



Computer Science 523 Advanced Programming

The College of Saint Rose
Summer 2014

Topic Notes: Classes

So far, our example programs have operated on data (*i.e.*, variables and parameters) that we can categorize in just two ways:

- primitive types, such as `int`, `double`, `char`
- object types, such as `String`, `Scanner`, `Random`, `DecimalFormat`

We next focus on these object types a bit more. In particular, we will introduce our own object types into our programs.

Objects and Classes

Every object in a Java program is an entity that can contain both *fields* – which are like the variables we’ve been using in our programs, and *methods*, which operate on the data in those fields.

The idea of an object is central to the *object-oriented programming* paradigm, which has been very popular since being introduced a few decades ago.

The idea is that we write program components, called *classes*, which represent templates for the *objects* we wish to represent in our program. For each object, we include fields that are used to represent the state of the object and methods that allow that state to be queried or modified.

The text includes an example of an alarm clock. They came up with a list of fields that can be used to describe the state of an alarm clock. It is similar to this list:

- the current hour (0-23)
- the current minute (0-59)
- the current second (0-59)
- the alarm hour (0-23)
- the alarm minute (0-59)
- the alarm status (on or off)

And some methods that can be used to modify the state of the alarm clock:

- set current time
- set alarm time
- disable alarm
- enable alarm
- stop currently sounding alarm

We might also consider some methods to query the current state of the alarm:

- get current time
- get alarm time
- get alarm status (on or off)

And then the alarm clock might have some other things it does “on its own” – its state changes as the time proceeds:

- increment time by 1 second
- start sounding alarm

In Java, the functionality of an object is described in a *class*. We have been writing classes all semester as containers for our `main`, and recently, a few other methods. It turns out that this is just a small fraction of what a Java class can be used to do.

We look at a mechanism to achieve this by a simpler example. Consider this example, which we could have written much earlier in the semester:

See Example: `RatiosNoClass`

This program maintains information about ratios of integer values. We create two ratios, each represented by 2 `int` variables, and print them, modify them, and compute their decimal equivalents.

But with a class that represents a `Ratio` object, we can *encapsulate* the numbers (the numerator and denominator) into fields, and provide methods to construct (the *constructor(s)*), access (the *accessor methods*), and modify (the *mutator methods*) the fields.

See

See Example: `Ratios`

and the extensive comments within for details.

Object Orientation

The key to successful object-oriented programming is to identify and design the objects in the problem. Looking first at the ratios example, it is essential to notice that the `Ratio` class defines the data and operations related to a single ratio. It does not know or care how or why the ratios are being used. It simply defines a ratio and methods that operate on one. It is up to other code – in this case the `main` method in the `Ratios` class – to determine when and why `Ratio` objects are instantiated and how they are used.

The text shows a popular metaphor for a class definition: the cookie cutter. I like another: a file cabinet of forms and instructions. Think of the class as a master copy of a form that contains information about an object, how to create one, and instructions for how the object works. Each time we need a new instance of the object, we make a copy of that form and use it to keep track of the details of that object.

As a further example, let's consider a program intended to keep track of a series of items purchased in a store. Each item has a name, a unit price, and a quantity purchased by a customer. Our program will read in a series of these items and track the most expensive, least expensive, the largest quantity purchased, and the largest total cost (unit price times quantity). For simplicity, we will break any ties by the first encountered. So if the most expensive item cost \$9.50 but there are two different items each priced at \$9.50, we will keep the first and ignore the second.

We will make use of a class that represents one of these items. We will then be able to create an instance of this class for each such item. We can add accessor methods to retrieve the information we need to implement our program.

See Example: `PurchaseTracker`

One additional item here is the use of a *class variable* in the `PurchasedItem` class. The `DecimalFormat` object we declare and construct with the `static` qualifier is shared among all instances of the class, no matter how many we create.

General Purpose Classes

The few custom classes we have seen so far have been developed for very specific situations. The `Ratio` class stores a ratio, but isn't useful for much else. The `PurchasedItem` class in the purchase tracker example contains a `String`, a `double`, and an `int`, all used for very specific purposes.

Part of a good object-oriented design is to find places where we can write some software (in our case, a Java class) that might be useful in situations beyond the one at hand. It turns out there are many simple and no-so-simple structures that arise in many contexts, and rather than developing a new Java class each time we need one, we strive to re-use ones that already exist. Or if one does not exist, develop one that will satisfy needs of future programmers as well.

Let's consider one common *data structure*: the *pair*. To begin, let's assume that we want a pair of floating-point values. These might be used to represent a coordinate in two dimensions, a latitude and longitude, or perhaps even a pair of corresponding data values from a science experiment such as an object's volume at a given temperature.

Without knowing the reason someone might want to use our pair of `double` values, we can write a class that encapsulates them:

See Example: `DoublePair`

Notice that we have only very general-purpose code here. While it might make sense, for example, to add up the two numbers in the pair in some cases, we don't add that to our general purpose class.

Before we move on, we notice a couple of other items of interest.

- We provide a method `equals` that returns whether this `DoublePair` is the same as another `DoublePair`. `equals` is another of those methods (like `toString`) that are provided by Java for any object type. But like `toString`, we normally will want to provide our own `equals` method that determines the equality or equivalence of two objects of the class we are defining in some meaningful way. Here, we define `equals` so that two `DoublePair` objects are equal only if each of the numbers this one contains are equal to the corresponding numbers in the other.

Note also that the method signature of the `equals` method includes a formal parameter of type `Object`, even though we know that for this to make any sense, it should be passed as an object of type `DoublePair`. It is necessary to define it with `Object`, however, as `equals` methods must have this same method signature for all classes to conform to Java's rules.

The main complication of this fact is that we need to tell Java that we expect this `Object` to be a `DoublePair`, so we can make the meaningful comparison. The first line of the body of our `equals` method is a *cast*, which tells Java that the object we passed in as `o` should subsequently be treated as a `DoublePair` called `other`. The main drawback of this is that if we pass in some other type of object, the program will crash when the cast is attempted, with a `ClassCastException`. Uncomment the last line of the provided `main` method to see this happen.

- We provide a `main` method that tests our class. This method would not be used by a “user” of this class – it is provided only as a convenient way to test the class. It would work exactly the same way if we placed the `main` into a separate class.

Making it More General Purpose

While it's nice to have a pair of `double` values, what if we need a pair of `ints`, or `Strings`, or `PurchasedItems`? Or maybe a pair where the “first” is a `String` and the “second” is a `double`?

Java's `Object` class, which we just saw as a parameter to the `equals` method, provides a mechanism we can use for this.

See Example: `ObjectPair`

The code is nearly identical, except now our “first” and “second” can be any object type we wish.

In the `main` method, we actually pass primitive types as well, but that is being facilitated by Java's *autoboxing* functionality, supported by Java versions 1.5 and up.

Since primitive types, like `int`, `double`, and `boolean` are not objects, they do not qualify as a valid item to be stored in an `Object` variable. But, since we often do want to treat such values as objects, Java provides a set of classes, including `Integer`, `Double`, and `Boolean`, that are objects whose sole purpose is to include one `int`, `double`, or `boolean`, respectively. In old Java versions, we would have had to create such objects explicitly (e.g., "`new Double(9.1)`") but the Java compiler will now insert code to do this for us automatically.

Making it Generic

Another relatively new feature of Java, also introduced in Java 1.5, is to allow class definitions to include *generic*, or *parameterized data types*. This means that we can write a definition of the structure using data types that are unspecified (much like the value of a method parameter is unspecified) until we create an instance of the class. But once we "bind" to a type, we have to stick with that type (unlike how our `ObjectPair`'s "first" was initially a `String` and later a `Double`).

Here is the generic version of our pair class:

See Example: `GenericPair`

The important feature here is the use of *type parameters* to specify the actual data types we want to use for the first and second entries of our pair. These are specified throughout in this example as `U` and `V`. We assume (and in fact, require) that all places we refer to `first` are of type `U` and of `second` of type `V`. This includes declarations of instance variables, formal parameters, and return types.

Only when we construct an actual `GenericPair` object do we specify the types we wish to use. See in the provided `main` method how this is done.

Associations

Let's consider a very simple example of a data structure called an `Association`.

As the name suggests, an `Association` is a way to associate pairs of objects, one of which is the `key` and one of which is the `value`. Unlike the "pair" structures we just considered, once created, the `key` of an `Association` cannot be changed, only the `value` can.

This is another "general purpose" structure, but one with the above restriction.

When developing any such structure, there are some questions to be answered first.

- What should such a structure look like?
- What instance variables will it need?
- What constructors should be provided?

- What methods will it need?

As we have seen, we have two main options when developing a generic class. We can develop our structures to hold references to `Objects` and then use them to store instances of any Java class, or use type parameters.

Before we make use of these structure soon, let's take a look at the implementation of each that is provided with the supplemental Bailey text's "structure package".

The `Object`-based implementation:

See Structure Source:

`structure/Association.java`

The generic implementation:

See Structure Source:

`structure5/Association.java`

This class is defined as part of package `structure`, meaning it can access protected entries of other classes in the structure, and those classes can access class `Association`'s protected items.

The actual implementation of the `Association` class is pretty straightforward. A couple of quick notes:

- We require a key to construct a new `Association`, but the value is optional. If not provided, the value part defaults to `null`.
- Two `Associations` are considered equal (by the `equals` method) if their keys are the same, regardless of their values.
- We have an accessor for the key (`getKey`) but no mutator. Once created, the key of an `Association` may not be modified.
- For the value, we have both an accessor (`getValue`) and a mutator (`setValue`).

We will see examples that use the `Association` class soon.