

# Multithreading

---

# Chapter Goals

- To understand how multiple threads can execute in parallel
- To learn how to implement threads
- To understand race conditions and deadlocks
- To be able to avoid corruption of shared objects by using locks and conditions
- To be able to use threads for programming animations

# Threads

- **Thread:** a program unit that is executed independently of other parts of the program
- The Java Virtual Machine executes each thread in the program for a short amount of time
- This gives the impression of parallel execution

# Running a Thread

- Implement a class that implements the `Runnable` interface:

```
public interface Runnable
{
    void run();
}
```

- Place the code for your task into the `run` method of your class:

```
public class MyRunnable implements Runnable
{
    public void run()
    {
        Task statements
        ...
    }
}
```

# Running a Thread

- Create an object of your subclass:

```
Runnable r = new MyRunnable();
```

- Construct a `Thread` object from the `Runnable` object:

```
Thread t = new Thread(r);
```

- Call the `start` method to start the thread:

```
t.start();
```

# Example

A program to print a time stamp and “Hello World” once a second for ten seconds:

```
Mon Dec 28 23:12:03 PST 2009 Hello, World!  
Mon Dec 28 23:12:04 PST 2009 Hello, World!  
Mon Dec 28 23:12:05 PST 2009 Hello, World!  
Mon Dec 28 23:12:06 PST 2009 Hello, World!  
Mon Dec 28 23:12:07 PST 2009 Hello, World!  
Mon Dec 28 23:12:08 PST 2009 Hello, World!  
Mon Dec 28 23:12:09 PST 2009 Hello, World!  
Mon Dec 28 23:12:10 PST 2009 Hello, World!  
Mon Dec 28 23:12:11 PST 2009 Hello, World!  
Mon Dec 28 23:12:12 PST 2009 Hello, World!
```

# GreetingRunnable Outline

```
public class GreetingRunnable implements Runnable
{
    private String greeting;

    public GreetingRunnable(String aGreeting)
    {
        greeting = aGreeting;
    }

    public void run()
    {
        Task statements
        ...
    }
}
```

## **Thread Action for `GreetingRunnable`**

- Print a time stamp
- Print the greeting
- Wait a second



## GreetingRunnable

- We can get the date and time by constructing a `Date` object:

```
Date now = new Date();
```

- To wait a second, use the `sleep` method of the `Thread` class:

```
sleep(milliseconds)
```

- A sleeping thread can generate an `InterruptedException`
  - *Catch the exception*
  - *Terminate the thread*

# Running Threads

- `sleep` puts current thread to sleep for given number of milliseconds:

```
Thread.sleep(milliseconds)
```

- When a thread is interrupted, most common response is to terminate `run`

# Generic `run` method

```
public void run()  
{  
    try  
    {  
        Task statements  
    }  
    catch (InterruptedException exception)  
    {  
    }  
    Clean up, if necessary  
}
```

# ch20/greeting/GreetingRunnable.java

```
1  import java.util.Date;
2
3  /**
4   * A runnable that repeatedly prints a greeting.
5   */
6  public class GreetingRunnable implements Runnable
7  {
8     private static final int REPETITIONS = 10;
9     private static final int DELAY = 1000;
10
11    private String greeting;
12
13    /**
14     * Constructs the runnable object.
15     * @param aGreeting the greeting to display
16     */
17    public GreetingRunnable(String aGreeting)
18    {
19        greeting = aGreeting;
20    }
21
```

***Continued***

## ch20/greeting/GreetingRunnable.java (cont.)

```
22     public void run()
23     {
24         try
25         {
26             for (int i = 1; i <= REPETITIONS; i++)
27             {
28                 Date now = new Date();
29                 System.out.println(now + " " + greeting);
30                 Thread.sleep(DELAY);
31             }
32         }
33         catch (InterruptedException exception)
34         {
35         }
36     }
37 }
```

# To Start the Thread

- Construct an object of your `Runnable` class:

```
Runnable t = new GreetingRunnable("Hello World");
```

- Then construct a thread and call the `start` method:

```
Thread t = new Thread(r);  
t.start();
```

# ch20/greeting/GreetingThreadRunner.java

```
1  /**
2     This program runs two greeting threads in parallel.
3  */
4  public class GreetingThreadRunner
5  {
6     public static void main(String[] args)
7     {
8         GreetingRunnable r1 = new GreetingRunnable("Hello, World!");
9         GreetingRunnable r2 = new GreetingRunnable("Goodbye, World!");
10        Thread t1 = new Thread(r1);
11        Thread t2 = new Thread(r2);
12        t1.start();
13        t2.start();
14    }
15 }
```

***Continued***

# ch20/greeting/GreetingThreadRunner.java (cont.)

## Program Run:

```
Mon Dec 28 12:04:46 PST 2009 Hello, World!  
Mon Dec 28 12:04:46 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:47 PST 2009 Hello, World!  
Mon Dec 28 12:04:47 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:48 PST 2009 Hello, World!  
Mon Dec 28 12:04:48 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:49 PST 2009 Hello, World!  
Mon Dec 28 12:04:49 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:50 PST 2009 Hello, World!  
Mon Dec 28 12:04:50 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:51 PST 2009 Hello, World!  
Mon Dec 28 12:04:51 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:52 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:52 PST 2009 Hello, World!  
Mon Dec 28 12:04:53 PST 2009 Hello, World!  
Mon Dec 28 12:04:53 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:54 PST 2009 Hello, World!  
Mon Dec 28 12:04:54 PST 2009 Goodbye, World!  
Mon Dec 28 12:04:55 PST 2009 Hello, World!  
Mon Dec 28 12:04:55 PST 2009 Goodbye, World!
```



# Thread Scheduler

- **Thread scheduler:** runs each thread for a short amount of time (a **time slice**)
- Then the scheduler activates another thread
- There will always be slight variations in running times - especially when calling operating system services (e.g. input and output)
- There is no guarantee about the order in which threads are executed

## Self Check 20.1

What happens if you change the call to the `sleep` method in the `run` method to `Thread.sleep(1)`?

**Answer:** The messages are printed about one millisecond apart.

## Self Check 20.2

What would be the result of the program if the `main` method called

```
r1.run();  
r2.run();
```

instead of starting threads?

**Answer:** The first call to `run` would print ten “Hello” messages, and then the second call to `run` would print ten “Goodbye” messages

# Terminating Threads

- A thread terminates when its `run` method terminates
- Do not terminate a thread using the deprecated `stop` method
- Instead, notify a thread that it should terminate:

```
t.interrupt();
```

- `interrupt` does not cause the thread to terminate – it sets a boolean variable in the thread data structure

# Terminating Threads

- The `run` method should check occasionally whether it has been interrupted
  - *Use the `interrupted` method*
  - *An interrupted thread should release resources, clean up, and exit:*

```
public void run()
{
    for (int i = 1;
        i <= REPETITIONS && !Thread.interrupted();
        i++)
    {
        Do work
    }
    Clean up
}
```

# Terminating Threads

- The `sleep` method throws an `InterruptedException` when a sleeping thread is interrupted
  - *Catch the exception*
  - *Terminate the thread :*

```
public void run()
{
    try
    {
        for (int i = 1; i <= REPETITIONS; i++)
        {
            Do work
            Sleep
        }
    }
    catch (InterruptedException exception)
    {
        Clean up
    }
}
```

# Terminating Threads

- Java does not force a thread to terminate when it is interrupted
- It is entirely up to the thread what it does when it is interrupted
- Interrupting is a general mechanism for getting the thread's attention

## Self Check 20.4

Consider the following `Runnable`:

```
public class MyRunnable implements Runnable
{
    public void run()
    {
        try
        {
            System.out.println(1);
            Thread.sleep(1000);
            System.out.println(2);
        }
        catch (InterruptedException exception)
        {
            System.out.println(3);
        }
        System.out.println(4);
    }
}
```



## Self Check 20.4 (cont.)

Suppose a thread with this `Runnable` is started and immediately interrupted.

```
Thread t = new Thread(new MyRunnable());  
t.start();  
t.interrupt();
```

What output is produced?

**Answer:** The `run` method prints the values 1, 3, and 4. The call to `interrupt` merely sets the interruption flag, but the `sleep` method immediately throws an `InterruptedException`.

# Race Conditions

- When threads share a common object, they can conflict with each other
- Sample program: multiple threads manipulate a bank account

# ch20/unsynch/BankAccountThreadRunner.java

```
1  /**
2     This program runs threads that deposit and withdraw
3     money from the same bank account.
4  */
5  public class BankAccountThreadRunner
6  {
7     public static void main(String[] args)
8     {
9         BankAccount account = new BankAccount();
10        final double AMOUNT = 100;
11        final int REPETITIONS = 100;
12        final int THREADS = 100;
13
14        for (int i = 1; i <= THREADS; i++)
15        {
16            DepositRunnable d = new DepositRunnable(
17                account, AMOUNT, REPETITIONS);
18            WithdrawRunnable w = new WithdrawRunnable(
19                account, AMOUNT, REPETITIONS);
20
21            Thread dt = new Thread(d);
22            Thread wt = new Thread(w);
23
```

***Continued***

## ch20/unsynch/BankAccountThreadRunner.java (cont.)

```
24         dt.start();
25         wt.start();
26     }
27 }
28 }
```

# ch20/unsynch/DepositRunnable.java

```
1  /**
2     A deposit runnable makes periodic deposits to a bank account.
3  */
4  public class DepositRunnable implements Runnable
5  {
6     private static final int DELAY = 1;
7     private BankAccount account;
8     private double amount;
9     private int count;
10
11    /**
12     Constructs a deposit runnable.
13     @param anAccount the account into which to deposit money
14     @param anAmount the amount to deposit in each repetition
15     @param aCount the number of repetitions
16    */
17    public DepositRunnable(BankAccount anAccount, double anAmount,
18        int aCount)
19    {
20        account = anAccount;
21        amount = anAmount;
22        count = aCount;
23    }
24
```

***Continued***

## ch20/unsynch/DepositRunnable.java (cont.)

```
25     public void run()  
26     {  
27         try  
28         {  
29             for (int i = 1; i <= count; i++)  
30             {  
31                 account.deposit(amount);  
32                 Thread.sleep(DELAY);  
33             }  
34         }  
35         catch (InterruptedException exception) {}  
36     }  
37 }
```

# ch20/unsynch/WithdrawRunnable.java

```
1  /**
2     A withdraw runnable makes periodic withdrawals from a bank account.
3  */
4  public class WithdrawRunnable implements Runnable
5  {
6     private static final int DELAY = 1;
7     private BankAccount account;
8     private double amount;
9     private int count;
10
11    /**
12     Constructs a withdraw runnable.
13     @param anAccount the account from which to withdraw money
14     @param anAmount the amount to withdraw in each repetition
15     @param aCount the number of repetitions
16    */
17    public WithdrawRunnable(BankAccount anAccount, double anAmount,
18        int aCount)
19    {
20        account = anAccount;
21        amount = anAmount;
22        count = aCount;
23    }
24
```

***Continued***

## ch20/unsynch/WithdrawRunnable.java (cont.)

```
25     public void run()
26     {
27         try
28         {
29             for (int i = 1; i <= count; i++)
30             {
31                 account.withdraw(amount);
32                 Thread.sleep(DELAY);
33             }
34         }
35         catch (InterruptedException exception) {}
36     }
37 }
```



# ch20/unsynch/BankAccount.java

```
1  /**
2     A bank account has a balance that can be changed by
3     deposits and withdrawals.
4  */
5  public class BankAccount
6  {
7     private double balance;
8
9     /**
10     Constructs a bank account with a zero balance.
11     */
12     public BankAccount()
13     {
14         balance = 0;
15     }
16
```

***Continued***

## ch20/unsynch/BankAccount.java (cont.)

```
17     /**
18         Deposits money into the bank account.
19         @param amount the amount to deposit
20     */
21     public void deposit(double amount)
22     {
23         System.out.print("Depositing " + amount);
24         double newBalance = balance + amount;
25         System.out.println(", new balance is " + newBalance);
26         balance = newBalance;
27     }
28
```

*Continued*

## ch20/unsynch/BankAccount.java (cont.)

```
29     /**
30         Withdraws money from the bank account.
31         @param amount the amount to withdraw
32     */
33     public void withdraw(double amount)
34     {
35         System.out.print("Withdrawing " + amount);
36         double newBalance = balance - amount;
37         System.out.println(", new balance is " + newBalance);
38         balance = newBalance;
39     }
40
41     /**
42         Gets the current balance of the bank account.
43         @return the current balance
44     */
45     public double getBalance()
46     {
47         return balance;
48     }
49 }
```

## ch20/unsynch/BankAccount.java (cont.)

### Program Run:

```
Depositing 100.0, new balance is 100.0
Withdrawing 100.0, new balance is 0.0
Depositing 100.0, new balance is 100.0
Withdrawing 100.0, new balance is 0.0
...
Withdrawing 100.0, new balance is 400.0
Depositing 100.0, new balance is 500.0
Withdrawing 100.0, new balance is 400.0
Withdrawing 100.0, new balance is 300.0
```

# Sample Application

- Create a `BankAccount` object
- Create two sets of threads:
  - *Each thread in the first set repeatedly deposits \$100*
  - *Each thread in the second set repeatedly withdraws \$100*
- `deposit` and `withdraw` have been modified to print messages:

```
public void deposit(double amount)
{
    System.out.print("Depositing " + amount);
    double newBalance = balance + amount;
    System.out.println(", new balance is "
        + newBalance);
    balance = newBalance;
}
```

# Sample Application

- The result should be zero, but sometimes it is not
- Normally, the program output looks somewhat like this:

```
Depositing 100.0, new balance is 100.0
Withdrawing 100.0, new balance is 0.0
Depositing 100.0, new balance is 100.0
Depositing 100.0, new balance is 200.0
Withdrawing 100.0, new balance is 100.0
...
Withdrawing 100.0, new balance is 0.0
```

- But sometimes you may notice messed-up output, like this:

```
Depositing 100.0Withdrawing 100.0, new balance is
100.0, new balance is -100.0
```

# Scenario to Explain Non- zero Result: Race Condition

1. A deposit thread executes the lines:

```
System.out.print("Depositing " + amount);  
double newBalance = balance + amount;
```

The `balance` variable is still 0, and the `newBalance` local variable is 100

2. The deposit thread reaches the end of its time slice and a withdraw thread gains control

3. The withdraw thread calls the `withdraw` method which withdraws \$100 from the `balance` variable; it is now -100

4. The withdraw thread goes to sleep

*Continued*

# Scenario to Explain Non- zero Result: Race Condition

5. The deposit thread regains control and picks up where it left off; it executes:

```
System.out.println(", new balance is " + newBalance);  
balance = newBalance;
```

The balance is now 100 instead of 0 because the deposit method used the OLD balance



# Corrupting the Contents of the balance Variable

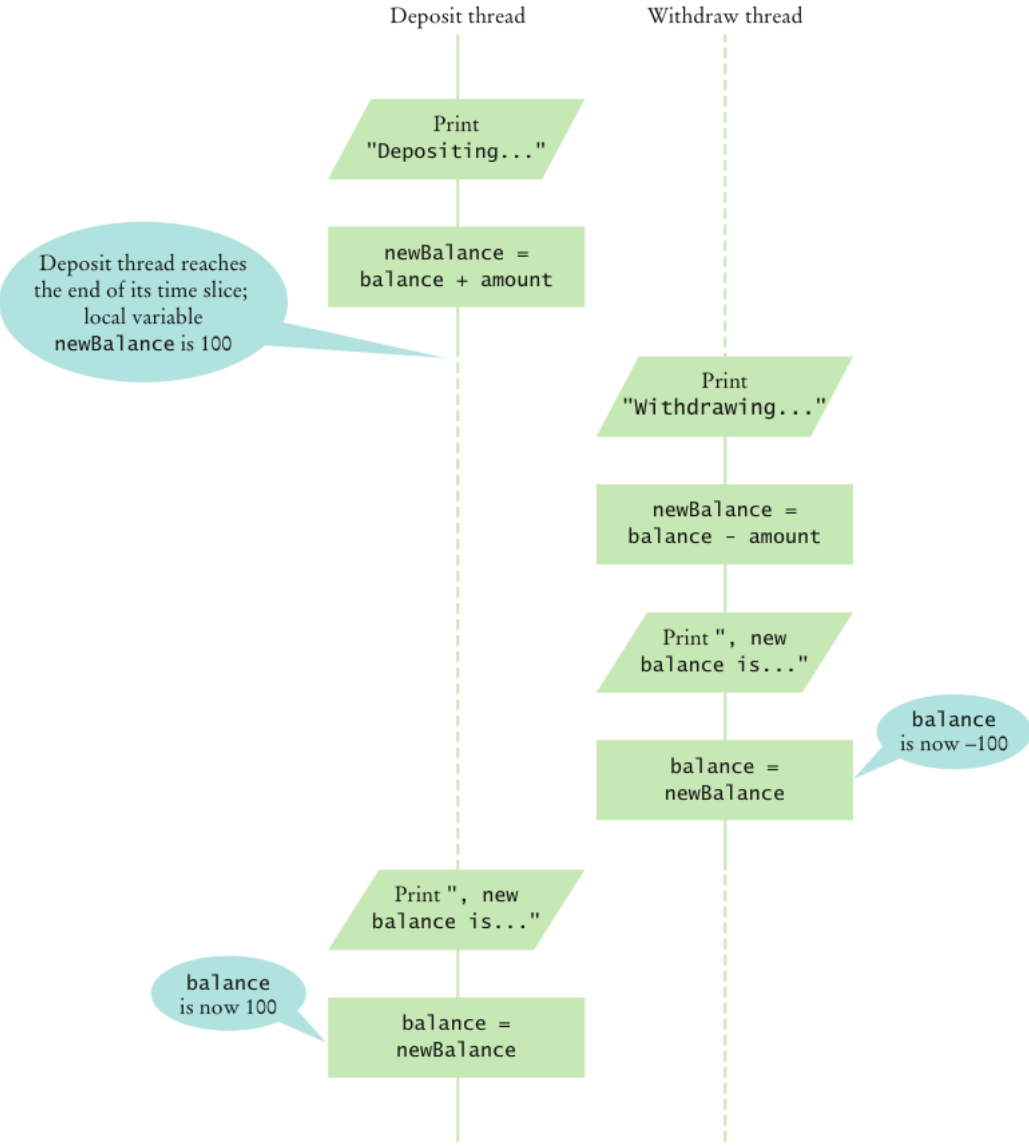


Figure 1 Corrupting the Contents of the balance Variable

# Race Condition

- Occurs if the effect of multiple threads on shared data depends on the order in which they are scheduled
- It is possible for a thread to reach the end of its time slice in the middle of a statement
- It may evaluate the right-hand side of an equation but not be able to store the result until its next turn:

```
public void deposit(double amount)
{
    balance = balance + amount;
    System.out.print("Depositing " + amount
        + ", new balance is " + balance);
}
```

- Race condition can still occur:

```
balance = the right-hand-side value
```

## Self Check 20.5

Give a scenario in which a race condition causes the bank balance to be -100 after one iteration of a deposit thread and a withdraw thread.

**Answer:** There are many possible scenarios. Here is one:

- *The first thread loses control after the first `print` statement.*
- *The second thread loses control just before the assignment `balance = newBalance`.*
- *The first thread completes the `deposit` method.*
- *The second thread completes the `withdraw` method.*

# Synchronizing Object Access

- To solve problems such as the one just seen, use a *lock object*
- **Lock object:** used to control threads that manipulate shared resources
- In Java: `Lock` interface and several classes that implement it
  - *ReentrantLock*: most commonly used lock class
  - Locks are a feature of Java version 5.0
  - Earlier versions of Java have a lower-level facility for thread synchronization

# Synchronizing Object Access

- Typically, a lock object is added to a class whose methods access shared resources, like this:

```
public class BankAccount
{
    private Lock balanceChangeLock;

    public BankAccount()
    {
        balanceChangeLock = new ReentrantLock();
        ...
    }
    ...
}
```

# Synchronizing Object Access

- Code that manipulates shared resource is surrounded by calls to `lock` and `unlock`:

```
balanceChangeLock.lock();
```

Manipulate the shared resource

```
balanceChangeLock.unlock();
```

- If code between calls to `lock` and `unlock` throws an exception, call to `unlock` never happens

# Synchronizing Object Access

- To overcome this problem, place call to `unlock` into a `finally` clause:

```
public void deposit(double amount)
{
    balanceChangeLock.lock();
    try
    {
        System.out.print("Depositing " + amount);
        double newBalance = balance + amount;
        System.out.println(", new balance is " +
            newBalance);    balance = newBalance;
    }
    finally
    {
        balanceChangeLock.unlock();
    }
}
```

# Synchronizing Object Access

- When a thread calls `lock`, it owns the lock until it calls `unlock`
- A thread that calls `lock` while another thread owns the lock is temporarily deactivated
- Thread scheduler periodically reactivates thread so it can try to acquire the lock
- Eventually, waiting thread can acquire the lock



## Self Check 20.7

If you construct two `BankAccount` objects, how many lock objects are created?

**Answer:** Two, one for each bank account object. Each lock protects a separate balance variable.

## Self Check 20.8

What happens if we omit the call `unlock` at the end of the `deposit` method?

**Answer:** When a thread calls `deposit`, it continues to own the lock, and any other thread trying to deposit or withdraw money in the same bank account is blocked forever.

# Avoiding Deadlocks

- A deadlock occurs if no thread can proceed because each thread is waiting for another to do some work first
- BankAccount example:

```
public void withdraw(double amount)
{
    balanceChangeLock.lock();
    try
    {
        while (balance < amount)
            Wait for the balance to grow
            ...
    }
    finally
    {
        balanceChangeLock.unlock();
    }
}
```

# Avoiding Deadlocks

- How can we wait for the balance to grow?
- We can't simply call `sleep` inside `withdraw` method; thread will block all other threads that want to use `balanceChangeLock`
- In particular, no other thread can successfully execute `deposit`
- Other threads will call `deposit`, but will be blocked until `withdraw` exits
- But `withdraw` doesn't exit until it has funds available
- DEADLOCK

# Condition Objects

- To overcome problem, use a condition object
- Condition objects allow a thread to temporarily release a lock, and to regain the lock at a later time
- Each condition object belongs to a specific lock object

## Condition Objects (cont.)

- You obtain a condition object with `newCondition` method of `Lock` interface:

```
public class BankAccount
{
    public BankAccount()
    {
        balanceChangeLock = new ReentrantLock();
        sufficientFundsCondition =
            balanceChangeLock.newCondition();
        ...
    }
    ...
    private Lock balanceChangeLock;
    private Condition sufficientFundsCondition;
}
```

# Condition Objects

- It is customary to give the condition object a name that describes condition to test
- You need to implement an appropriate test

## Condition Objects (cont.)

- As long as test is not fulfilled, call `await` on the condition object:

```
public void withdraw(double amount)
{
    balanceChangeLock.lock();
    try
    {
        while (balance < amount)
            sufficientFundsCondition.await();
        ...
    }
    finally
    {
        balanceChangeLock.unlock();
    }
}
```



# Condition Objects

- Calling `await`
  - *Makes current thread wait*
  - *Allows another thread to acquire the lock object*
- To unblock, another thread must execute `signalAll` *on the same condition object* :

```
sufficientFundsCondition.signalAll();
```

- `signalAll` unblocks all threads waiting on the condition
- `signal`: randomly picks just one thread waiting on the object and unblocks it
- `signal` can be more efficient, but you need to know that every waiting thread can proceed
- Recommendation: always call `signalAll`

# ch20/synch/BankAccountThreadRunner.java

```
1  /**
2     This program runs threads that deposit and withdraw
3     money from the same bank account.
4  */
5  public class BankAccountThreadRunner
6  {
7     public static void main(String[] args)
8     {
9         BankAccount account = new BankAccount();
10        final double AMOUNT = 100;
11        final int REPETITIONS = 100;
12        final int THREADS = 100;
13
14        for (int i = 1; i <= THREADS; i++)
15        {
16            DepositRunnable d = new DepositRunnable(
17                account, AMOUNT, REPETITIONS);
18            WithdrawRunnable w = new WithdrawRunnable(
19                account, AMOUNT, REPETITIONS);
20
21            Thread dt = new Thread(d);
22            Thread wt = new Thread(w);
23
```

***Continued***

## ch20/synch/BankAccountThreadRunner.java (cont.)

```
24         dt.start();
25         wt.start();
26     }
27 }
28 }
```

# ch20/synch/BankAccount.java

```
1  import java.util.concurrent.locks.Condition;
2  import java.util.concurrent.locks.Lock;
3  import java.util.concurrent.locks.ReentrantLock;
4
5  /**
6   * A bank account has a balance that can be changed by
7   * deposits and withdrawals.
8   */
9  public class BankAccount
10 {
11     private double balance;
12     private Lock balanceChangeLock;
13     private Condition sufficientFundsCondition;
14
15     /**
16      * Constructs a bank account with a zero balance.
17      */
18     public BankAccount()
19     {
20         balance = 0;
21         balanceChangeLock = new ReentrantLock();
22         sufficientFundsCondition = balanceChangeLock.newCondition();
23     }
24
```

## ch20/synch/BankAccount.java (cont.)

```
25     /**
26         Deposits money into the bank account.
27         @param amount the amount to deposit
28     */
29     public void deposit(double amount)
30     {
31         balanceChangeLock.lock();
32         try
33         {
34             System.out.print("Depositing " + amount);
35             double newBalance = balance + amount;
36             System.out.println(", new balance is " + newBalance);
37             balance = newBalance;
38             sufficientFundsCondition.signalAll();
39         }
40         finally
41         {
42             balanceChangeLock.unlock();
43         }
44     }
45
```

***Continued***

## ch20/synch/BankAccount.java (cont.)

```
46     /**
47         Withdraws money from the bank account.
48         @param amount the amount to withdraw
49     */
50     public void withdraw(double amount)
51         throws InterruptedException
52     {
53         balanceChangeLock.lock();
54         try
55         {
56             while (balance < amount)
57                 sufficientFundsCondition.await();
58             System.out.print("Withdrawing " + amount);
59             double newBalance = balance - amount;
60             System.out.println(", new balance is " + newBalance);
61             balance = newBalance;
62         }
63         finally
64         {
65             balanceChangeLock.unlock();
66         }
67     }
68
```

***Continued***

## ch20/synch/BankAccount.java (cont.)

```
69     /**
70         Gets the current balance of the bank account.
71         @return the current balance
72     */
73     public double getBalance()
74     {
75         return balance;
76     }
77 }
```

*Continued*

## ch20/synch/BankAccount.java (cont.)

### Program Run:

```
Depositing 100.0, new balance is 100.0
```

```
Withdrawing 100.0, new balance is 0.0
```

```
Depositing 100.0, new balance is 100.0
```

```
Depositing 100.0, new balance is 200.0
```

```
...
```

```
Withdrawing 100.0, new balance is 100.0
```

```
Depositing 100.0, new balance is 200.0
```

```
Withdrawing 100.0, new balance is 100.0
```

```
Withdrawing 100.0, new balance is 0.0
```



## Self Check 20.9

What is the essential difference between calling `sleep` and `await`?

**Answer:** A sleeping thread is reactivated when the sleep delay has passed. A waiting thread is only reactivated if another thread has called `signalAll` or `signal`.

## Self Check 20.10

Why is the `sufficientFundsCondition` object an instance variable of the `BankAccount` class and not a local variable of the `withdraw` and `deposit` methods?

**Answer:** The calls to `await` and `signal/signalAll` must be made *to the same object*.