFER

```
public void put(  
    ElementType element) {  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
  
}  
public ElementType get() {  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    return retVal;  
}
```

Puts cannot be executed concurrently

Gets cannot be executed concurrently

Puts and gets can be executed concurrently

ATOMIC METHODS?

```
public synchronized void put(  
    ElementType element) {  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
  
}  
public synchronized ElementType get() {  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    return retVal;  
}
```



Puts and gets can be executed concurrently

CONDITIONS?

```
public void put(
    ElementType element) {
    if (putLocked) {
        wax(');
        buffer[nextIn] = element;
        nextIn = (nextIn + 1) % MAX_SIZE;
        putLocked = false;
        notifx'l();
    }
}

public ElementType get() {
    ElementType retVal =
        (ElementType) buffer[nextOut];
    if (getLocked) {
        wait();
        getLocked = true;
        nextOut = (nextOut + 1) % MAX_SIZE;
        getLocked = false;
    }
}
```

Condition synchronizer operations and associated checks should occur in atomic methods



SYNCHRONIZED STATEMENT BLOCK



putLock



getLock

```
public void put(
    ElementType element) {
    synchronized (putLock) {
        buffer[nextIn] = element;
        nextIn = (nextIn + 1) % MAX_SIZE;
    }
}

public ElementType get() {
    ElementType retVal =
        (ElementType) buffer[nextOut];
    synchronized (getLock) {
        getLocked = true;
    }
}
```

putLock and getLock are monitors with no atomic methods associated with statement block


synchronize(object) { <statement list> } gets monitor lock at start of statement list and releases it on end, without entering monitor

LOCK DECLARATIONS

```
public class ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int nextIn = 0;
    int nextOut = 0;
    Object putLock = new Object();
    Object getLock = new Object();
```


putLock and getLock are monitors with no atomic methods: Object instances

SYNCHRONIZED CALLING ANOTHER SYNCHRONIZED METHOD IN SAME MONITOR



Once a thread has the monitor lock, it can execute any number of atomic(and non atomic) methods before returning from initial method

SYNCHRONIZED CALLING ANOTHER SYNCHRONIZED METHOD IN SAME MONITOR




When synchronized method in one monitor calls synchronized method in another monitor, it does not release lock of first monitor

Does not have to re-get lock and restore monitor invariant before calling external method, can execute both atomically


Can lead to deadlocks

DEADLOCK (DEADLY EMBRACE)




Some designs release lock when calling external method

WAITING IN SYNCHRONIZED METHOD



Lock released when wait called on this in a synchronized method

WAITING IN EXTERNAL MONITOR IN SYNCHRONIZED METHOD



In Java: No-op (not illegal) when wait called on another monitor in a synchronized method

WAITING IN EXTERNAL MONITOR IN SYNCHRONIZED METHOD

```
synchronized(putLock) {  
    try {  
        putLock.wait();  
    } catch (Exception e) {}  
    ...  
}
```

In Java: Synchronized statement block can wait on monitor associated with block

Can wait on or notify condition queue of monitor iff you have that monitor lock

SLEEPING IN SYNCHRONIZED METHOD



Monitor 1

S method1()

`Thread.sleep(t);`

In Java: all waiting calls keep the lock

Even though waiting method can do nothing and Java knows it is waiting

Better semantics: require invariant to be restore when lock is gotten



HIGHER-LEVEL SOLUTION?

Entry Queue

Single Condition Queue

Cons. 2

Urgent Queue

```
public synchronized void put(  
    ElementType element) {  
    while (size >= max)  
        try{  
            this.wait()  
        } catch(Exception e, {}  
    }  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % max;  
    size++;  
    this.notify();  
}
```

Declarative for mutual exclusion

Procedural for (other) synchronization

Declarative Synchronization?

Hint: FSA vs. Regular Expressions?

Vocabulary is methods

REMOVE ALL SYNCHRONIZATION CODE

```
public synchronized void put(
    ElementType element) {
    while (size >= MAX_SIZE) {
        this.wait();
    }
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
    size++;
    this.notify();
}

public synchronized ElementType get() {
    while (size == 0) {
        this.wait();
    }
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    size--;
    this.notify();
    return retVal;
}
```

No synchronized keyboard

No wait and notify

No check for ok to proceed

No increment/decrement
of size

BOUNDED BUFFER (DATA STRUCTURES)

```
public class ABoundedBuffer<ElementType>
    implements BoundedBuffer<ElementType>{
    public static final int MAX_SIZE = 10;
    Object[] buffer = new Object[MAX_SIZE];
    int size = 0;
    int nextIn = 0;
    int nextOut = 0;
```

No size field

REGULAR VS. PATH EXPRESSIONS

Regular

Vocabulary = { $a_1, \dots a_N$ }

Usually characters

Defines legal sequences of vocabulary elements

$a \mid d$

either a or d is acceptable

Path

Vocabulary = { $m_1, \dots m_N$ }

Method calls

Defines legal sequences of method calls in a module

get, put

either get or put can be called

REGULAR VS. PATH EXPRESSIONS

Regular

$R = a_i$

$R = R_1 \ R_2$

R_1 must be followed by R_2

$R = R_1 \mid R_2$

R_1 or R_2 is legal

R^*

Zero or more occurrences of R can be matched next

Path

$P = m_i$

$P = P_1;P_2$

P_1 must be called after P_2

$P = P_1,P_2$

Call to P_1 or P_2 can be made

$P = [P]$

Zero or more occurrences of call to P can be executed

$N: (P)$

Up to N concurrent executions of P can be executed



REGULAR VS. PATH EXPRESSIONS (REVIEW)

Regular

Vocabulary = $\{a_1, \dots a_N\}$

Usually characters

Defines legal sequences of vocabulary elements

a | d

either a or d is acceptable

Path

Vocabulary = $\{m_1, \dots m_N\}$

Method calls

Defines legal sequences of method calls in a module

get, put

either get or put can be called

REGULAR VS. PATH EXPRESSIONS (REVIEW)

Regular

$R = a_i$

$R = R_1 \ R_2$

R_1 must be followed by R_2

$R = R_1 \mid R_2$

R_1 or R_2 is legal

R^*

Zero or more occurrences of R can be matched next

Path

$P = m_i$

$P = P_1;P_2$

P_1 must be called after P_2

$P = P_1,P_2$

Call to P_1 or P_2 can be made

$P = [P]$

Zero or more occurrences of call to P can be executed

$N: (P)$

Up to N concurrent executions of P can be executed

SPECIFYING MUTUAL EXCLUSION

```
public void put(  
    ElementType element) {  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
}  
  
public ElementType get() {  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    return retVal;  
}
```

put and get can be executed concurrently, as no shared size variable

[get, put]

[get], [put]

get, put

An arbitrary number of activations of get or put and get can be executed concurrently, [] implicit on outermost expression

1: (get, put)

Only one activation of get or put can be executed concurrently

1: (get) , 1: (put)

Only one execution of get or put can be active at any one time

SPECIFYING SYNCHRONIZATION

```
public void put(
    ElementType element) {
    buffer[nextIn] = element;
    nextIn = (nextIn + 1) % MAX_SIZE;
}
public ElementType get() {
    ElementType retVal =
        (ElementType) buffer[nextOut];
    nextOut = (nextOut + 1) % MAX_SIZE;
    return retVal;
}
```

Must define an order between
get and put

SPECIFYING SYNCHRONIZATION

```
public void put(  
    ElementType element) {  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
}  
  
public ElementType get() {  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    return retVal;  
}
```

Puts can execute concurrently
with another put

No mutual exclusion!

[put; get]

put; get

Each activation of get preceded
by put, an infinite number of
concurrent put; get sequences

1 :(put; get)

A single put; get sequence can
be active at any one time

The two methods alternate, BB
of size 1

N :(put; get)

Up to N activations of (put;get)
can be active at any one time

SPECIFYING MUTUAL EXCLUSION AND SYNCHRONIZATION

```
public void put(  
    ElementType element) {  
    buffer[nextIn] = element;  
    nextIn = (nextIn + 1) % MAX_SIZE;  
}  
  
public ElementType get() {  
    ElementType retVal =  
        (ElementType) buffer[nextOut];  
    nextOut = (nextOut + 1) % MAX_SIZE;  
    return retVal;  
}
```

N : (path expression) gives global restriction, not within subexpression

Mutual Exclusion

1: (get), 1: (put)

Only one execution of get or put can active

Synchronization

N :(put; get)

N activations of (put;get) can be active at any one time

Both

N :(1:(put); 1:(get))

N activations of (put;get) can be active but only one activation of put and one activation of get



READER WRITER PROBLEM

There can be an arbitrary number of readers or a single writer active at any one time

1: ([read], write)

Fair readers writers: writer does not wait for a reader who comes after it

Path expression cannot handle all synchronization schemes

Regular expression cannot handle all languages

PATH EXPRESSIONS

Method calls are matched according to specified path expressions

Method calls that do not match are blocked

After top-level path expression is matched , it can be matched again

1: (get, put)

Implementation unspecified: when is unblocking checking checked and which waiting method unblocked

Not practical but helps understand synchronization better