Object Semantics

6.170 Lecture 2

The objectives of this lecture are to:

- to help you become familiar with the basic runtime mechanism common to all object-oriented languages (but with a particular focus on Java): variables, object references, assignments, mutability, and so on;

- to introduce a diagrammatic notation, *object diagrams* for describing 'snapshots' (that is, particular configurations of objects in the heap);[1]

- to make you aware, in passing, of some tricky issues that we'll return to later in more detail, and which turn out to be of fundamental importance: equality, rep invariants and exposure.

When you've completed this material, you should have a solid grasp of what happens when Java code executes, so that you can predict what some code will do without running it.

1   Variables, References and Objects

Some types of objects can be created with *literals*. What happens when you run this?

```
1.      String a = "zeeb";
2.      String b = a.toUpperCase ();
3.      System.out.println (b);
```

It prints ZEEB. Statement 2 is a call to the method `toUpperCase`. The method has a 'receiver' `a`. The call results in the creation of a fresh string that is then bound to the variable `b` .

We can draw the result of the first two statements as an *object diagram*, showing that `a` is a *reference* to an object of type `String` (and not the string itself, or a slot that holds it); similarly for `b`.
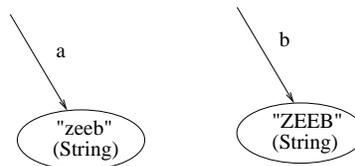
What does this do?

```
1.      String a = "zeeb";
2.      a.toUpperCase ();
3.      System.out.println (a);
```

---

[1]Later, we will introduce *object models* that are far more useful and which describe sets of snapshots.

It prints `zeeb`. Statement 2 creates a fresh string that gets thrown away since it is bound to no variable. The string object referenced by `a` is not changed: strings are immutable.

What about this?

```
String a = "zeeb";
a = a.toUpperCase ();
System.out.println (a);
```

Again, no object changes. But `a` is made to refer to the new object by the assignment in Statement 2, so the result is `ZEEB`.
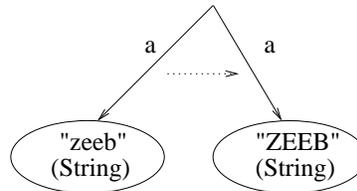
2   Aliasing, Mutability and Reference Equality

Java provides a variety of collections as part of its standard library. A vector is like an array, but it can grow and shrink dynamically. An example of using a vector:

```
1.        Vector v = new Vector ();
2.        Vector k = v;
3.        String a = "zeeb";
4.        v.add (a);
5.        k.add (a.toUpperCase());
6.        System.out.println (v.lastElement ());
```

The method `lastElement` is like the string method `toUpperCase`: it has no effect on the receiver and returns a reference to an object (in this case, the last element of the vector `v`). The method `add`, on the other hand, takes an argument the string `a`. Unlike `toUpperCase`, the method `add` does not return an object, but changes or *mutates* its receiver object. Vectors, unlike strings, can change, and are thus said to be *mutable*.

The code above prints `ZEEB`, since after Statement 2 the two variables `k` and `v` are names for the same vector object: they are said to be *aliases*. The calls to `add` mutate the one vector object, first adding the object for the lower case string, then the upper case string. The changes are visible through both names, so in Statements 4 to 6 we could actually permute the names `k` and `v` without any change in behavior.

Here is the object diagram:
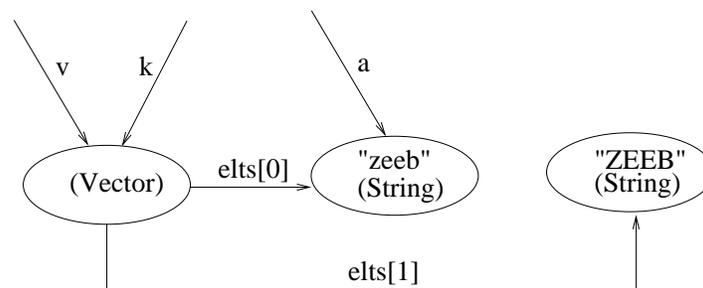
Aliasing is pervasive in languages like Java, and very useful. But it adds a lot of complexity. For one thing, it breaks the rule that a statement 'affects only the variables it mentions'. Just because v isn't mentioned in Statement 5 doesn't mean that it won't affect the result of statement 6 which mentions only v and not k.

How can we observe the aliasing more directly? By testing equality:

```
Vector v = new Vector ();
Vector k = v;
if (v == k)
    System.out.println ("same");
else
    System.out.println ("different");
```

which results in same being printed. The built-in == test tells you whether two references are for the same object, so it's often called a test of 'reference equality'.

What does this do?

```
Vector v = new Vector ();
Vector k = new Vector ();
if (v == k)
    System.out.println ("same");
else
    System.out.println ("different");
```

It prints different, because v and k are distinct objects. It's a fundamental property of constructors that the objects they return really are fresh. In fact, the garbage collector can recycle an object, but only if there is no reference to it still around. This ensures that even if objects are recycled, we can never tell.

A puzzle: what does this do?

```
String a = "zeeb";
String b = "zeeb";
if (a == b)
    System.out.println ("same");
else
    System.out.println ("different");
```

Strangely, this prints same, because the Java virtual machine automatically 'interns' string literals: if it can tell that two string literals have the same sequence of characters, it only allocates one object. You'd be right to think this is a bit confusing; it's a performance optimization.

In fact, it's very bad form to test reference equality of immutable objects, unless you're doing something subtle with memory management. Arguably it's a design defect of Java that you can even observe whether two immutable objects are the same or not. So how should you compare two immutable objects? With an equals method.

The String class provides a method equals that tells you whether two strings contain the same sequence of characters or not. This code

```
String a = "zeeb";
String b = a.toUpperCase ();
```

```
    if (b.equals ("ZEEB"))
        System.out.println ("same characters");
    else
        System.out.println ("different characters");
```

prints `same characters`.

When we study inheritance, you'll learn that every class automatically inherits an `equals` method, so you might think you don't need to write one. But it's almost never what you want, so whenever you design a class, one of the first things you'll need to figure out is when two objects of the class should be considered equal to one another.

Here are some questions:

- Would you expect that generally `x == y` implies `x.equals (y)`? Yes, it should. Because the `equals` method can be user-defined, just like any other method, you could make it behave in any way you wanted. On a mutable type, it might even mutate the object! But that would be disastrous: there's a generic contract that clients expect equals to obey. More on this later.

- Why would a language have immutable types? Because aliasing is complicated, and when you use immutable types, the issue doesn't arise. Also, code built with immutable types can sometimes be more efficient.

- So if immutable types are so much simpler, why have mutable types? Because mutation gives a very useful form of modularity: it allows you to make local changes to a structure. And mutation is often a natural way to model entities in the real world: a transaction on a bank account changes it; it doesn't produce a new bank account.

3 Null References

What does this do?

```
    String a = null;
    System.out.println (a);
```

It prints `null`. The keyword `null` denotes a value that can be taken on by an object reference. It means that the reference does not in fact refer to any object. There is no null object!

But note that this code

```
1.          String a = null;
2.          String b = a.toUpperCase ();
3.          System.out.println (b);
```

behaves quite differently. It throws a `NullPointerException` on Statement 2. We'll learn about exceptions later, but for now, all you need to understand is that Statement 2 failed, when the expression `a.toUpperCase()` was evaluated.

What's the difference? The receiver to a method call can never be null, because it identifies the object that 'receives' the call—and that has to be some object. So `a.toUpperCase()` fails when `a` is null. But in the previous example, `System.out.println(a)` is OK when `a` is null, since reference is an argument, and this is a special kind of method (really just a plain old procedure) that doesn't have a receiver.

You can write a method that tests whether an argument is null and does something appropriate. Dereferencing null is a common programming mistake in Java. To avoid it, you can check whether a reference is null before you attempt to call a method.

In general, rather than catching nulls and treating them specially, it's better to avoid creating null references in the first place. You'll learn about that when we discuss representation invariants. Sometimes you can't avoid it, and then it's important to document where the null references may occur. That's one reason specifications are important: they can spare you runtime errors and unnecessary checks.

Here are some questions:

- In general, would you expect `a.equals (b)` to be substitutable for `b.equals (a)` ? No, because when `a` is null and `b` is not, the first will throw an exception, and the second will (usually) return false.

- OK, smarty pants, leave nulls alone. Would you then expect `a.equals (b)` and `b.equals (a)` to have the same effect? Yes, you would. In fact, this property of the `equals` method – called *symmetry* – is demanded by Java's 'object contract'. We'll see later when we study equality in depth what other properties are required, and how easy it is to mess up and write an `equals` method that does *not* have these properties.

4   User-defined Classes and Fields

Let's make some objects of our own:

```
class Trans {
        int amount;
        Date date;
}
```

This code declares a *class*, a kind of template for making objects. These objects are going to represent transactions in a banking system. Each object has an integer amount (which may be negative for a withdrawal), and a date/time stamp to mark the moment at which the transaction occurred.

The class declares two *fields* or *instance variables*, `amount` and `date`. Each object of the class will contain two references, one to an integer and one to a date. The type `Date` is a class from the Java library; it's predefined like `String` (but not part of the language definition the way `String` is).

The type `int` is a rather strange beast. It's not a class at all, but a *primitive type*. Variables or fields of type `int` don't hold references to integer objects; they hold the integers themselves. You may think it a bit jarring that an object-oriented language has this rather unobject-oriented notion in it (and many people share your opinion). Sometimes we'll actually need an integer that's an object, and in that case we can use the class `Integer` from the Java library. How do you get from an `int` to an `Integer` and back? To create an `Integer`, you use a constructor:

```
int_i = 5;
Integer obj_i = new Integer (i);
```

and to extract the primitive integer from an integer object, you call a method:

```
int i = obj_i.intValue();
```

A little cumbersome, so that's one reason people don't like this design.
If we run this code:

```
1.      Trans t = new Trans ();
2.      t.amount = 20;
3.      t.date = new Date ();
```

a fresh object gets created, and its fields are set, resulting in this configuration:
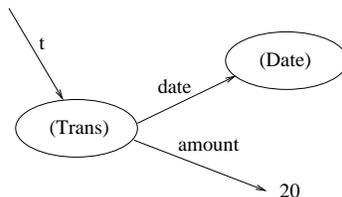
The expression on the right-hand side of Statement 1 is a call to a constructor: it makes a new object with default values for the fields. In this case, it creates a `Trans` object with zero for the `amount` and null for `date`. The Statement 2 is called a *setter*: it sets the value of the field `amount` of the object referred to by `t`. Statement 3 has another constructor call on the right; creating a new `Date`, which by default creates a `Date` object representing the moment at which the object is itself created. But it's also a setter: it sets the `date` field of `t` to point to this new date.

5 User-defined Constructors

So we've succeeded in making a `Trans` object representing a deposit of twenty dollars at this moment in time. The way we did it – creating an uninitialized object and then setting its fields – is not a good one, however. We'll want our transactions to be well-formed; for example we won't want to have transactions that don't have dates. And perhaps we'll want every transaction to have a non-zero amount. Later, we'll study these kinds of *invariants* in much more depth.

For now, just observe that immediately after Statement 1 we have a transaction object that is *not* well formed. When an object is created, its fields are initialized to default values: null for object references, and zero for integers. So `t.amount` will be zero, and `t.date` will be null. These default values are rarely what you want; after all, which values make sense will depend on the problem we're trying to solve. In this case, you'd have to know something about banking to know that a transaction of zero dollars is ill-formed.

Is it a big deal that there's a bogus transaction object hanging around between Statements 1 and 2? Yes, it is, and here's why. We'd like the responsibility for ensuring that objects of the class `Trans` are well formed to be handled entirely within the `Trans` class. In our program, you need to check not only the code of the class, but also the code that uses the class. At this scale, it's not a disaster. But in a much larger program, you need as much *modularity* as you can get, confining tricky aspects of the program as much as possible to small areas of the code.

To solve this problem, we declare our own constructor:

```
class Trans {
      int amount;
```

```
        Date date;
        Trans (int a, Date d) {amount = a; date = d;}
    }
```

This constructor takes an amount and a date as arguments, and creates a transaction object with that amount and date. If I'd wanted to ensure that the amount of a transaction is non-zero, I could have added a check that threw an exception if the amount was zero; we'll see how to do that later.

A peculiar, but useful, property of constructors is that having defined our own constructor, the default constructor – the one that just initializes each field to default values – becomes no longer available. So Statement 1 will no longer compile. Instead, we can write

```
    Trans t = new Trans (20, new Date ());
```

and that one line of code will have the effect that Statements 1 to 3 had previously.

(We haven't fully solved the problem of modularizing the invariant of `Trans`, by the way. You can still mess up a `Trans` object by setting its `date` field to null from outside, for example. The first step to prevent this is to make use of Java's accessibility mechanisms: we can make the fields *private* so they can't be read or written from outside the class. In fact, this turns out not to be enough, as we'll learn when we study data abstraction.)

6   Conclusion

Now we've seen all the key notions for how objects are manipulated in an object-oriented language. We've seen how they're created; how references are bound to objects; and how their fields are set. We've mentioned the two fundamental kinds of equality – reference equality (tested with `==`) and object equality (tested with an `equals` method – about which we'll have much more to say later. These mechanisms comprise one important part of what it means to be 'object oriented'.