Concurrency in Java

COMS W4995-02

Prof. Stephen A. Edwards
Fall 2002
Columbia University
Department of Computer Science

The Java Language

Developed by James Gosling et al. at Sun Microsystems in the early 1990s

Originally called "Oak"

First intended application was as an OS for TV set top boxes

Main goals were portability and safety

Originally for embedded consumer software

The Java Language

Set-top boxes: nobody cared at the time

Next big application: "applets," little programs dynamically added to web browsers

A partial failure despite agressive Sun marketing.

- Incompatible Java implementations
- Few users had enough bandwidth
- Fantastically slow Java interpreters

Javascript, a high-level scripting language that has nothing to do with the Java language, but leverages Sun's marketing, has taken over this role.

Aside: The JVM

```
int gcd(int a, int b)
{
  while (a != b) {
    if (a > b)
      a -= b;
    else
      b -= a;
  }
  return a;
}
```

A stack-based machine langauge. Mostly low-level operations (e.g., add top two stack elements), but also virtual method dispatch.

```
# javap -c Gcd

Method int gcd(int, int)
   0 goto 19

   3 iload_1          // Push a
   4 iload_2          // Push b
   5 if_icmple 15     // if a <= b goto 15

   8 iload_1          // Push a
   9 iload_2          // Push b
  10 isub             // a - b
  11 istore_1         // Store new a
  12 goto 19

  15 iload_2          // Push b
  16 iload_1          // Push a
  17 isub             // b - a
  18 istore_2         // Store new b

  19 iload_1          // Push a
  20 iload_2          // Push b
  21 if_icmpne 3      // if a != b goto 3

  24 iload_1          // Push a
  25 ireturn          // Return a
```

Aside: The JVM

Advantages:

- Trivial translation of expressions
- Trivial interpreters
- No problems with exhausting registers
- Often compact

Disadvantages:

- Semantic gap between stack operations and modern register machines
- Hard to see what communicates with what
- Difficult representation for optimization

JVM: The Lesson

> *If you're going to interpret something, its level of abstraction must be sufficiently higher than assembly to offset the cost of interpretation.*

Java bytecode is stack-based and therefore fairly close to assembly. Usual interpreters run at 1/10th the speed of equivalent C.

Just-in-time compilers, which translate Java bytecode to native assembly *and cache the results*, are the current solution.

The Java Language

Where does Java succeed?

Corporate programming

- E.g., dynamic web page generation from large corporate databases in banks
- Environment demands simpler language
- Unskilled programmers, unreleased software
- Speed, Space not critical
- Tends to be run on very large servers
- Main objective is reduced development time

The Java Language

Where does Java succeed?

Education

- Well-designed general-purpose programming language
- Spares programmer from many common pitfalls:
- Uninitialized pointers
- Memory management
- Widely known and used, not just a teaching language

Embedded Systems?

- Jury is still out

Concurrency in Java

Language supports threads

Multiple contexts/program counters running within the same memory space

All objects shared among threads by default

Fundamentally nondeterministic

Language provide synchronization facilities (monitors) to help avoid nondeterminsm

Still a difficult paradigm in which to program

Sun's libraries reputed to still not be thread-safe

Thread Basics

A thread is a separate program counter

... and stack, local variables, etc.

Not an object or a collection of things

Classes, objects, methods, etc. do not belong to a thread

But a thread may hold a lock on an object

Any method may be executed by one or more threads, even simultaneously

Thread Basics

How to create a thread:

```
class MyThread extends Thread {
  public void run() { // A thread's "main" method
    /* thread body */
  }
}

MyThread mt = new MyThread; /* Create thread */
mt.start();   /* Start thread at run() */
// Returns immediately
```

implements Runnable vs. extends Thread

An alternative:

```
class MyRunnable implements Runnable {
  public void run() {
    /* thread body */
  }
}

Thread t = new Thread(new MyRunnable());
t.start();   /* Starts thread at run() */
// Returns immediately
```

Advantage: class implementing Runnable can be derived from some other class (i.e., not just from Thread).

Disadvantage: "this" is not a Thread so, e.g., sleep() must be called with Thread.currentThread().sleep().

The Sleep Method

```
class Sleeper
  extends Thread {
  public void run() {
    for (;;) {
      try {
        sleep(1000); // Pause for at least a second
      } catch (InterruptedException e) {
        return;   // caused by thread.interrupt()
      }
      System.out.println("tick");
    }
  }
}
```

Does this print "tick" once a second?

No: the sleep() delay is merely a lower bound, and it's not clear how much time the rest of the loop takes.

A Clock?

```
class PrintingClock implements Runnable {
  public void run() {
    for (;;) {
      java.util.Date now = new java.util.Date();
      System.out.println(now.toString());
      try {
        Thread.currentThread().sleep(1000);
      } catch (InterruptedException e) {}
    }
  }
}
public class Clock {
  public static void main(String args[]) {
    Thread t = new Thread(new PrintingClock());
    t.start();
  }
}
```

What does the clock print?

```
$ java Clock
Sat Sep 14 13:04:27 EDT 2002
Sat Sep 14 13:04:29 EDT 2002
Sat Sep 14 13:04:30 EDT 2002
Sat Sep 14 13:04:31 EDT 2002
```

What happened to 13:04:28?

Synchronization

Motivation for Synchronization

Something you might want to implement:

```
class ScoreKeeper {
  int _score = 0;
  void score(int v) {
    int tmp = _score;
    tmp += v;
    _score = tmp;
  }
}
```

What could the final score be if two threads simultaneously call `score(1)` and `score(2)`?

Non-atomic Operations

Java guarantees 32-bit reads and writes to be atomic

64-bit operations may not be

Therefore,

```
int i;
double d;
```

Thread 1	Thread 2
`i = 10;` `d = 10.0;`	`i = 20;` `d = 20.0;`

`i` will contain 10 or 20.

`d` may contain 10, 20, or garbage.

Per-Object Locks

Each Java object has a lock that may be owned by at least one thread

A thread waits if it attempts to obtain an already-obtained lock

The lock is a counter: one thread may lock an object more than once

The Synchronized Statement

A synchronized statement gets an object's lock before running its body

```
 Counter mycount = new Counter;
synchronized(mycount) {
  mycount.count();
 }
```

Releases the lock when the body terminates.

Choice of object to lock is by convention.

Synchronized Methods

```
class AtomicCounter {

  private int _count;

  public synchronized void count() {
    _count++;
  }
}
```

`synchronized` attribute equivalent to enclosing body with `synchronized (this)` statement.

Most common way to achieve object-level atomic operations.

Implementation guarantees at most one thread can run a synchronized method on a particular object at once

Deadlock

```
synchronized(Foo) {
  synchronized(Bar) {
    // Asking for trouble
  }
}
```

```
synchronized(Bar) {
  synchronized(Foo) {
    // Asking for trouble
  }
}
```

Rule: always acquire locks in the same order.

Synchronization

Say you want a thread to wait for a condition before proceeding.

An infinite loop may deadlock the system (e.g., if it's using cooperative multitasking).

```
while (!condition) {}
```

Calling `yield` avoids deadlock, but is inefficient:

```
while (!condition) yield();
```

Scheduler may choose to run this thread often, even though condition has not changed.

Java's Solution: wait() and notify()

`wait()` like `yield()`, but requires other thread to reawaken it

```
while (!condition) wait();
```

Thread that changes the condition calls `notify()` to resume the thread.

Programmer responsible for ensuring each `wait()` has a matching `notify()`.

Wait() and Notify() in Real Life

I often have books delivered to the CS department office.

This operation consists of the following steps:

1. Place the order
2. Retrieve the book from the CS department office
3. Place book on bookshelf

Obviously, there's a delay between steps 1 and 2.

The Implementation in the CS Office

This "order, retrive, file" thread is running in me, and it needs to wait for the book to arrive.

I could check the department office every minute, hour, day, etc. to see if the book has come in, but this would be waste of time (but possibly good exercise).

Better approach would be to have the "receiving" process alert me when the book actually arrives.

This is what happens: Alice in the front office sends me email saying I have a book.

Implementing the CS Office

```
class Mailbox {}

public class BookOrder {
  static Mailbox m = new Mailbox();

  public static void main(String args[]) {
    System.out.println("Ordering book");
    Thread t = new Delivery(m);
    t.start();
    synchronized (m) {
      try {
         m.wait(); // Must own lock on m to wait
      } catch (InterruptedException e)
    }
    System.out.println("Book Arrived");
  }
}
```

Implementing the CS Office

```
class Delivery extends Thread {
  Mailbox m;
  Delivery(Mailbox mm) { m = mm; }
  public void run() {
    try {
      sleep(1000);   // one-second delivery
    } catch (InterruptedException e) {}
    synchronized (m) {
      m.notify(); // Must own lock on m to notify it
    }
  }
}
```

Harder Problem

Sometimes I have a number of books on order, but Alice only tells me a book has arrived. How do I handle waiting for a number of books?

Last solution assumed a single source of `notify()`*; not true in general.*

Two rules:

1. Use `notifyAll()` when more than one thread may be waiting
2. Wait in a `while` if you could be notified early, e.g.,

```
while (!condition) {
  wait();
}
```

The multi-book problem

```
class Mailbox {
  int book;
  Mailbox() {book = -1; }
  void receive_book(int b) {book = b; }
  int which_book() {return book; }
}
```

This is not thread-safe: we'll need to synchronize all access to it.

The multi-book problem

The Delivery class also tells the mailbox which book it got.

```
class Delivery extends Thread {
  Mailbox m;
  int book;
  Delivery(Mailbox mm, int b) {
     m = mm; book = b;
  }
  public void run() {
    try {sleep(1000); }
    catch (InterruptedException e) {}
    synchronized (m) {
      m.receive_book(book);
      m.notifyAll();
    }
  }
}
```

The multi-book problem

```
class BookOrder extends Thread {
  int book;
  Mailbox m;
  BookOrder(Mailbox mm, int b) {
     m = mm; book = b;
  }
  public void run() {
    System.out.println("Ordering book " +
                 Integer.toString(book) );
    Thread t = new Delivery(m, book);
    t.start();
    synchronized (m) {
      while (m.which_book() != book) {
        try {m.wait(); }
        catch (InterruptedException e) {}
      }
    }
    System.out.println("Book " +
     Integer.toString(book) + " Arrived");
  }
}
```

The multi-book problem

Finally, the main routine kicks off two ordering threads.

```
public class MultiOrder {
  static Mailbox m = new Mailbox();
  public static void main(String args[]) {
    BookOrder t1 = new BookOrder(m, 1);
    BookOrder t2 = new BookOrder(m, 2);
    t1.start();
    t2.start();
  }
}
```

```
$ java MultiOrder
Ordering book 1
Ordering book 2
Book 1 Arrived
Book 2 Arrived
```

A Better Solution

Last solution relied on threads to synchronize their own access to the Mailbox. Mailbox should be doing this itself:

```
class Mailbox {
  int book;
  Mailbox() { book = -1; }
  synchronized void deliver(int b) {
    book = b;
    notifyAll();
  }
  synchronized void wait_for_book(int b) {
    while (book != b) {
      try { wait(); }
      catch (InterruptedException e) {}
    }
  }
}
```

A Better Solution

```
class Delivery extends Thread {
  Mailbox m;
  int book;
  Delivery(Mailbox mm, int b) {
    m = mm; book = b;
  }
  public void run() {
    try { sleep(1000); }
    catch (InterruptedException e) {}
    m.deliver(book);
  }
}
```

A Better Solution

```
class BookOrder extends Thread {
  int book;
  Mailbox m;
  BookOrder(Mailbox mm, int b) {
     m = mm; book = b;
  }
  public void run() {
    System.out.println("Ordering book " +
      Integer.toString(book) );
    Thread t = new Delivery(m, book);
    t.start();
    m.wait_for_book(book);
    System.out.println("Book " +
     Integer.toString(book) + " Arrived");
  }
}
```

A Better Solution

```
public class MultiOrder2 {
  static Mailbox m = new Mailbox();
  public static void main(String args[]) {
    BookOrder t1 = new BookOrder(m, 1);
    BookOrder t2 = new BookOrder(m, 2);
    t1.start();
    t2.start();
  }
}

$ java MultiOrder2
Ordering book 1
Ordering book 2
Book 1 Arrived
Book 2 Arrived
```

Building a Blocking Buffer

Problem: Build a single-place buffer for Objects that will block on write if the buffer is not empty and on read if the buffer is not full.

```
interface OnePlace {
  public void write(Object o);
  public Object read();
}
```

Blocking Buffer: Write Method

```
class MyOnePlace implements OnePlaceBuf {
  Object o;

  public synchronized // ensure atomic updates
  void write(Object oo) {

    try {
      while (o != null)
        wait();   // Block while buffer is full
    } catch (InterruptedException e) {}

    o = oo;       // Fill the buffer
    notifyAll();  // Awaken any waiting processes
  }
}
```

Blocking Buffer: Read Method

```
class MyOnePlace implements OnePlaceBuf {
  Object o;

  public synchronized // ensure atomic updates
  Object read() {

    try {
      while (o == null)
        wait();   // Block while buffer is empty
    } catch (InterruptedException e) {}

    Object oo = o; // Get the object
    o = null;      // Empty the buffer
    notifyAll();   // Awaken any waiting processes
    return oo;
  }
}
```

Blocking Buffer: Reader and Writer

```
class Writer extends Thread {
  OnePlaceBuf b;
  Writer(OnePlaceBuf bb) {b = bb; }
  public void run() {
    for (int i = 0 ; ; i++ ) {
      b.write(new Integer(i)); // Will block
    }
  }
}
class Reader extends Thread {
  OnePlaceBuf b;
  Reader(OnePlaceBuf bb) {b = bb; }
  public void run() {
    for (;;) {
      Object o = b.read(); // Will block
      System.out.println(o.toString());
    }
  }
}
```

Blocking Buffer: Main Routine

```
public class OnePlace {
  static MyOnePlace b = new MyOnePlace();
  public static void main(String args[]) {
    Reader r = new Reader(b);
    Writer w = new Writer(b);
    r.start();
    w.start();
  }
}
```

Thread Priorities

Priorities

Each thread has a priority from 1 to 10 (5 typical)

Scheduler's job is to keep highest-priority threads running

`thread.setPriority(5)`

What the Language Spec. Says

From *The Java Language Specification*

> Every thread has a priority. When there is competition for processing resources, threads with higher priority are *generally* executed in preference to threads with lower priority. Such preference is *not, however, a guarantee* that the highest priority thread will always be running, and thread *priorities cannot be used* to reliably implement mutual exclusion.

Vague enough for you?

Multiple threads at same priority?

Language definition gives implementer freedom

Calling yield() suspends current thread to allow other at same priority to run ... maybe

Solaris implementation runs threads until they stop themselves with wait(), yield(), etc.

Solaris uses co-operative, application-level threads

Windows implementation timeslices because it uses native threads

Starvation

Java does not demand a fair scheduler

Higher-priority threads can consume all resources, prevent lower-priority threads from running

This is called *starvation*

Timing dependent: function of program, hardware, and Java implementation

How do you know if your program suffers from starvation?

Implementing Threads

Processes and Threads

Many operating systems now distinguish between processes and threads:

Process A thread of control running with its own address space, stack, etc. Uses inter-process communication (e.g., Pipes) with other processes.

Thread A thread of control sharing an address space with another thread, but with its own stack, registers, etc. Communicates through shared memory.

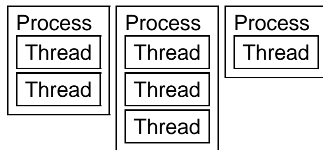

Typical Thread Implementations

"Native threads" Implemented by the operating system

- Scheduled by the OS
- Preemptive
- Context switching costly (jump to OS land)
- Can provide true parallelism on multiprocessors

"Green threads" Implemented by the process

- Scheduled by the process
- Cooperative: control only relinquished voluntarily
- Do not permit true parallelism
- Usually faster

Implementing Java Threads

Many-to-one: Java VM in a single process/OS thread

Scheduler implemented directly in the JVM

+ Cheaper context switches (no need to involve the OS)
+ Does not rely on particular OS API
– Must carefully wrap all OS calls to avoid them blocking the process
– Can't call other libraries since they may make blocking OS calls
– Often difficult to support per-thread system objects, e.g., multiple network connections

Implementing Java Threads

One-to-one: Each Java thread mapped gets its own OS thread

+ Can exploit multiprocessors (OS can schedule different threads on different processors)
+ No need to wrap OS calls or other libraries
– More expensive context switching and thread control. Everything is a system call.
– Less portable

Java Thread Implementations

Solaris Supports Light-Weight Processes (OS threads) and Application Threads (in-process threads).

- Java VM uses Application-level threads by default,
- Application thread scheduler can move threads to other LWPs
- Has 2^{31} priority levels.
- Java thread implementation is non-preemptive

Java Thread Implementations

Windows 95 etc. Supports OS threads only

- Java VM uses native threads only
- Has only 7 priority levels for threads
- Java thread implementation is preemptive

Linux Supports kernel-level POSIX threads

- Java VM uses native threads

Disturbing Conclusion

Since it is very easy to write a threaded Java program whose behavior differs depending on the scheduling policy, Java programs are not the exemplars of portability as touted by Sun.

For example, consider

```
bool waiting = true;
while (waiting) {}  // wait for waiting=false
```

This deadlocks under a (non-preemptive) green threads implementation, but it might work fine with (preemptive) native threads.

Thread Miscellany

Thread-Related Methods

Object.wait()	Wait indefinitely to be notified
Object.wait(long t)	Wait at most *t* milliseconds
Object.wait(long t, int n)	Wait *t* plus *n* nanoseconds
Object.notify()	Release one *waiting* thread
Object.notifyAll()	Release all *waiting* threads
Thread.interrupt()	Break from wait, sleep, etc.
Thread.sleep(long t)	Sleep for *t* milliseconds
Thread.yield()	Pass control to another runnable thread
Thread.join()	Wait for given thread to terminate
Thread.join(long t)	Wait for termination or timeout
Thread.setPriority(int)	Set thread's scheduling priority
Thread.getPriority()	Get thread's priority

Java Thread States

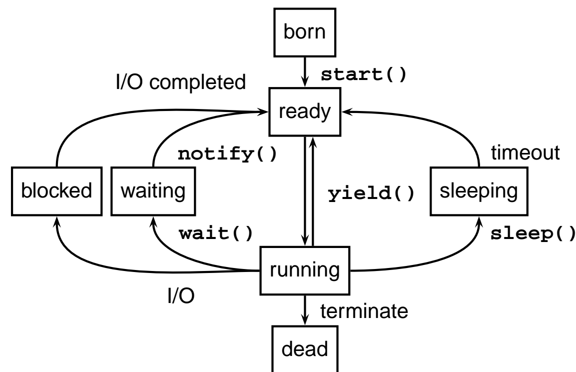

Deprecated Thread methods

Before JDK 1.2, Thread had three other methods:

Thread.stop()	Terminate a thread and release locks
Thread.suspend()	Suspend thread without releasing locks
Thread.resume()	Resume a suspended thread

These were terribly unsafe and a recipe for disaster:

- stop() would release locks, but could easily leave objects in an inconsistent state. Catching the exception it threw would be possible, but insanely messy.
- suspend() did not release locks, so it could cause a deadlock if the thread meant to resume() it needed locks before it could proceed.

Common Mistakes

- Forgetting to catch the **`InterruptedException`** potentially thrown by **`wait()`**, **`sleep()`**, and other blocking thread operations.
- Forgetting the **`notify()`** for a corresponding **`wait()`** (The "Charlie Brown at the Mailbox" error.)
- Waiting on an object without having its lock (i.e., without having **`synchronized`** on the object). Throws an **`IllegalMonitorStateException`**.