


## A Brief Introduction to Object-Oriented Programming Using Java

Java's basic language syntax is similar to C/C++, but there are many differences too. Java is **strongly typed** like C++, meaning that a Java compiler detects type inconsistencies at compile-time. A Java compiler translates Java source code into Java **bytecodes**, which are instructions to a Java **virtual machine** (Java VM).

A Java VM executes a Java program by *interpreting* the byte codes as virtual machine instructions, which in turn cause the VM to execute real machine instructions. Interestingly, the java interpreter is mostly written in Java! So, when you execute the interpreter, you're running a java program. There is a small run-time part that is written in C, and the rest of it is implemented in Java. For example, the Java compiler is written in Java! This is also the case with other interpreted programming languages, such as Smalltalk, Lisp, Scheme, Dylan, ML, and others.



Java bytecodes are machine independent. This means that you can compile Java source code on one hardware platform and transfer the bytecodes to another platform and execute them using a Java VM that has been ported to that platform. This machine independence of the bytecode instructions is what makes Java suitable for “programming the Internet.” Recently, **native-code** Java compilers have been developed for performance reasons.

## Packages, Classes, Fields and Methods

Java requires that for every class, instance variable, and method you define, that you declare an access qualifier as part of the definition, or accept the default access (which is not private access) Compare this to what you must write in C++. A public Java class must correspond to a file of the same name. So, the Hello class must be defined in a Hello.java file, which is part of the 'test' package. Technically, the main method is called:

```
edu.utexas.cs.example.Hello.main() // a fully qualified method name
```

NOTE: Unlike C++ which uses the '::' operator for scope resolution, Java uses the '.' operator for both scope resolution and class member access.

A *fully qualified name* is of the form <package>.<class>.<field or method>. You might organize your packages into subpackages, in which case you write <package>.<subpackage>.<class>.<field or method>. A fully qualified name corresponds to a directory path in the filesystem, which is how Java is able to guarantee unique naming. So the name test.Hello might correspond to a file:

```
/home/user/edu/utexas/cs/example/Hello.java
```

A single .java file may contain multiple class definitions for a package, but only one may be public. The name of the public class must correspond to the name of the file. The java compiler will compile a .java file with multiple classes into multiple .class files: one for each class.

The CLASSPATH environment variable is used to tell the java interpreter where to look for packages and classes. Rather than always using a fully qualified name all the time, you can "open" a package using the **import** statement. When you specify an import statement, it opens the package, and allows you to omit the package name qualifiers. A package name may be a compound name: For example, the following input/output (io) subpackage is part of the standard Java SDK package:

```
import java.io.*;
```

The import statement causes the java.applet package to be opened, and the '\*' says to import all classes from that package.

## A Simple Java Program

A Java application program consists of one or more classes. One class should have a “main” method, which is the entry point for the Java program. Unlike C/C++, there are no global functions or procedures. All functions/procedures are defined on classes or objects. Note the similarities and differences with C++.

```
package edu.utexas.cs.example; // define class as part of your organization's test package
import java.io.*; // import the java I/O package
public class Hello {
    private static String s = new String("Hello, world");

    public static void main(String[] args) { // note the difference in the array arg
        System.out.println(s);
    }
}
```

1. The **package** statement defines a *namespace*, called edu.utexas.cs.example. The public class Hello belongs to this namespace. The package statement must be the first Java statement in a file. If not specified, the class is put in the default “global package”. Note that package names usually correspond to Internet domain names, for uniqueness.
2. There are no #include directives for including other class definitions, because there is no preprocessor in Java. The **import** statement imports a set of class definitions from another package. In this case, the java.io package.
3. A class may be either public or private, with respect to a package. A public class is visible to users of a package containing the class. A private class is only visible to other classes defined as part of the same package. There are no public/private sections as defined in C++. Instead, each instance variable and method is defined to be **public**, **private**, **protected**, or it has **default access**. Default access is like protected, but restricts access to only those classes in the same package. Default access is also sometimes called “friendly” access.
4. A variable or method declared static is global to all instances of the class, just as in C++. Static methods have no implicit **this reference**, which is how each runtime object is associated with its methods and local data. Note that when you declare a variable, you should initialize it as the same time as it is declared. All objects are created with the **new** operator and the result is assigned to an **object reference variable**. *There are no pointers in Java!*

## Java Data Types

Java has three categories of data types: **primitive types**, **arrays**, and **reference types**.

Because java programs execute on a virtual machine, the primitive types are defined to be the same across all real machine architectures. Note that this is different from C/C++, where the real machine architecture has an impact on the precision. Q: What is the max n for fact(n) in C/C++?

Variables defined as a primitive **boolean** type default to **false** on declaration within a class (i.e., heap allocated), but must be explicitly initialized within methods (i.e., stack allocated):

boolean	true or false
---------	---------------

Variables defined as one of the following primitive integral types default to a value of zero on declaration with a class (i.e., heap allocated), but must be explicitly initialized within methods (i.e., stack allocated):

byte	-128..127
short	-32768..32767
int	-2147483648..2147483647
long	-922337203685477508..9223372036854775807
char	0..65535

## Arrays

Arrays are created by calling `new`, which returns a reference to the array. Arrays are automatically garbage collected when no longer referenced in any scope.

```
byte buffer[] = new byte[1024];
```

Each byte in the array 'buffer' is initialized to its default value, which is zero for integral types.

```
Button button[] = new Button[10];
```

An array of type `Button` results in an array of `Button` references, all of which are initialized to null.

If you attempt to reference an element of an array, and invoke an operation on a `button` object, a `NullPointerException` will be generated.

```
button[i].push();
```

So, you must explicitly initialize each array reference with a new object.

```
for (int i = 0; i < button.length; i++)  
    button[i] = new Button();
```

Note that every array has a builtin public "constant" `length` field, which is a read-only variable that returns the length of the array as an integer. The length is set to the extent of the array when the array is first created by a call to `new`.

## Arrays

You also can initialize an array with a static initializer:

```
String commands[] = { "File", "Edit", "Format", "Help" };
```

Which is equivalent to:

```
String[] commands = { new String("File"), new String("Edit"),  
                      new String("Format"), new String("Help") };
```

Note that in Java, you can place the '[]' syntax, indicating an array type, next to the typename T instead of next to the variable name. This syntax is perhaps more intuitive since what you are declaring is an array of some type T. The syntax of using the [] after the variable name in an array declaration is in the language as a courtesy to long-time C/C++ programmers.

For example:

```
String[] s; // declare an uninitialized reference to an array of Strings  
s = new String[100];  
s = new String[10];
```

The second time the reference variable 's' is initialized, the previous array becomes garbage and is subject to garbage collection on the next sweep of the garbage collector.

## Some Simple Facts about Strings

Strings have somewhat special treatment in Java because they are so often used. You can write the following:

```
String s = "Hello, world";
```

instead of

```
String s = new String("Hello, world");
```

The Java compiler turns all string literals, like "Hello, world", into heap allocated, garbage collected, String objects automatically.

Once a String object is initialized, it is treated as a constant string value.

Unlike C++, user-defined operator overloading is not allowed in Java. However, Java does provide all the standard C operators, and includes a few builtin overloads. For example, the '+' operator is overloaded for String objects, so that you can perform concatenation.

```
System.out.println("hello," + " world");
```

Many methods are defined in Java classes that take String objects as arguments and return String objects as results. It is very common for objects to define a **toString()** method, which implements a conversion to a String object. If an object is used where a String is expected, the Java run-time automatically invoked the toString method to implicitly convert the object to a String.

## What does it mean to program with Java?

Java is an object-oriented programming language. This means that computation is achieved via **interacting objects**. Objects interact by invoking operations on one another. Operations effect local state changes on objects, hopefully leading to some useful computation, the results of which are typically presented as output to some device.

The operations associated with a particular object must be well defined via the object's **interface**. The interface effectively defines a *contract* between the user of the object (i.e., some other object in the program) and the implementer of the object (i.e., the programmer). So, there are two perspectives to every object in an object-oriented programmer: The “user view” and the “implementer view.”

The interface, or contract, for a particular category of objects is defined using the **class** syntactic construct. A Java source program consists entirely of a set of classes that completely specify a set of object interfaces and their implementation at compile-time. A Java run-time program consists of a set of compiled bytecodes representing dynamic instantiations of those classes in the form of run-time objects that execute some computation when the bytecodes are interpreted by a Java virtual machine.

The task of a Java programmer is to define new classes in terms of existing classes and builtin types. This consists of writing **class definitions**, which define new types of objects, the specification of variables related to the implementation of the public methods that provide the interface for an object, and possibly other variables and methods that are internal (private) to the implementation.

So, in Java, a program consists of lots of different class definitions. Some have to be written from scratch, others are simply reused by importing their definitions from pre-defined **packages** of class definitions (i.e., libraries). Designing a computation as a collection of compile-time class definitions each representing a category of objects that interact at run-time to perform computation is called object-oriented programming.

## How do we program in Java?

We take the “user view” and study existing packages of class definitions to learn how to interact with pre-defined categories of objects. We also take the “implementor view” and write new packages of class definitions that define new types of objects and specifies how they interact with other objects.

A **Java program** is then constructed by instantiating objects and causing them to invoke operations on one another to effect a useful computation.

Since classes are the key to understanding pre-defined object categories and for writing new ones, we have to master the structure, syntax, and object-oriented programming conventions used by Java classes.

### Java class structure

<b>data</b>	instance variables and class variables
<b>operations</b>	object constructors object methods class methods
<b>optional things</b>	optional class variable initialization block
	optional nested class definition(s)

## Objects

All objects by default inherit from a base object class called `java.lang.Object`. So, in Java all objects are subtypes of the type `Object`. That is, anywhere a type `Object` is called for (e.g., in a type signature of some procedure), ANY object can be used since all objects are of type `Object`.

This means that objects of any type can all be treated generically as instance of class `Object`, which is useful for defining heterogeneous collections of objects, such as a vector, stack, etc. In fact, the `java.util.vector` and `java.util.stack` classes are both defined to operate on objects of type `Object`.

Java objects are allocated from the heap and a garbage collector takes care of reaping objects from the heap when there are no more references to the object. When you create a new object, the result of the `new` operator is a reference to the heap allocated object that is assigned to an object reference variable. You use this object reference variable to access the public fields and methods of an object.

The `Object` class provides some high-level methods that can be called on an object of any type. For example, the `Object` class defines a `clone` method that makes a copy of an object:

You can copy objects using a special method called `java.lang.Object.clone()`, which is defined for all objects since all objects inherit from the `java.lang.Object`.

```
Vector v = new Vector();  
Vector z = v;  
Vector c = z.clone();
```

The `java.lang.Object.clone()` method does a field by field copy of the object, or what is called a “shallow” copy. It does not do a “deep” copy.

## The “this” Reference

Every object in Java is uniquely identified by its **this** reference. The **this** reference is the “handle” (pointer) to the object that is used to invoke methods or explicitly reference instance variables.

When we create an object by calling `new`, a reference variable is initialized with the value of the object’s location in memory. When we use a reference variable to invoke an operation on an object, we are really invoking a method that is associated with some class, and we dynamically associate the state of some particular object with the method.

```
String s = new String("This is a string");  
int x = s.length(); // compute length of string
```

The compiler turns the invocation of `s.length` into the following procedure call:

```
int x = String.length(this=s);
```

So, the `length` method of the `String` class is invoked, and it will have a **this** reference that is initialized to the memory address of the object pointed to by `s`. Conceptually, what you have to keep in mind in an object oriented language like Java or C++, is that an object at run-time is just a chunk of memory somewhere in the heap that contains the data elements of the object, along with some bookkeeping information generated by the compiler for the object. At run-time, the object consists of data in the heap, and a single copy of the methods for all objects of a particular class are separate.

When you invoke a method on an object, the method and the object’s state are dynamically “tied together” for the duration of the method’s execution. The “**this**” reference is the mechanism that is used to tie them together at run-time. The **this** reference as a hidden first argument to every non-static class method. Static “class” methods do not have a **this** reference passed to them.

## A typical Java class schemata

```
package package-name ;
public class class-name {
    // instance variable declarations, with optional inline initialization
    private type-name variable-name ;
    private type-name variable-name = initialization-expr ;

    // class variable declaration of a constant with inline initialization
    private final static type-name variable-name = initialization-expr ;

    // object constructors
    public class-name () { stmt_1 ; ... ; stmt_n }
    public class-name ( formal-parameters ) { stmt_1 ; ... ; stmt_n ; }

    // object method definition
    public return-type method-name ( formal-parameters ) {
        stmt_1 ; ... ; stmt_n ;
    }
    // class method definition
    public static return-type method-name ( formal-parameters ) {
        stmt_1 ; ... ; stmt_n ;
    }
}
```

Type-names and return-types are taken from the set of builtin types, and previously defined class-names. A special return-type is **void** (the “un-type”). Formal parameters are a (possibly empty) comma-separated-list of name pairs: *type-name* *variable-name*.

## A fragment of a typical class definition

```
public class String {
    // private instance variables, each object has its own copy
    private char value[];
    private int offset = 0;
    private int count;

    // private class variable, shared by all String objects
    private static final long serialVersionUID = -6849794470754667710L;

    // constructors
    public String() { value = new char[0]; }

    public String(char value[]) {
        this.count = value.length;
        this.value = new char[count];
        System.arraycopy(value, 0, this.value, 0, count);
    }

    // public object method
    public int length() { return count; }

    // public class method
    public static String valueOf(int i) { return Integer.toString(i, 10); }
    ...
}
```

## Use of 'this' in methods

```
public class StringBuffer {
    private char value[];
    private int count;
    private boolean shared;

    public StringBuffer() { this(16); }
    public StringBuffer(int length) {
        this.value = new char[length];
        this.shared = false;
    }
    public StringBuffer(String str) {
        this(str.length() + 16);
        append(str);
    }
    public StringBuffer append(String str) {
        if (str == null) str = String.valueOf(str);
        int len = str.length();
        ensureCapacity(this.count + len);
        copyWhenShared();
        str.getChars(0, len, this.value, this.count);
        this.count += len;
        return this;
    }
}
```

**NOTE: this is not defined in a class method (i.e., a method declared static)**

## Reference variables and object instantions

Objects are used in Java by declaring and using reference variables that are initialized using constructors. Constructors are special class methods that do not return any value.

```
class-name variable-name ;  
class-name variable-name = null ;  
class-name variable-name = new class-name ( optional-argument-list ) ;  
class-name variable-name = initialization-expression ;
```

Examples:

```
String a ; // uninitialized String reference variable  
String b = null ; // explicitly setting a null object reference  
String c = new String() ; // the "empty" string object  
String d = new String ("This is a string literal") ;  
String e = s ; // initialize e using String object s ;  
a = e ; // initialize a to refer to the same String object as e
```

Strings have a special kind of initialization expression, involving string literals.

```
String f = "this is a string literal" ;  
String g = "Hello" + ", " + "world" ;
```

These are equivalent to the initialization expressions:

```
String f = new String ("this is a string literal") ;  
String g = new String ("Hello, world") ;
```

## Object and class method invocations

Once an object of a class has been instantiated and a reference variable to that object is initialized, the object is typically used in some expression as the target for some kind of method invocation using the reference variable.

```
reference-var . method-name ( optional-argument-list );
```

For example:

```
String hello = new String("Hello, world");  
int x = hello.length();
```

Class methods, which are globally accessible operations defined on classes not on objects, are invoked similarly, but using a class name instead of a reference variable.

```
class-name . class-method-name ( optional-arguments );
```

For example, an initialization expression can invoke a method that constructs an object of the appropriate type to perform the initialization.

```
String five = String.valueOf(5); // String class method to "stringify" an int
```

The `String.valueOf(int)` method is a class method that constructs a new `String` object by converting an integer value to a `String` value, assuming a base 10 radix.

## Some special class definition forms

```
public class class-name {  
  
    // publicly accessible class variables  
    public static type-name variable-name = initialization-expr;  
    public static type-name variable-name;  
  
    // a class variable initialization block  
    static {  
        stmt_1; ...; stmt_n;  
    }  
  
    // special "main" class method used as a program entry point  
    public static void main (String[] args) {  
        stmt_1; ...; stmt_n;  
    }  
}
```

Class variables are very often declared public and are referenced using an expression of the form:

```
class-name . class-variable-name
```

For example, the class `java.lang.Math` class defines a public class variable representing a 64-bit approximation of  $\pi$  defined as `PI = 3.14159265358979323846` and used as follows:

```
double area = Math.PI * radius * radius;
```

## Static Data Initialization

Static data is usually initialized at the point of declaration within a class:

```
class Program {
    static String errorMessage = "fatal error";
    static int errorCode = -1;
    ...
}
```

Alternatively, you can use a **static block** when the initialization is more complicated.

```
class Program {
    static String[] errMessages;
    static int errorCode;

    static {
        errorCode = -1;
        errMessages = new String[10];
        errMessages[0] = "fatal error";
        ...
    }
    ...
}
```

A static block is NOT the same thing as a constructor. Static variables are initialized when either a static method first is called on the class that contains the static data, or when the first object of the class is first instantiated.

## Using class variables and methods

Class variables and methods can be used before any objects are ever instantiated in a Java program. The keyword **static** is used to denote class variables and methods because they are considered to be “statically” instantiated attributes of the class, in contrast to objects, which are all dynamically instantiated.

```
class Hello {
    public static String username = null;

    static {
        try {
            username = System.getProperty("user.name");
        } catch (Exception e) { System.err.println(e); }
    }

    public static void main (String[] args) {
        System.out.println("Username is: " + username);
    }
}
```

A static initialization block is typically used to initialize class variables that require more than a simple inline initialization expression, or simply to reduce “declaration clutter” and collect together all the static initializations in one place. The special class method **main** is used as the starting point for a program. It is possible for multiple classes in the same Java program to have distinct main methods. The main method called to start a Java program is determined by which class the Java VM is asked to load first.

## More special class definition forms

```
public final class class-name {  
  
    // public, globally accessible constant declaration  
    public final static type-name variable-name = initialization-expr;  
  
    // a "final" object method, which cannot be overridden by a subclass  
    public final return-type method-name ( formal-parameters ) {  
        stmt_1; ...; stmt_n;  
    }  
  
    // a "native" object method, which has an external implementation  
    public native return-type method-name ( formal-parameters );  
    public final native return-type method-name ( formal-parameters );  
  
}
```

A final class cannot be extended. That is to say, a new class cannot be defined as an extension of a final class.

A final static variable is the same thing as a constant. A final method is one that cannot be overridden by a subclass. In a final class, all methods are implicitly final. In a non-final class, methods can be selectively made final. Only object methods can be final. Static methods are implicitly final.

A native method consists of just a method declaration with no method body. The method implementation is typically written in C/C++ and is linked into the Java run-time, usually for efficiency reasons.

## Example usage of “final” and “native” qualifiers

```
public final class Math {  
  
    public static final double E = 2.7182818284590452354;  
    public static final double PI = 3.14159265358979323846;  
  
    public static native double sin (double a);  
    public static native double cos (double a);  
    public static native double tan (double a);  
    ...  
    public static native double exp(double a);  
    public static native double log(double a);  
    public static native double sqrt(double a);  
    public static native double pow(double a, double b);  
    ...  
    public static int abs(int a) { return (a < 0) ? -a : a; }  
    public static long abs(long a) { return (a < 0) ? -a : a; }  
    public static float abs(float a) { return (a < 0) ? -a : a; }  
    public static double abs(double a) { return (a < 0) ? -a : a; }  
    ...  
    public static native double random();  
}
```

Note that it is necessary to **overload** the absolute value method for each the builtin numeric types. In Java, you can define many methods with the same name, as long as the type signatures are unique.

## Declaration of methods that throw exceptions

```
public class class-name {  
    ...  
    // a "final" object method, which cannot be overridden by a subclass  
    public return-type method-name ( formal-parameters )  
    throws exception-type-list  
    {  
        stmt_1; ...; stmt_n;  
    }  
};
```

An exception-type-list is a comma-separated list of previously defined exception types (i.e., class names). In general, you should invoke any method that will throw an exception inside the context of a **try block**, but not necessarily in the same scope level. A catch block is defined for each specific exception type that you want to catch. An optional finally block can be defined to contain a set of statements that will be executed whether or not the exception occurs.

```
try {  
    ...; variable-name.method-name (arguments); ...;  
}  
catch (exception-type1 e1) { stmt_1; ...; stmt_n; }  
catch (exception-type2 e2) { stmt_1; ...; stmt_n; }  
...  
finally {  
    stmt_1; ...; stmt_n;  
}
```

## Example: a method that throws an exception

```
public final class Byte {
    public static final byte    MIN_VALUE = -128;
    public static final byte    MAX_VALUE = 127;

    public static byte parseByte(String s) throws NumberFormatException {
        return parseByte(s, 10);
    }
    public static byte parseByte(String s, int radix)
throws NumberFormatException {
        int i = Integer.parseInt(s, radix);
        if (i < MIN_VALUE || i > MAX_VALUE)
            throw new NumberFormatException();
        return (byte)i;
    }
}

public static void main (String[] args) {
    int err_code = 0;
    try {
        byte b = Byte.parseByte(args[0]);
    }
    catch (ArrayIndexOutOfBoundsException e) { err_code = -1; }
    catch (NumberFormatException e) { err_code = -2; }
    finally { System.exit(err_code); }
}
```

## Inheritance in Java

Inheritance is a compile-time mechanism in Java that allows you to extend a class (called the **base class** or **superclass**) with another class (called the **derived class** or **subclass**). In Java, inheritance is used for two purposes:

1. **class inheritance** - create a new class as an extension of another class, primarily for the purpose of **code reuse**. That is, the **derived class** inherits the public methods and public data of the **base class**. Java only allows a class to have one immediate base class, i.e., single class inheritance.
2. **interface inheritance** - create a new class to implement the methods defined as part of an **interface** for the purpose of **subtyping**. That is a class that implements an interface “conforms to” (or is constrained by the type of) the interface. Java supports multiple interface inheritance.

In Java, these two kinds of inheritance are made distinct by using different language syntax. For class inheritance, Java uses the keyword **extends** and for interface inheritance Java uses the keyword **implements**.

```
public class derived-class-name extends base-class-name {  
    // derived class methods extend and possibly override  
    // those of the base class  
}
```

```
public class class-name implements interface-name {  
    // class provides an implementation for the methods  
    // as specified by the interface  
}
```

## Composition vs Inheritance

A fundamental design choice that we have to make when considering whether or not to use inheritance is to examine the relationship that exists between two classes. The simple way to do this is the ‘Is-a’ and ‘has-a’ rule. We say an object X is-a Y, if everywhere you can use an object of type Y, you can use instead of object of type X. In other words, X is a **proper subtype** of Y. Usually, this means that X implements the same interface as Y. So, you know that X and Y conform to the same set of method type signatures, however, their implementation may be different. We say an object X has-a Y, if Y is a part-of X. So, you typically think of X containing an instance of Y, not X inheriting from Y. For example:

1. Would you say that a Stack is-a Vector or a Stack has-Vector?
2. Would you say that Circle is-a Shape or a Circle has-a Shape?

In each case, you need to consider whether or not the relationship between two objects is simple one of “using the object” or “being the object”. If you just need to use an object, then that implies a composition relationship. If an object behaves like another object, then that implies a subtype relationship.

Subtyping, and subtype polymorphism, are one of the most important contributions of object-oriented programming to programming language theory in general. Java implements a very general form of subtype polymorphism, because every object is-a Object, since every class inherits implicitly from the class Object. You can always treat any object in Java, not matter what its class type is, as a subtype of Object. However, this is a very generic way to treat ALL objects in a program, which may in fact be too general, as it forces the programmer to do a lot of **upcasting** and **downcasting**. You cast up from a concrete type to a more general type and you cast down from a more general type of a more concrete type. Improper downcasting can however result in run-time type errors.

## Example of class inheritance

```
package MyPackage;

class Base {
    private int x;
    public int f() { ... }
    protected int g() { ... }
}

class Derived extends Base {
    private int y;
    public int f() { /* new implementation for Base.f() */ }
    public void h() { y = g(); ... }
}
```

In Java, the **protected** access qualifier means that the protected item (field or method) is visible to a any derived class of the base class containing the protected item. It also means that the protected item is visible to methods of other classes in the same package. This is different from C++.

Q: What is the base class of class Object? I.e., what would you expect to get if you executed the following code?

```
Object x = new Object();
System.out.println(x.getClass().getSuperclass());
```

## Order of Construction under Inheritance

Note that when you construct an object, the default base class constructor is called implicitly, before the body of the derived class constructor is executed. So, objects are constructed top-down under inheritance. Since every object inherits from the `Object` class, the `Object()` constructor is always called implicitly. However, you can call a superclass constructor explicitly using the builtin **super** keyword, as long as it is the *first* statement in a constructor.

For example, most Java exception objects inherit from the `java.lang.Exception` class. If you wrote your own exception class, say `SomeException`, you might write it as follows:

```
public class SomeException extends Exception {

    public SomeException() {
        super(); // calls Exception(), which ultimately calls Object()
    }

    public SomeException(String s) {
        super(s); // calls Exception(String), to pass argument to base class
    }

    public SomeException (int error_code) {
        this("error"); // class constructor above, which calls super(s)
        System.err.println(error_code);
    }
}
```

## Abstract Base Classes

An **abstract class** is a class that leaves one or more method implementations unspecified by declaring one or more methods *abstract*. An abstract method has no body (i.e., no implementation). A subclass is required to override the abstract method and provide an implementation. Hence, an abstract class is incomplete and cannot be instantiated, but can be used as a base class.

```
abstract public class abstract-base-class-name {  
    // abstract class has at least one abstract method  
  
    public abstract return-type abstract-method-name ( formal-params );  
  
    ... // other abstract methods, object methods, class methods  
}  
  
public class derived-class-name extends abstract-base-class-name {  
  
    public return-type abstract-method-name (formal-params) { stmt-list; }  
  
    ... // other method implementations  
}
```

It would be an error to try to instantiate an object of an abstract type:

```
abstract-class-name obj = new abstract-class-name(); // ERROR!
```

That is, operator `new` is invalid when applied to an abstract class. Why?

## Example abstract class usage

```
abstract class Point {
    private int x, y;
    public Point(int x, int y) { this.x = x; this.y = y; }
    public void move(int dx, int dy)
    { x += dx; y += dy; plot(); }
    public abstract void plot(); // has no implementation
}

abstract class ColoredPoint extends Point {
    private int color;
    protected public ColoredPoint(int x, int y, int color)
    { super(x, y); this.color = color; }
}

class SimpleColoredPoint extends ColoredPoint {
    public SimpleColoredPoint(int x, int y, int color) { super(x,y,color); }
    public void plot() { ... } // code to plot a SimpleColoredPoint
}
```

Since ColoredPoint does not provide an implementation of the plot method, it must be declared abstract. The SimpleColoredPoint class does implement the inherited plot method. It would be an error to try to instantiate a Point object or a ColoredPoint object. However, you can declare a Point reference and initialize it with an instance of a subclass object that implements the plot method:

```
Point p = new SimpleColoredPoint(a, b, red); p.plot();
```

## Interfaces

An abstract class mixes the idea of mutable data in the form of instance variables, non-abstract methods, and abstract methods. An abstract class with only static final instance variables and all abstract methods is called an **interface**. An interface is a *specification*, or contract, for a set of methods that a class that implements the interface must conform to in terms of the type signature of the methods. The class that implements the interface provides an implementation for each method, just as with an abstract method in an abstract class.

So, you can think of an interface as an abstract class with all abstract methods. The interface itself can have either public, package, private or protected access defined. All methods declared in an interface are implicitly abstract and implicitly public. It is not necessary, and in fact considered redundant to declare a method in an interface to be abstract.

You can define data in an interface, but it is less common to do so. If there are data fields defined in an interface, then they are implicitly defined to be:

- public.
- static, and
- final

In other words, any data defined in an interface are treated as public constants.

Note that a class and an interface in the same package cannot share the same name.

Methods declared in an interace cannot be declared final. **Why?**

## Interface declaration

Interface names and class names in the same package must be distinct.

```
public interface interface-name {  
  
    // if any data are defined, they must be constants  
    public static final type-name var-name = constant-expr;  
  
    // one or more implicitly abstract and public methods  
    return-type method-name ( formal-params );  
}
```

An interface declaration is nothing more than a specification to which some class that purports to implement the interface must conform to in its implementation. That is, a class that implements the interface must have defined implementations for each of the interface methods.

The major reason for interfaces being distinct elements in Java is that you often want to define some operation to operate on objects that all conform to the same interface. So, you can define some code in a very general way, with a guarantee that by only using the methods defined in the interface, that all objects that implement the interface will have defined implementations for all the methods.

For example, you might define a Shape interface that defines an area() method, and so you would expect that any class that implements the Shape interface, would define an area method. So, if I have a list of references to objects that implement the Shape interface, I can legitimately invoke the area method, abstractly, on each of these objects and expect to obtain as a result a value that represents the area of some shape.

## Separation of Interface from Implementations

Interfaces are specifications for many possible implementations. Interfaces are used to define a contract for how you interact with an object, independent of the underlying implementation. The objective of an object-oriented programmer is to separate the specification of the interface from the hidden details of the implementation.

Consider the specification of a common LIFO stack.

```
public interface StackInterface {  
    boolean empty();  
    void push(Object x);  
    Object pop() throws EmptyStackException;  
    Object peek() throws EmptyStackException;  
}
```

Note that the methods in this interface are defined to operate on objects of type `Object`. Since a stack is a “container type”, it is very common to use the base class for all objects, namely `Object`, as the type of the arguments and results to a container type. Since all objects ultimately inherit from class `Object`, we can always generically refer to any object in the system using an `Object` reference. However, when we pop from a stack, we have to explicitly type case from the very general type `Object` to a concrete type, so that we can manipulate the object concretely.

**Q: How many different ways are there to implement a stack? If we are using just using a `Stack` object (as opposed to implementing it ourselves) should we care?**

## Stack implementation of the StackInterface

```
public class Stack implements StackInterface {  
  
    private Vector v = new Vector(); // use java.util.Vector class  
  
    public boolean empty() { return v.size() == 0; }  
  
    public void push(Object item) { v.addElement(item); }  
  
    public Object pop() {  
        Object obj = peek();  
        v.removeElementAt(v.size() - 1);  
        return obj;  
    }  
  
    public Object peek() throws EmptyStackException {  
        if (v.size() == 0)  
            throw new EmptyStackException();  
        return v.elementAt(v.size() - 1);  
    }  
}
```

## Should a stack use or inherit from a vector?

The `java.util.Stack` class is defined as a subclass of the `java.util.Vector` class, rather than using a `Vector` object as in the previous example. This sort of inheritance is not subtype inheritance, because the interface of a `Stack` object can be violated because a `Vector` has a “wider” interface than a `Stack`, i.e., a vector allows insertion into the front and the rear, so it is possible to violate the stack contract by treating a stack object as a vector, and violating the LIFO specification of a stack.

```
public class Stack extends Vector {
    public Object push(Object item) {addElement(item); return item;}
    public Object pop() {
        Object obj;
        int len = size();
        obj = peek();
        removeElementAt(len - 1);
        return obj;
    }
    public Object peek() {
        int len = size();
        if (len == 0) throw new EmptyStackException();
        return elementAt(len - 1);
    }
    public boolean empty() { return size() == 0;}
}
```

```
Vector v = new Stack(); // legal - base class reference to subclass object
v.insertElementAt(x, 2); // insert object x into Vector slot 2!!
```

## When to use an Interface vs when to use an abstract class

Having reviewed their basic properties, there are two primary differences between interfaces and abstract classes:

- an abstract class can have a mix of abstract and non-abstract methods, so some default implementations can be defined in the abstract base class. An abstract class can also have static methods, static data, private and protected methods, etc. In other words, a class is a class, so it can contain features inherent to a class. The downside to an abstract base class, is that since there is only single inheritance in Java, you can only inherit from one class.
- an interface has a very restricted use, namely, to declare a set of public abstract method signatures that a subclass is required to implement. An interface defines a set of type constraints, in the form of type signatures, that impose a requirement on a subclass to implement the methods of the interface. Since you can inherit multiple interfaces, they are often a very useful mechanism to allow a class to have different behaviors in different situations of usage by implementing multiple interfaces.

It is usually a good idea to implement an interface when you need to define methods that are to be explicitly overridden by some subclass. If you then want some of the methods implemented with default implementations that will be inherited by a subclass, then create an implementation class for the interface, and have other class inherit (extend) that class, or just use an abstract base class instead of an interface.

## Example of default interface implementations

```
interface X {
    void f();
    int g();
}

class XImpl implements X {
    void g() { return -1; } // default implementation for g()
}

class Y extends XImpl implements X {
    void f() { ... } // provide implementation for f()
}
```

Note that when you invoke an abstract method using a reference of the type of an abstract class or an interface, the method call is dynamically dispatched.

```
X x = new Y();
x.f();
```

The call `x.f()` causes a run-time determination of the actual type of object that 'x' refers to, then a method lookup to determine which implementation of `f()` to invoke. In this case, `Y.f(x)` is called, but the type of `x` is first converted to `Y` so that the 'this' reference is initialized to a reference of type `Y` inside of `f()`, since the implicit type signature of `Y.f()` is really `Y.f(final Y this)`;

## Subtype Polymorphism

The Java class `Object` class is the base class for every object. This idea of a common base class `Object` is borrowed from Smalltalk. The `Object` class permits a form of **subtype polymorphism** in Java that is not found in C++. Since every object “is-a” `Object`, then it is possible to define heterogeneous collections of objects. For example, the `java.util.Vector` class, is implemented as a resizable array of type `Object`. That means that you can insert any type of java object into a single `Vector`. However, when you do this, you “lose” the type of the object.

For example:

```
Vector v = new Vector();
v.addElement(new String("hello, world"));
v.addElement(new Integer(5));
Object s = v.firstElement();
Object i = v.lastElement();
```

In this simple example, you have to know at compile-time that ‘s’ is really a reference to an object of type `String`, and ‘i’ is a reference to an element of type `Integer`. What you really want to write is an explicit downcast:

```
String s = (String) v.firstElement();
Integer i = (Integer) v.lastElement();
```

Polymorphism in Java permits this type of explicit **downcasting** of an `Object` reference to a reference of the object’s real type. If you get it wrong, the java run-time will raise an exception called the `ClassCastException`. This is a case where you should use the **instanceof** language feature to check the real type of a subtype of the `Object` class at runtime before performing the type cast if you do not know for sure that you are casting to the correct type.

## Finding the Class of an Object at Runtime

Because every instance of a class is also an instance of the Object class, you can use an Object reference to determine at run-time the class of an object. In fact, Java defines a class called “java.lang.Class” that allows you to do some interesting things that are not easily achievable in C++. For example, you can query an object to determine its type:

```
Object x = vector.removeLast();
Class type = x.getClass(); // get the class of the thing that was in the vector
System.out.println(type.getName());
```

You can then ask the Class object for certain kinds of information about the class for an object, such as its superclass (there’s only one superclass in java since there is only single inheritance) and the names of all the interfaces implemented by the class.

```
Class superclass = type.getSuperclass();
System.out.println(superclass.getName());
Class[] interfaces = type.getInterfaces();
for (int i = 0; i < interfaces.length; i++)
    System.out.println(interfaces[i].getName());
```

This type of flexibility is useful for obtaining type information at run-time, but it is not as useful as it could be, since you (the programmer) still have to ultimately know what type you are dealing with since you will need to eventually **downcast** an Object reference to a reference to the real type in order to use the type. The Object class mostly provides a convenient handle that allows you to treat all objects generically, for example, in a container type such as a Vector, Stack, Queue, etc.

## More on the class `java.lang.Class`

An interesting feature of the `Class` class is a static public method call `Class.forName` that allows you to construct a `Class` object given a `String`, then ask the class object to construct an object that is an instance of the class name given in the string. This is useful for constructing objects from user input. For example:

```
public static void main(String[] args)
{
    Object[] objects = new Object[args.length];
    try {
        for(int i = 0; i < args.length; i++) {
            Class type = Class.forName(args[i]); // construct Class using args
            objects[i] = type.newInstance();
        }
    }
    catch (Exception e) { System.err.println(e); System.exit(-1); }
}
```

Of course, you have to eventually cast each constructed object down to the type of the `Class`, which you need to know in advance when you write the program. So, Java provides some dynamic type programming features, but the strong typing isn't gone.

## Exceptions arising from previous example

The types of exceptions that can arise are:

1. **ClassNotFoundException** thrown by `Class.forName` because the String given does not match a known class.
2. **NoSuchMethodError** thrown by `Class.newInstance` because the Class does not define a constructor that takes no arguments. The `Class.newInstance` method constructs an object using the default constructor.
3. **IllegalAccessException** thrown by `Class.newInstance` because the default constructor is not accessible.
4. **InstantiationException** is thrown by `Class.newInstance` because the class is abstract or is an interface.