

# Equality

6.170 Lecture 8

Spring 2004

*I like pigs. Dogs look up to us. Cats look down on us. Pigs treat us as equals.*

Winston Churchill

## 1 The Object Contract

In Java, every class extends `Object`, and therefore inherits all of its methods. Two of these are particularly important and consequential in all programs: the method for testing equality, and the method for generating a hash code.

```
public boolean equals (Object o);  
public int hashCode ();
```

Like any other methods of a superclass, these methods can be overridden. We have mentioned in class that a subclass should be a *subtype*. This means that it should behave according to the specification of the superclass, so that an object of the subclass can be placed in a context in which a superclass object is expected, and still behave appropriately.

The specification of the `Object` class given in its documentation is rather abstract and may seem abstruse. But failing to obey it has dire consequences, and tends to result in horribly obscure bugs. Worse, if you do not understand this specification and its ramifications, you are likely to introduce flaws in your code that have a pervasive effect and are hard to eliminate without major reworking.

The specification of the `Object` class is so important that it is often referred to as ‘The Object Contract’. The contract can be found in the method specifications for `equals` and `hashCode` in the Java API documentation. It states that:

- ▷ `equals` must define an equivalence relation – that is, a relation that is reflexive, symmetric, and transitive;
- ▷ `equals` must be consistent: repeated calls to the method must yield the same result unless the arguments are modified in between;
- ▷ for a non-null reference `x`, `x.equals (null)` should return false; and
- ▷ `hashCode` must produce the same result for two objects that are deemed equal by the `equals` method.

## 2 Equality and inheritance

For the moment, let us ignore the `hashCode` method and look first at the properties of the `equals` method. *Reflexivity* means that an object always equals itself; *symmetry* means that when `a` equals `b`, `b` equals `a`; *transitivity* means that when `a` equals `b` and `b` equals `c`, `a` also equals `c`.

These may seem like obvious properties, and indeed they are. If they did not hold, it is hard to imagine how the `equals` method would be used: you would have to worry about whether to write `a.equals(b)` or `b.equals(a)`, for example, if it were not symmetric.

What is much less obvious, however, is how easy it is to break these properties inadvertently. The following example (taken from Joshua Bloch's excellent *Effective Java: Programming Language Guide*, one of the course's recommended texts) shows how symmetry and transitivity can be broken in the presence of inheritance. Consider a simple class that implements a two-dimensional point:

```
public class Point {
    private final int x;
    private final int y;
    public Point (int x, int y) {
        this.x = x; this.y = y;
    }
    public boolean equals (Object o) {
        if (!(o instanceof Point))
            return false;
        Point p = (Point) o;
        return p.x == x && p.y == y;
    }
    ...
}
```

The `equals` method takes an `Object` as its argument; this is mandated by the object contract. The method returns `true` if the object passed is a `Point` and stores the same coordinates. Now suppose we add the notion of a color in a derived class:

```
public class ColorPoint extends Point {
    private Color color;
    public ColorPoint (int x, int y, Color color) {
        super (x, y);
        this.color = color;
    }
    ...
}
```

What should the `equals` method of `ColorPoint` look like? We could just inherit `equals` from `Point`, but then two `ColorPoints` will be deemed equal even if they have different colors. We could override it like this:

```
public boolean equals (Object o) {
    if (!(o instanceof ColorPoint))
        return false;
    ColorPoint cp = (ColorPoint) o;
    return super.equals (o) && cp.color.equals (color);
}
```

This seemingly inoffensive method actually violates the requirement of symmetry. To see why, consider a point and a color point:

```
Point p = new Point (1, 2);
ColorPoint cp = new ColorPoint (1, 2, Color.RED);
```

Now `p.equals(cp)` will return true, but `cp.equals(p)` will return false! The problem is that these two expressions use different `equals` methods: the first uses the method from `Point`, which ignores color, and the second uses the method from `ColorPoint`.

We could try to fix this by having the `equals` method of `ColorPoint` ignore color when comparing against a non-colored `Point`:

```
public boolean equals (Object o) {
    if (!(o instanceof Point))
        return false;
    // if o is a normal Point, do a color-blind compare
    if (!(o instanceof ColorPoint))
        return super.equals (o);
    ColorPoint cp = (ColorPoint) o;
    return super.equals (o) && cp.color.equals (color);
}
```

This solves the symmetry problem, but now equality isn't transitive! To see why, consider constructing these points:

```
ColorPoint p1 = new ColorPoint(1, 2, Color.RED);
Point p2 = new Point (1, 2);
ColorPoint p3 = new ColorPoint (1, 2, Color.BLUE);
```

The calls `p1.equals(p2)` and `p2.equals(p3)` will both return true, but `p1.equals(p3)` will return false.

It turns out there is no solution to this problem: it is a fundamental flaw in inheritance. You cannot write a good `equals` method for `ColorPoint` if it inherits from `Point` as written. However, there are two workarounds. The first is to change `Point`'s `equals` method so that it rejects equality with any of its subclasses:

```
public class Point {
    ...
    public boolean equals (Object o) {
        if (o == null || !o.getClass().equals(getClass ()))
            return false;
        Point p = (Point) o;
        return p.x == x && p.y == y;
    }
    ...
}
```

Explicit comparison of classes is a stronger test than `instanceof`. A `ColorPoint` is an `instanceof Point`, but `ColorPoint.getClass() != Point.getClass()`. The drawback of this approach is that you lose the ability for harmless `Point` subclasses to compare for equality to a `Point`. For

example, you might write a subclass `PolarPoint` that doesn't add any new attributes to `Point`, but merely offers some new methods for obtaining the point in polar coordinates. With the `getClass` workaround, a `PolarPoint` can never equal a `Point`, even though it has the same (x,y) value internally. The second workaround is to implement `ColorPoint` using `Point` in its representation, rather than inheriting it:

```
public class ColorPoint {
    Point pt;
    Color color;
    ...
}
```

Since `ColorPoint` no longer extends `Point`, the problem goes away. See Bloch's book for more details about this approach, which is called composition. Bloch's book also give some hints on how to write a good `equals` method, and he points out some common pitfalls. For example, what happens if you write something like this:

```
public boolean equals (Point p)
```

where you've substituted another type for `Object` in the declaration of `equals`?

### 3 Hash tables

To understand the part of the `Object` contract relating to the `hashCode` method, you'll need to have some idea of how hash tables work. Hash tables are a fantastic invention – one of the best ideas of computer science. A hash table is a representation for a mapping: an abstract data type that maps keys to values. Hash tables offer constant time lookup, so they tend to perform better than trees or lists. Keys don't have to be ordered, or have any particular property, except for offering `equals` and `hashCode`.

Here's how a hash table works. It contains an array that is initialized to a size corresponding the number of elements that we expect to be inserted. When a *key* and a *value* are presented for insertion, we compute the hashcode of the key, and convert it into an index in the array's range (e.g., by a modulo division). The value is then inserted at that index.

The rep invariant of a hash table includes the fundamental constraint that keys are in the slots determined by their hash codes. We'll see later why this is important.

Hashcodes are designed so that the keys will be spread evenly over the indices. But occasionally a *conflict* occurs, and two keys are placed at the same index. So rather than holding a single value at an index, a hash table actually holds a list of key/value pairs (usually called 'hash buckets'), implemented in Java as objects from class with two fields. On insertion, you add a pair to the list in the array slot determined by the hash code. For lookup, you hash the key, find the right slot, and then examine each of the pairs until one is found whose key matches the given key.

Now it should be clear why the `Object` contract requires equal objects to have the same hash key. If two equal objects had distinct hash keys, they might be placed in different slots. So if you attempt to lookup a value using a key equal to the one with which it was inserted, the lookup may fail.

A simple and drastic way to ensure that the contract is met is for `hashCode` to always return some constant value, so every object's hash code is the same. This satisfies the `Object` contract, but it would have a disastrous performance effect, since every key will be stored in the same slot, and every lookup will degenerate to a linear search along a long list.

The standard way to construct a more reasonable hash code that still satisfies the contract is to compute a hash code for each component of the object that is used in the determination of equality (usually by calling the `hashCode` method of each component), and then combining these, throwing in a few arithmetic operations. Look at Joshua Bloch's book *Effective Java* for details.

Most crucially, note that if you don't override `hashCode` at all, you'll get the one from `Object`, which is based on the address of the object. If you have overridden `equals`, this will mean that you will have almost certainly violated the contract. So as a general rule:

*Always override hashCode when you override equals.*

(This is one of Bloch's aphorisms. The whole book is a collection of aphorisms like it, each nicely explained and illustrated.)

Some years ago, so the legend goes, a 6.170 student spent hours tracking down a bug in a project that amounted to nothing more than misspelling `hashCode` as `hashcode`. This created a method that didn't override the `hashCode` method of `Object` at all, and strange things happened...

## 4 Element and Container Equality

When are two containers equal? If they are immutable, they should be equal if they contain the same elements. For example, two strings should be equal if they contain the same characters (in the same order). Otherwise, if we just kept the default `Object` equality method, a string entered at the keyboard, for example, would never match a string in a list or table, because it would be a new string object and therefore not the same object as any other. And indeed, this is exactly how `equals` is implemented in the Java `String` class, and if you want to see if two strings `s1` and `s2` contain the same character sequence, you should write

```
s1.equals (s2)
```

and not

```
s1 == s2
```

which will return false when `s1` and `s2` denote different string objects that contain the same character sequences.

So much for strings – sequences of characters. Let's consider lists now, which are sequences of arbitrary objects. Should they be treated the same way, so that two lists are equal if they contain the same elements in the same order?

Suppose I'm planning a party at which my friends will sit at several different tables, and I've written a program to help me create a seating plan. I represent each table as a list of friends, and the party as a whole as a set of these lists. The program starts by creating empty lists for the tables and inserting them into the hash set:

```
List t1 = new LinkedList ();
List t2 = new LinkedList ();
...
Set s = new HashSet ();
s.add (t1);
s.add (t2);
...
```

Think of the hashset as a hash table with no values. At some later point, the program will add friends to the various lists; it may also create new lists and replace existing lists in the set with them. Finally, it iterates over the contents of the set, printing out each list.

This program will fail, because the initial insertions will not have the expected effect. Even though the empty lists represent conceptually distinct table plans, they will be equal according to the `equals` method of `LinkedList`. Since `Set` uses the `equals` method on its elements to reject duplicates, all insertions but the first will have no effect, since all of the empty lists will be deemed duplicates.

How can we solve this problem? You might think that `Set` should have used `==` to check for duplicates instead, so that an object is regarded as a duplicate only if that very object is already in the set. But that wouldn't work for strings; it would mean that after:

```
Set set = new HashSet ();
String lower = "hello";
String upper = "HELLO";
set.add (lower.toUpperCase());
```

the test `set.contains (upper)` would evaluate to `false`, since the `toUpperCase` method creates a new string.

## 5 The Liskov Approach

In our course text, Professor Liskov presents a systematic approach to dealing with this problem. You provide two distinct methods: `equals`, which returns true when two objects in a class are behaviorally equivalent, and `similar`, which returns true when two objects are observationally equivalent.

Here's the difference. Two objects are *behaviorally equivalent* if there is no sequence of operations that can distinguish them. On these grounds, the empty lists `t1` and `t2` from above are not equivalent, since if you insert an element into one, you can see that the other doesn't change. But two distinct strings that contain the same sequence of characters are equivalent, since you can't modify them and thus discover that they are different objects. (We're assuming you're not allowed to use `==` in this experiment.)

Two objects are *observationally equivalent* if you can't tell the difference between them using observer operations (and no mutations). On these grounds, the empty lists `t1` and `t2` from above are equivalent, since they have the same size, contain the same elements, etc. And two distinct strings that contain the same sequence of characters are also equivalent.

Here's how you code `equals` and `similar`. For a mutable type, you simply inherit the `equals` method from `Object`, but you write a `similar` method that performs a field-by-field comparison. For an immutable type, you override `equals` with a method that performs a field-by-field comparison, and have `similar` call `equals` so that they are the same.

This solution, when applied uniformly, is easy to understand and works well. But it's not always ideal. Suppose you want to write some code that interns objects. This means mutating a data structure so that references to objects that are structurally identical become references to the very same object. This is often done in compilers; the objects might be program variables, for example, and you want to have all references to a particular variable in the abstract syntax tree point to the same object, so that any information you store about the variable (by mutating the object) is effectively propagated to all sites in which it appears.

To do the interning, you might try to use a hash table. Every time you encounter a new object in the data structure, you look that object up in the table to see if it has a canonical representative. If it does, you replace it by the representative; otherwise you insert it as both key and value into the table.

Under the Liskov approach, this strategy would fail, because the equality test on the keys of the table would never find a match for distinct objects that are structurally equivalent, since the `equals` method of a mutable object only returns true on the very same object. You need to wrap the key with an object whose `equals` method is the `similar` method on the key, or write your own class. Of course interning mutable objects in this way is dangerous and you have to be very careful to avoid bugs. A better way would be to create an immutable ‘snapshot’ of the object to be stored, and store *that* in the collection. Then, the Liskov approach would work.

## 6 Java Approach

The designer of the Java collections API did *not* follow the Liskov approach. There is no `similar` method, and `equals` is observational equivalence.

This has some convenient consequences. The interning table will work, for example. But it also has some unfortunate consequences. The seating plan program will break, because two distinct empty lists will be deemed equal.

In the Java `List` specification, two lists are equal not only if they contain the same elements in the same order, but also if they contain `equal` elements in the same order. In other words, the `equals` method is called recursively. To maintain the `Object` contract, the `hashCode` method is also called recursively on the elements. This results in a very nasty surprise. The following code, in which a list is inserted into itself, will actually fail to terminate!

```
List l = new LinkedList ();
l.add (l);
int h = l.hashCode ();
```

This is why you’ll find warnings in the Java API documentation about inserting containers into themselves, such as this comment in the specification of `List`:

*Note: While it is permissible for lists to contain themselves as elements, extreme caution is advised: the equals and hashCode methods are no longer well defined on such a list.*

There are some other, even more subtle, consequences of the Java approach, such as the risk of representation exposure. risk of rep exposure. Suppose part of `Set`’s rep invariant is the requirement that it contain no duplicates, i.e., that for all  $i \neq j$  `elt[i].equals(elt[j])` always returns false. If `equals` is defined as observational equivalence, then a client can mutate an object in the set to make it equal to another object in the set, thereby breaking `Set`’s rep invariant. The Liskov approach doesn’t suffer from this problem. Behavioral equivalence ensures that if two objects are equal or unequal at one point, they continue to be equal or unequal forever after; there is nothing you can do to two unequal objects that will make them become equal.

This leaves you with two choices, both of which are acceptable in 6170:

- ▷ You can follow the Java approach, in which case you’ll get the benefits of its convenience, but you’ll have to deal with the complications that can arise.

- ▷ Alternatively, you can follow the Liskov approach, but in that case you'll need to figure out how to incorporate into your code the Java collection classes (such as `LinkedList` and `HashSet`).

In general, when you have to incorporate a class whose `equals` method follows a different approach from the program as a whole, you can write a wrapper around the class that replaces the `equals` method with a more suitable one.