

CMSC 838F: Programming Language Techniques for Concurrent and Distributed Systems

Welcome!

- Class mechanics:
 - Class participation (35%)
 - Homework and programming projects
 - Paper presentation during the semester
 - Discussion of papers during class
 - Project (40%)
 - Final exam (25%)
- No class on Monday, Jan 29

2

What is this Class About?

- Concurrent and distributed programming becoming mainstream
 - Multicore processors
 - The Internet
- How do we *program* these systems?
- What do the programs *mean*?
- How can we *reason* about such programs?

3

Why Concur. and Distrib. Prog?

- Performance
 - Exploit more resources
 - Reduce data movement
- Design
 - Some tasks can be divided up naturally
- Failure isolation
- Heterogeneous trust

4

The Plan


- Get our feet wet with threading in Java
 - Project due Wed, Jan 31
- Next time go over advanced threading idioms in Java
- Then back to the basics
 - Memory consistency models
 - Programming models and language design
 - Transactions
 - Analysis

5

CMSC 838F: Programming Language Techniques for Concurrent and Distributed Systems

Java Threads


Computation Abstractions



A computer

7

Processes vs. Threads



Processes do not share data

Threads share data within a process

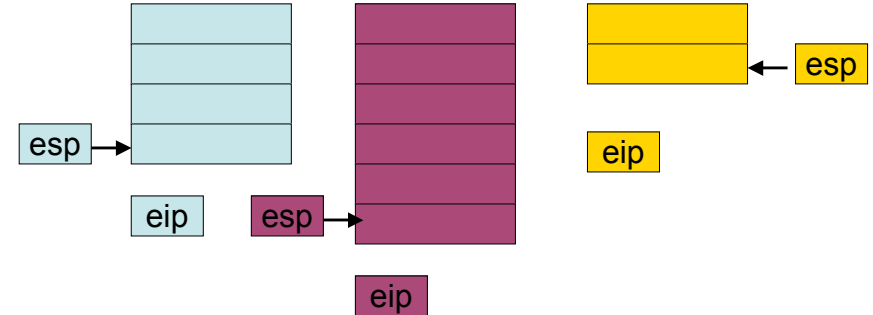
8

So, What Is a Thread?

- **Conceptually:** it is a parallel computation occurring within a process
- **Implementation view:** it's a program counter and a stack. The heap and static area are shared among all threads
- All programs have at least one thread (main)

9

Implementation View



- Per-thread stack and instruction pointer
 - Saved in memory when thread suspended
 - Put in hardware esp/eip when thread resumes

10

Tradeoffs

- Threads can increase performance
 - Parallelism on multiprocessors
 - Concurrency of computation and I/O
- Natural fit for some programming patterns
 - Event processing
 - Simulations
- But increased complexity
 - Need to worry about safety, liveness, composition
- And higher resource usage

11

Programming Threads

- Threads are available in many languages
 - C, C++, Objective Caml, Java, SmallTalk ...
- In many languages (e.g., C and C++), threads are a platform specific add-on
 - Not part of the language specification
- They're part of the Java language specification

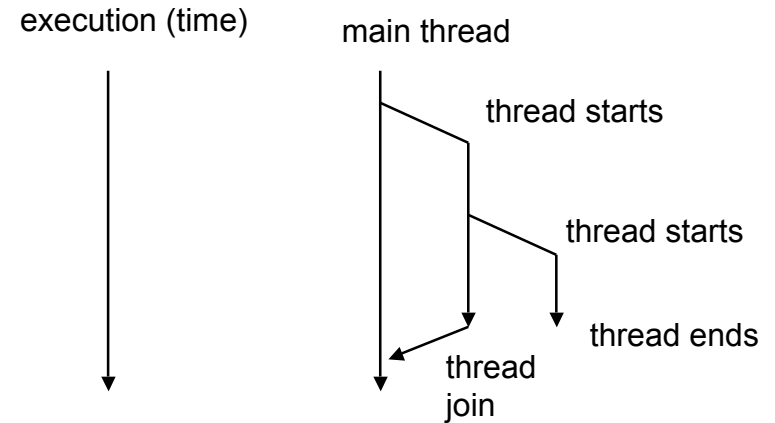
12

Java Threads

- Every application has at least one thread
 - The “main” thread, started by the JVM to run the application’s `main()` method
- `main()` can create other threads
 - Explicitly, using the `Thread` class
 - Implicitly, by calling libraries that create threads as a consequence
 - RMI, AWT/Swing, Applets, etc.

13

Thread Creation



14

Thread Creation in Java

- To explicitly create a thread:
 - Instantiate a `Thread` object
 - An object of class `Thread` or a subclass of `Thread`
 - Invoke the object’s `start()` method
 - This will start executing the `Thread`’s `run()` method concurrently with the current thread
 - Thread terminates when its `run()` method returns

15

Running Example: Alarms

- Goal: let’s set alarms which will be triggered in the future
 - Input: time `t` (seconds) and message `m`
 - Result: we’ll see `m` printed after `t` seconds

16

Example: Synchronous alarms

```
while (true) {
    System.out.print("Alarm> ");

    // read user input
    String line = b.readLine();
    parseInput(line); // sets timeout

    // wait (in secs)
    try {
        Thread.sleep(timeout * 1000);
    } catch (InterruptedException e) { }
    System.out.println("(" + timeout + ") " + msg);
}
```

17

Making It Threaded (1)

```
public class AlarmThread extends Thread {
    private String msg = null;
    private int timeout = 0;

    public AlarmThread(String msg, int time) {
        this.msg = msg;
        this.timeout = time;
    }

    public void run() {
        try {
            Thread.sleep(timeout * 1000);
        } catch (InterruptedException e) { }
        System.out.println("(" + timeout + ") " + msg);
    }
}
```

18

Making It Threaded (2)

```
while (true) {
    System.out.print("Alarm> ");

    // read user input
    String line = b.readLine();
    parseInput(line);
    if (m != null) {
        // start alarm thread
        Thread t = new AlarmThread(m, tm);
        t.start();
    }
}
```

19

Alternative: The Runnable Interface

- Extending `Thread` prohibits a different parent
- Instead implement `Runnable`
 - Declares that the class has a `void run()` method
- Construct a `Thread` from the `Runnable`
 - Constructor `Thread(Runnable target)`
 - Constructor `Thread(Runnable target, String name)`

20

Thread Example Revisited

```
public class AlarmRunnable implements Runnable {
    private String msg = null;
    private int timeout = 0;

    public AlarmRunnable(String msg, int time) {
        this.msg = msg;
        this.timeout = time;
    }

    public void run() {
        try {
            Thread.sleep(timeout * 1000);
        } catch (InterruptedException e) { }
        System.out.println("(" + timeout + ") " + msg);
    }
}
```

21

Thread Example Revisited (2)

```
while (true) {
    System.out.print("Alarm> ");

    // read user input
    String line = b.readLine();
    parseInput(line);
    if (m != null) {
        // start alarm thread
        Thread t = new Thread(
            new AlarmRunnable(m, tm));
        t.start();
    }
}
```

22

Notes: Passing Parameters

- `run()` doesn't take parameters
- We "pass parameters" to the new thread by storing them as private fields
 - In the extended class
 - Or the `Runnable` object
 - Example: the time to wait and the message to print in the `AlarmThread` class


23

Concurrency

- A *concurrent* program is one that has multiple threads that may be active at the same time
 - Might run on one CPU
 - The CPU alternates between running different threads
 - The *scheduler* takes care of the details
 - Switching between threads might happen *at any time*
 - Might run *in parallel* on a *multiprocessor* machine
 - One with more than one CPU
 - May have multiple threads per CPU
- Multiprocessor machines are becoming more common
 - Multi-CPU machines aren't that expensive any more
 - Dual-core CPUs are available now

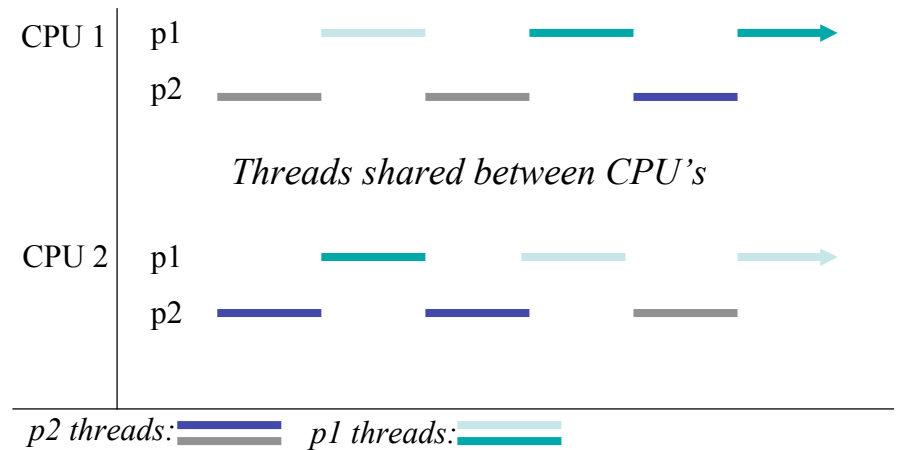
24

Scheduling Example (1)



25

Scheduling Example (2)



26

Concurrency and Shared Data

- Concurrency is easy if threads don't interact
 - Each thread does its own thing, ignoring other threads
 - Typically, however, threads need to communicate with each other
- Communication is done by *sharing* data
 - In Java, different threads may access the heap simultaneously
 - But the scheduler might interleave threads arbitrarily
 - Problems can occur if we're not careful.

27

Data Race Example

```
public class Example extends Thread {
    private static int cnt = 0; // shared state
    public void run() {
        int y = cnt;
        cnt = y + 1;
    }
    public static void main(String args[]) {
        Thread t1 = new Example();
        Thread t2 = new Example();
        t1.start();
        t2.start();
    }
}
```

28

Data Race Example

```
static int cnt = 0;
t1.run() {
  int y = cnt;
  cnt = y + 1;
}
t2.run() {
  int y = cnt;
  cnt = y + 1;
}
```

Shared state cnt = 0

Start: both threads ready to run. Each will increment the global cnt.

29

Data Race Example

```
static int cnt = 0;
t1.run() {
  int y = cnt;
  cnt = y + 1;
}
t2.run() {
  int y = cnt;
  cnt = y + 1;
}
```

Shared state cnt = 0

y = 0



T1 executes, grabbing the global counter value into its own y.

30

Data Race Example

```
static int cnt = 0;
t1.run() {
  int y = cnt;
  cnt = y + 1;
}
t2.run() {
  int y = cnt;
  cnt = y + 1;
}
```

Shared state cnt = 1

y = 0



T1 executes again, storing its value of y + 1 into the counter.

31

Data Race Example

```
static int cnt = 0;
t1.run() {
  int y = cnt;
  cnt = y + 1;
}
t2.run() {
  int y = cnt;
  cnt = y + 1;
}
```

Shared state cnt = 1

y = 0



y = 1

T1 finishes. T2 executes, grabbing the global counter value into its own y.

32

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```

Shared state **cnt = 2**

y = 0

```
t2.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```



y = 1

T2 executes, storing its incremented cnt value into the global counter.

33

But When it's Run Again?

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```

Shared state **cnt = 0**

```
t2.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```

Start: both threads ready to run. Each will increment the global count.

35

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```

Shared state **cnt = 0**

y = 0

```
t2.run() {  
    int y = cnt;  
    cnt = y + 1;  
}
```



T1 executes, grabbing the global counter value into its own y.

36

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```

Shared state cnt = 0

y = 0

```
t2.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```



y = 0

T1 is preempted. T2 executes, grabbing the global counter value into its own y.

37

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```

Shared state cnt = 1

y = 0

```
t2.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```



y = 0

T2 executes, storing the incremented cnt value.

38

Data Race Example

```
static int cnt = 0;
```

```
t1.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```

Shared state cnt = 1

y = 0

```
t2.run() {  
  int y = cnt;  
  cnt = y + 1;  
}
```



y = 0

T2 completes. T1 executes again, storing the incremented original counter value (1) rather than what the incremented updated value would have been (2)!

39

What Happened?

- Different schedules led to different outcomes
 - This is a *data race* or *race condition*
- A thread was preempted in the middle of an operation
 - Reading and writing `cnt` was supposed to be *atomic* to happen with no interference from other threads
 - But the schedule (interleaving of threads) which was chosen allowed atomicity to be violated
 - These bugs can be extremely hard to reproduce, and so hard to debug
 - Depends on what scheduler chose to do, which is hard to predict

40

Question

- If instead of

```
int y = cnt;
cnt = y+1;
```
- We had written

```
- cnt++;
```
- Would the result be any different?
- Answer: NO!
 - Don't depend on your intuition about atomicity

41

Question

- If you run a program with a race condition, will you always get an unexpected result?
 - No! It depends on the scheduler, i.e., which JVM you're running, and on the other threads/processes/etc, that are running on the same CPU
- Race conditions are hard to find

42

Synchronization

- Refers to mechanisms allowing a programmer to control the execution order of some operations across different threads in a concurrent program.
- Different languages have adopted different mechanisms to allow the programmer to synchronize threads.
- Java has several mechanisms; we'll look at locks first.

43

Locks (Java 1.5)

```
interface Lock {
    void lock();
    void unlock();
    ... /* Some more stuff, also */
}
class ReentrantLock implements Lock { ... }
```

- Only one thread can hold a lock at once
 - Other threads that try to acquire it *block* (or become suspended) until the lock becomes available
- *Reentrant lock* can be reacquired by same thread
 - As many times as desired
 - No other thread may acquire a lock until has been released same number of times it has been acquired

44

Avoiding Interference: Synchronization

```
public class Example extends Thread {
    private static int cnt = 0;
    static Lock lock = new ReentrantLock();
    public void run() {
        lock.lock();
        int y = cnt;
        cnt = y + 1;
        lock.unlock();
    }
    ...
}
```

Lock, for protecting the shared state

Acquires the lock; Only succeeds if not held by another thread

Releases the lock

45

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

Shared state cnt = 0



T1 acquires the lock

46

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

Shared state cnt = 0

y = 0



T1 reads cnt into y

47

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

Shared state cnt = 0

y = 0



T1 is preempted.
T2 attempts to acquire the lock but fails because it's held by T1, so it blocks

48

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

y = 0

Shared state cnt = 1



T1 runs, assigning
to cnt

49

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

y = 0

Shared state cnt = 1



T1 releases the lock
and terminates

50

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

y = 0

Shared state cnt = 1



T2 now can acquire
the lock.

51

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

y = 0

Shared state cnt = 1



T2 reads cnt into y.

y = 1

52

Applying Synchronization

```
int cnt = 0;
t1.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
t2.run() {
    lock.lock();
    int y = cnt;
    cnt = y + 1;
    lock.unlock();
}
```

Shared state **cnt = 2**

y = 0



T2 assigns cnt,
then releases the lock

y = 1

53

Different Locks Don't Interact

```
static int cnt = 0;
static Lock l =
    new ReentrantLock();
static Lock m =
    new ReentrantLock();

void inc() {
    l.lock();
    cnt++;
    l.unlock();
}
```

```
void inc() {
    m.lock();
    cnt++;
    m.unlock();
}
```

- This program has a race condition
 - Threads only block if they try to acquire a lock held by another thread

54

What's Wrong with the Following?

```
static int cnt = 0;
static int x = 0;
```

```
Thread 1
while (x != 0);
x = 1;
cnt++;
x = 0;
```

```
Thread 2
while (x != 0);
x = 1;
cnt++;
x = 0;
```

- Threads may be interrupted after the **while** but before the assignment **x = 1**
 - Both may think they “hold” the lock!
- This is *busy waiting*
 - Consumes lots of processor cycles

55

Reentrant Lock Example

```
static int cnt = 0;
static Lock l =
    new ReentrantLock();

void inc() {
    l.lock();
    cnt++;
    l.unlock();
}
```

```
void returnAndInc() {
    int temp;

    l.lock();
    temp = cnt;
    inc();
    l.unlock();
}
```

- Reentrancy is useful because each method can acquire/release locks as necessary
 - No need to worry about whether callers have locks
 - Discourages complicated coding practices

56

Deadlock

- *Deadlock* occurs when no thread can run because all threads are waiting for a lock
 - No thread running, so no thread can ever release a lock to enable another thread to run

```
Lock l = new ReentrantLock();
Lock m = new ReentrantLock();
```

Thread 1

```
l.lock();
m.lock();
...
m.unlock();
l.unlock();
```

Thread 2

```
m.lock();
l.lock();
...
l.unlock();
m.unlock();
```


57

Deadlock (cont'd)


- Some schedules work fine
 - Thread 1 runs to completion, then thread 2
- But what if...
 - Thread 1 acquires lock **l**
 - The scheduler switches to thread 2
 - Thread 2 acquires lock **m**
- **Deadlock!**
 - Thread 1 is trying to acquire **m**
 - Thread 2 is trying to acquire **l**
 - And neither can, because the other thread has it

58

Wait Graphs



Thread T1 holds lock **l**




Thread T2 attempting to acquire lock **m**

Deadlock occurs when there is a cycle in the graph

59

Wait Graph Example



T1 holds lock on **l**

T2 holds lock on **m**

T1 is trying to acquire a lock on **m**

T2 is trying to acquire a lock on **l**

60

Another Case of Deadlock

```
static Lock l = new ReentrantLock();

void f () throws Exception {
    l.lock();
    FileInputStream f =
        new FileInputStream("file.txt");
    // Do something with f
    f.close();
    l.unlock();
}
```

- `l` not released if exception thrown
 - Likely to cause deadlock some time later

61

Solution: Use Finally

```
static Lock l = new ReentrantLock();

void f () throws Exception {
    l.lock();
    try {
        FileInputStream f =
            new FileInputStream("file.txt");
        // Do something with f
        f.close();
    }
    finally {
        // This code executed no matter how we
        // exit the try block
        l.unlock();
    }
}
```

62

Synchronized

- This pattern is really common
 - Acquire lock, do something, release lock under any circumstances after we're done
 - Even if exception was raised etc.
- Java has a language construct for this
 - `synchronized (obj) { body }`
 - Every Java object has an implicit associated lock
 - Obtains the lock associated with `obj`
 - Executes `body`
 - Release lock when scope is exited
 - Even in cases of exception or method return

63

Example


```
static Object o = new Object();

void f() throws Exception {
    synchronized (o) {
        FileInputStream f =
            new FileInputStream("file.txt");
        // Do something with f
        f.close();
    }
}
```

- Lock associated with `o` acquired before body executed
 - Released even if exception thrown

64

Discussion



- An object and its associated lock are different!
 - Holding the lock on an object does not affect what you can do with that object in any way
 - Ex:

```
synchronized(o) { ... } // acquires lock named o
o.f ();                // someone else can call o's methods
o.x = 3;               // someone else can read and write o's fields
```

65

Example: Synchronizing on this

```
class C {
    int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }
}
```

```
C c = new C();
```

```
Thread 1
c.inc();
```

```
Thread 2
c.inc();
```

- Does this program have a data race?
 - No, both threads acquire locks on the same object before they access shared data

66

Example: Synchronizing on this (cont'd)

```
class C {
    int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }

    void dec() {
        synchronized (this) {
            cnt--;
        }
    }
}
```

```
C c = new C();
```

```
Thread 1
c.inc();
```

```
Thread 2
c.dec();
```

- Data race?
 - No, threads acquire locks on the same object before they access shared data

67

Example: Synchronizing on this (cont'd)

```
class C {
    int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }
}
```

```
C c1 = new C();
C c2 = new C();
```

```
Thread 1
c1.inc();
```

```
Thread 2
c2.inc();
```

- Does this program have a data race?
 - No, threads acquire different locks, but they write to different objects, so that's ok

68

Synchronized Methods

- Marking method as synchronized same as synchronizing on this in body of the method
 - The following two programs are the same

```
class C {
    int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }
}
```

```
class C {
    int cnt;

    synchronized void inc() {
        cnt++;
    }
}
```

69

Synchronized Methods (cont'd)

```
class C {
    int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }

    synchronized void dec() {
        cnt--;
    }
}
```

```
C c = new C();
```

```
Thread 1
c.inc();
```

```
Thread 2
c.dec();
```

- Data race?
 - No, both acquire same lock

70

Synchronized Static Methods

- Warning: Static methods lock class object
 - There's no `this` object to lock

```
class C {
    static int cnt;

    void inc() {
        synchronized (this) {
            cnt++;
        }
    }

    static synchronized void dec() {
        cnt--;
    }
}
```

```
C c = new C();
```

```
Thread 1
c.inc();
```

```
Thread 2
C.dec();
```

71

Thread Scheduling

- When multiple threads share a CPU...
 - When should the current thread stop running?
 - What thread should run next?
- A thread can voluntarily `yield()` the CPU
 - Call to `yield` may be ignored; don't depend on it
- *Preemptive schedulers* can de-schedule the current thread at any time
 - Not all JVMs use preemptive scheduling, so a thread stuck in a loop may *never* yield by itself. Therefore, put `yield()` into loops
- Threads are de-scheduled whenever they block (e.g., on a lock or on I/O) or go to sleep

72

Thread Lifecycle

- While a thread executes, it goes through a number of different phases
 - **New**: created but not yet started
 - **Runnable**: is running, or can run on a free CPU
 - **Blocked**: waiting for I/O or on a lock
 - **Sleeping**: paused for a user-specified interval
 - **Terminated**: completed

73

Which Thread to Run Next?

- Look at all runnable threads
 - A good choice to run is one that just became unblocked because
 - A lock was released
 - I/O became available
 - It finished sleeping, etc.
- Pick a thread and start running it
 - Can try to influence this with `setPriority(int)`
 - Higher-priority threads get preference
 - But you probably don't need to do this

74

Some Thread Methods

- `void join()` throws `InterruptedException`
 - Waits for a thread to die/finish
- `static void yield()`
 - Current thread gives up the CPU
- `static void sleep(long milliseconds)` throws `InterruptedException`
 - Current thread sleeps for the given time
- `static Thread currentThread()`
 - Get `Thread` object for currently executing thread

75

Example: Threaded, Sync Alarm

```
while (true) {
    System.out.print("Alarm> ");

    // read user input
    String line = b.readLine();
    parseInput(line);

    // wait (in secs) asynchronously
    if (m != null) {
        // start alarm thread
        Thread t = new AlarmThread(m, tm);
        t.start();
        // wait for the thread to complete
        t.join();
    }
}
```

76

Daemon Threads

- `void setDaemon(boolean on)`
 - Marks thread as a daemon thread
 - Must be set before thread started
- By default, thread acquires status of thread that spawned it
- Program execution terminates when no threads running except daemons

77

Key Ideas

- Multiple threads can run simultaneously
 - Either truly in parallel on a multiprocessor
 - Or can be scheduled on a single processor
 - A running thread can be pre-empted at any time
- Threads can share data
 - In Java, only fields can be shared
 - Need to prevent interference
 - Rule of thumb 1: You must hold a lock when accessing shared data
 - Rule of thumb 2: You must not release a lock until shared data is in a valid state
 - Overuse of synchronization can create deadlock
 - Rule of thumb: No deadlock if only one lock

78

Producer/Consumer Design


- Suppose we are communicating with a shared variable
 - E.g., some kind of a buffer holding messages
- One thread *produces* input to the buffer
- One thread *consumes* data from the buffer
- How do we implement this?
 - Use condition variables

79

Conditions (Java 1.5)

```
interface Lock { Condition newCondition(); ... }  
interface Condition {  
    void await();  
    void signalAll(); ... }
```

- **Condition** created from a **Lock**
- **await** called with lock held
 - Releases the lock
 - But not any other locks held by this thread
 - Adds this thread to wait set for lock
 - Blocks the thread
- **signalAll** called with lock held
 - Resumes all threads on lock's wait set
 - Those threads must reacquire lock before continuing
 - (This is part of the function; you don't need to do it explicitly)



80

Producer/Consumer Example

```
Lock lock = new ReentrantLock();
Condition ready = lock.newCondition();
boolean valueReady = false;
Object value;
```

```
void produce(Object o) {      Object consume() {
    lock.lock();              lock.lock();
    while (valueReady)        while (!valueReady)
        ready.await();        ready.await();
    value = o;                Object o = value;
    valueReady = true;        valueReady = false;
    ready.signalAll();        ready.signalAll();
    lock.unlock();            lock.unlock();
}                              }
```

81

Use This Design

- This is the right solution to the problem
 - Tempting to try to just use locks directly
 - Very hard to get right
 - Problems with other approaches often very subtle
 - E.g., double-checked locking is broken

82

Broken Producer/Consumer Example

```
Lock lock = new ReentrantLock();
boolean valueReady = false;
Object value;
```

```
void produce(object o) {      Object consume() {
    lock.lock();              lock.lock();
    while (valueReady);        while (!valueReady);
    value = o;                Object o = value;
    valueReady = true;        valueReady = false;
    lock.unlock();            lock.unlock();
}                              }
```

Threads wait with lock held – no way to make progress

83

Broken Producer/Consumer Example

```
Lock lock = new ReentrantLock();
boolean valueReady = false;
Object value;
```

```
void produce(object o) {      Object consume() {
    while (valueReady);        while (!valueReady);
    lock.lock();              lock.lock();
    value = o;                Object o = value;
    valueReady = true;        valueReady = false;
    lock.unlock();            lock.unlock();
}                              }
```

valueReady accessed without a lock held – race condition

84

Broken Producer/Consumer Example

```
Lock lock = new ReentrantLock();
Condition ready = lock.newCondition();
boolean valueReady = false;
Object value;
```

```
void produce(object o) {      Object consume() {
    lock.lock();              lock.lock();
    if (valueReady)           if (!valueReady)
        ready.await();        ready.await();
    value = o;                Object o = value;
    valueReady = true;        valueReady = false;
    ready.signalAll();        ready.signalAll();
    lock.unlock();            lock.unlock();
}                               }
```

what if there are multiple producers or consumers?

85

More on the Condition Interface

```
interface Condition {
    void await();
    boolean await (long time, TimeUnit unit);
    void signal();
    void signalAll();
    ... }
```

- `await(t, u)` waits for time `t` and then gives up
 - Result indicates whether woken by signal or timeout
- `signal()` wakes up only *one* waiting thread
 - Tricky to use correctly
 - Have all waiters be equal, handle exceptions correctly
 - Highly recommended to just use `signalAll()`

86

Await and SignalAll Gotcha's

- `await` *must* be in a loop
 - Don't assume that when wait returns conditions are met
- Avoid holding other locks when waiting
 - `await` only gives up locks on the object you wait on

87

Blocking Queues in Java 1.5

- Interface for producer/consumer pattern

```
interface Queue<E> extends Collection<E> {
    boolean offer(E x); /* produce */
    /* waits for queue to have capacity */


    E remove(); /* consume */
    /* waits for queue to become non-empty */
    ... }
```

- Two handy implementations
 - `LinkedBlockingQueue` (FIFO, may be bounded)
 - `ArrayBlockingQueue` (FIFO, bounded)
 - (plus a couple more)

88

Wait and NotifyAll (Java 1.4)

- Recall that in Java 1.4, use synchronize on object to get associated lock



- Objects also have an associated wait set

89

Wait and NotifyAll (cont'd)

- o.wait()**
 - Must hold lock associated with **o**
 - Release that lock
 - And no other locks
 - Adds this thread to wait set for lock
 - Blocks the thread
- o.notifyAll()**
 - Must hold lock associated with **o**
 - Resumes all threads on lock's wait set
 - Those threads must reacquire lock before continuing
 - (This is part of the function; you don't need to do it explicitly)

90

Producer/Consumer in Java 1.4

```
public class ProducerConsumer {
    private boolean valueReady = false;
    private Object value;

    synchronized void produce(Object o) {
        while (valueReady) wait();
        value = o; valueReady = true;
        notifyAll();
    }

    synchronized Object consume() {
        while (!valueReady) wait();
        valueReady = false;
        Object o = value;
        notifyAll();
        return o;
    }
}
```

91

InterruptedException

- Exception thrown if interrupted on certain ops
 - wait**, **await**, **sleep**, **join**, and **lockInterruptibly**
 - Also thrown if call one of these with interrupt flag set
- Not thrown* when blocked on 1.4 lock or I/O

```
class Object {
    void wait() throws IE;
    ... }
interface Lock {
    void lock();
    void lockInterruptibly() throws IE;
    ... }
interface Condition {
    void await() throws IE;
    void signalAll();
    ... }
```

92