Object-Oriented Design	Lecture 18
CS 3500 Spring 2010 (Pucella)	Tuesday, Mar 23, 2010

18 Mutation

We have been avoiding mutations until now. But in most languages, including Java, they exist, for better or for worse — especially in the libraries. So we need to know how to deal with them.

Recall that a class is *immutable* is the state of an instance of the class never changes after it has been created (where the state of an instance includes the value of all its fields). In contrast, a class is *mutable* if its instances may change state during their lifetime.

So let's move towards understanding mutation. Mutation can be problematic because an object can change under you without you noticing, leading to hard to find bugs. This is the root of my advocacy for immutability — it is just plain easier to reason about immutable code. Immutable code is also easier to parallelize, but that's a bit beyond what we're looking at here.

For the purpose of this lecture, we shall consider the following two classes. First, a class `Point` representing two-dimensional points:

```
public class Point {
  private int xpos;
  private int ypos;

  private Point (int x, int y) {
   xpos = x;
   ypos = y;
  }

  public static Point create (int x, int y) {
   return new Point(x,y);
  }
  public int xPos () {
   return xpos;
  }
  public int yPos () {
   return ypos;
  }
  public String toString () {
```

```
    return "(" + xpos + "," + ypos + ")";
  }
}
```

Second, a class `Line` representing lines:

```
public class Line {
  private Point start;
  private Point end;

  protected Line (Point s, Point e) {
    start = s;
    end = e;
  }

  public static Line create (Point s, Point e) {
    return new Line(s,e);
  }
  public Point start () {
    return start;
  }
  public Point end () {
    return end;
  }
  public String toString () {
    return "[" + start + "," + end + "]";
  }
}
```

A class can be made mutable in two ways: by making its fields public (so that instances of the class can be modified by the outside world), or by having methods that change the value of fields.

For the time being, let's make `Point` mutable by making its fields public, and illustrate why mutation is hard to reason about:

```
public class Point {
  public int xpos;
  public int ypos;

  private Point (int x, int y) {
    xpos = x;
```

```
    ypos = y;
  }

  public static Point create (int x, int y) {
    return new Point(x,y);
  }
  public int xPos () {
    return xpos;
  }
  public int yPos () {
    return ypos;
  }
  public String toString () {
    return "(" + xpos + "," + ypos + ")";
  }
}
```

Mutation means that invoking a method on an instance can return multiple values, depending on when the method is invoked (before or after a mutation). For instance:

```
Point p1 = Point.create(0,0);
System.out.println("Point p1 = " + p1.toString());
p1.xpos = 99;
System.out.println("Point p1 = " + p1.toString());
```

which yields:

```
Point p1 = (0,0)
Point p1 = (99,0)
```

Here, the call to `toString()` returns different results, even though it is invoked on the same value `p1`. One difficulty with mutation is an update may be hidden in some other method, which provides no indication that it is changing the instance. Suppose that we have a function

```
public static void someRandomFunction (Point p) {
  p.xpos = 99;
}
```

and write:

```
Point p1 = Point.create(0,0);
```

```
System.out.println("Point p1 = " + p1);
someRandomFunction(p1);
System.out.println("Point p1 = " + p1);
```

Again, this mutates point `p1`, but there is no indication that the call to `someRandomFunction()` does a mutation:

```
Point p1 = (0,0)
Point p1 = (99,0)
```

This is made worse, as we shall soon see, by the fact that point `p1` may be shared between many different objects.

Another difficulty with mutability is that it is a *contagious property*: a class that looks immutable may in fact be mutable if it relies on classes that are themselves mutable. Consider the mutable implementation of `Point`s above. What about `Line` though? Its fields are private, it never changes the value of its fields, so as far as one can tell by looking at the class definition, it is an immutable class. Unfortunately, the following snippet of code shows that an instance of `Line` can indeed change:

```
Point p1 = Point.create(0,0);
Point p2 = Point.create(10,10);
Line l = Line.create(p1,p2);
System.out.println("Line l = " + l);
p1.xpos = 99;
System.out.println("Line l = " + l);
```

which outputs:

```
Line l = [(0,0),(10,10)]
Line l = [(99,0),(10,10)]
```

That's something to keep in mind: a class may be mutable even though it looks like it is not. As soon as part of your program is mutable, it will have a tendency to make the rest of your program immutable as well.

Part of the point of this lecture is to provide a model of mutation so that you can understand exactly what is happening in the examples above.

18.1 Detour: Setters for Fields

So let's suppose that you have understood this issue about mutability, and that you accept the ensuing risks. In other words, you decided to have some classes with mutable state, which generally means having mutable fields.

Even if you are okay with having mutable fields, you still want to keep your fields private, and provide selectors and setters for each field that can mutate. Why? Because this lets us enforce *invariants*. Suppose we only want to work with points in the positive quadrant, that is, points whose coordinates are nonnegative. That's easy to enforce with the above `Point` class by changing the creator or constructor:

```
public static Point create (int x, int y) {
  if (x < 0 || y < 0)
    throw new IllegalArgumentException("negative");
  return new Point(x,y);
}
```

The creator enforces the invariant that the coordinates of a point are always nonnegative.

If we allow unrestricted field access, however, then anyone can just change one of the bounds and break the invariant. Which, in this case, can lead to points having negative coordinates, invalidating the invariant we want to preserve.

By forcing users to use setters, we can check that the invariant is maintained whenever state is changed:

```
public void setXPos (int x) {
  if (x < 0)
    throw new IllegalArgumentException("negative");
  xpos = x;
}

public void setYPos (int y) {
  if (y<0)
    throw new IllegalArgumentException("negative");
  ypos = y;
}
```

Therefore, I will expect you all to keep fields private and use explicit setters, instead of making fields public when you want a class to be mutable.

Therefore, we shall use the following (mutable) implementation of `Points` in the rest of this lecture:

```
public class Point {
  private int xpos;
  private int ypos;

  private Point (int x, int y) {
```

```
    xpos = x;
    ypos = y;
  }

  public static Point create (int x, int y) {
    return new Point(x,y);
  }
  public int xPos () {
    return xpos;
  }
  public int yPos () {
    return ypos;
  }
  public void setXPos (int x) {
    xpos = x;
  }
  public void setYPos (int y) {
    ypos = y;
  }
  public String toString () {
    return "(" + xpos + "," + ypos + ")";
  }
}
```

Our implementation of `Line` remains the same, but as I alluded to earlier, it is also mutable by virtue of `Point` being mutable.

18.2 Object Creation and Manipulation

To work with mutation correctly, and help you track down bugs, you need to have a good understanding of what changes when something changes! You update some field in some class, what actually gets changed, and who else can see it? The problem is that when the state of an object can change, it becomes very important to understand when objects are *shared* between other objects, so that we can track when a change can be seen from another object.

To pick a silly example, if we write:

```
Point p1 = Point.create(0,0);
Point p2 = Point.create(0,0);
```

Then computing with `p1` and `p2` both give the same result, and if we mutate `p1` by calling `setXPos()`, `p2` is unaffected. Contrast that to:

```
Point p1 = Point.create(0,0);
Point p2 = p1;
```

where again computing with `p1` and `p2` both give the same result, but if we mutate `p1` anywhere, then `p2` is changed as well. Tracking this kind of sharing is what makes working with mutation error-prone.

My claim is that to understand mutation, you need to have a working model of how Java represents objects internally. It does not need to be an accurate model; it just needs to have good predictive power. Let me describe a basic model that answers the question: where do variables live? (The question 'where do methods live?' is less interesting because methods are not mutable.)

Recall that when you create an object with a `new` statement, a block of memory is allocated in memory (in the heap) representing the new object, which gets filled by the constructor. Here is how we represent an object in the heap:

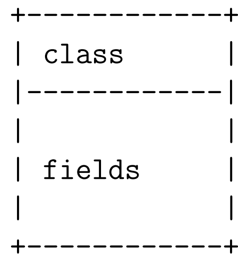

This representation does not have include the methods, because that's not what I want to focus on for now. The value returned by a constructor is actually an address, namely, the address where the object lives in memory. (This is sometimes called an object reference.) Thus, for instance, when you write

```
Point p1 = Point.create(0,10);
```

a new object is created in memory, say living at some address *addr*, and variable `p1` holds value *addr*. We can represent this as follows:

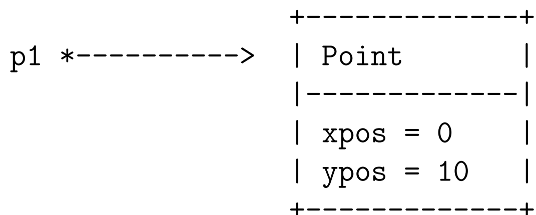

Now, when you pass an object as an argument to a method, what you end up passing is the *address* of that object (that is, the value returned by the constructor). That is how objects get manipulated.

Consider lines. When you create a `Line` object, passing in two points, you get as values of the fields `start` and `end` in the created line the addresses of the two points that were used to create the line in the first place.

```
Point p1 = Point.create(0,10);
Point p2 = Point.create(0,20);
Line l1 = Line.create(p2,p1);
```

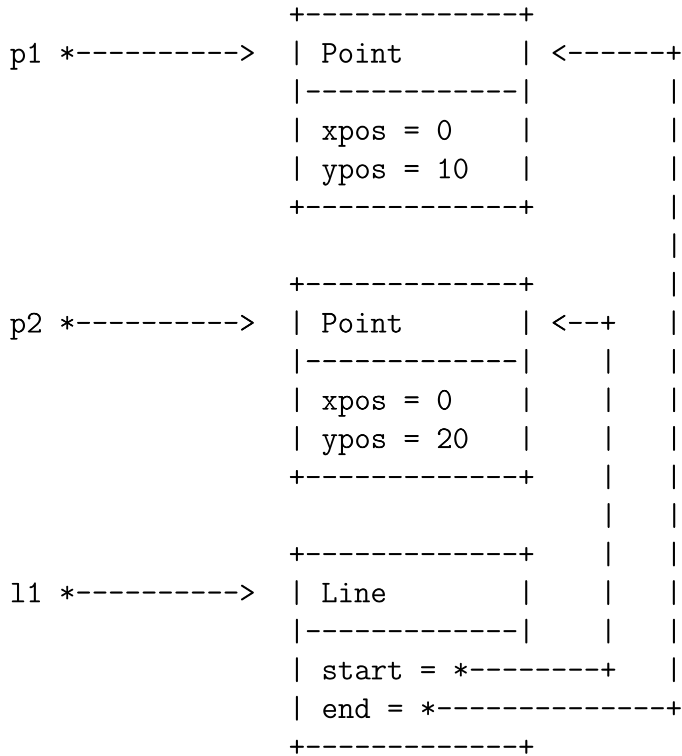

To find the value of a field, you follow the arrows to the object that holds the field you are trying to access. Thus, `p1.xPos()` looks up the `xpos` field in the object pointed to by `p1`. Similarly, `l1.end().yPos()` access the `ypos` field of object returned by `l1.end()`, which itself can be found as field `end` in the object pointed to by `l1`.

In particular, you see why if after creating the above we write

```
p1.setYPos(5);
System.out.println("Line l1 = " + l1);
```

we get `Line l1 = [(0,20),(0,5)]`; intuitively, because `p1` is the same object as the point stored as the end point in `l1`. We call this phenomenon *sharing*. It is reflected by the fact that there are two arrows pointing to an object in a diagram.

Compare the above by what gets constructed if we write:

```
Point p1 = Point.create(0,10);
Line l1 = Line.create(p1,p1);
```

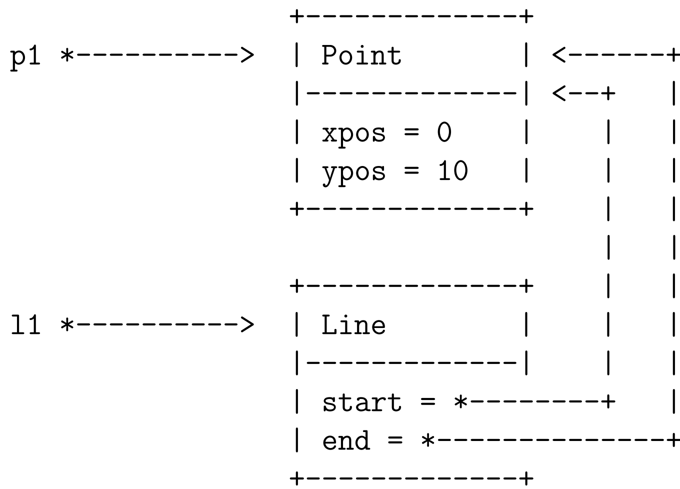

Here, the *same* point is used as first and second. Meaning, in particular, that if we update field `xpos` of `p1`, both the start and end points of `l1` are also changed.

18.3 Static Fields

Static fields are fields annotated with the `static` qualifier. We haven't discussed them until now, because frankly they make no sense unless we have mutability. Intuitively, a static field is associated with a class, as opposed to the instances of a class. More precisely, a static field is a field that is shared amongst all the instances of a class. So if one instance of the class updates the fields, that update is seen by all other instances.

Suppose we create an alternate class `LineCount` instead of `Line` that keeps track of the number of lines that has been created. The code is straightforward:

```
public class LineCount {
  private Point start;
  private Point end;
  private static int count = 0;

  private LineCount (Point s, Point e) {
    start = s;
    end = e;
    count = count + 1;
  }

  public static LineCount create (Point s, Point e) {
```

```
    return new LineCount(s,e);
  }

  public Point start () {
    return start;
  }
  public Point end () {
    return end;
  }
  public int count () {
    return count;
  }
  public String toString () {
    return "[" + start + "," + end + "] @ " + count;
  }
}
```

So what happens when this executes? Think of the static field `count` as somehow associated with the class rather than with instances of the class. In other words, there is a single copy of that field around, and all the instances of `LineCount` refer to that field. We are going to model this by having the class field of an object contain a link to a structure holding all the static fields of a class. This structure, which is shared amongst all the objects of the class, is created automatically by Java when you execute your program, before it yields control to your `main` method.

```
Point p1 = Point.create(0,10);
Point p2 = Point.create(0,20);
LineCount l1 = LineCount.create(p2,p1);
LineCount l2 = LineCount.create(p2,p1);
```

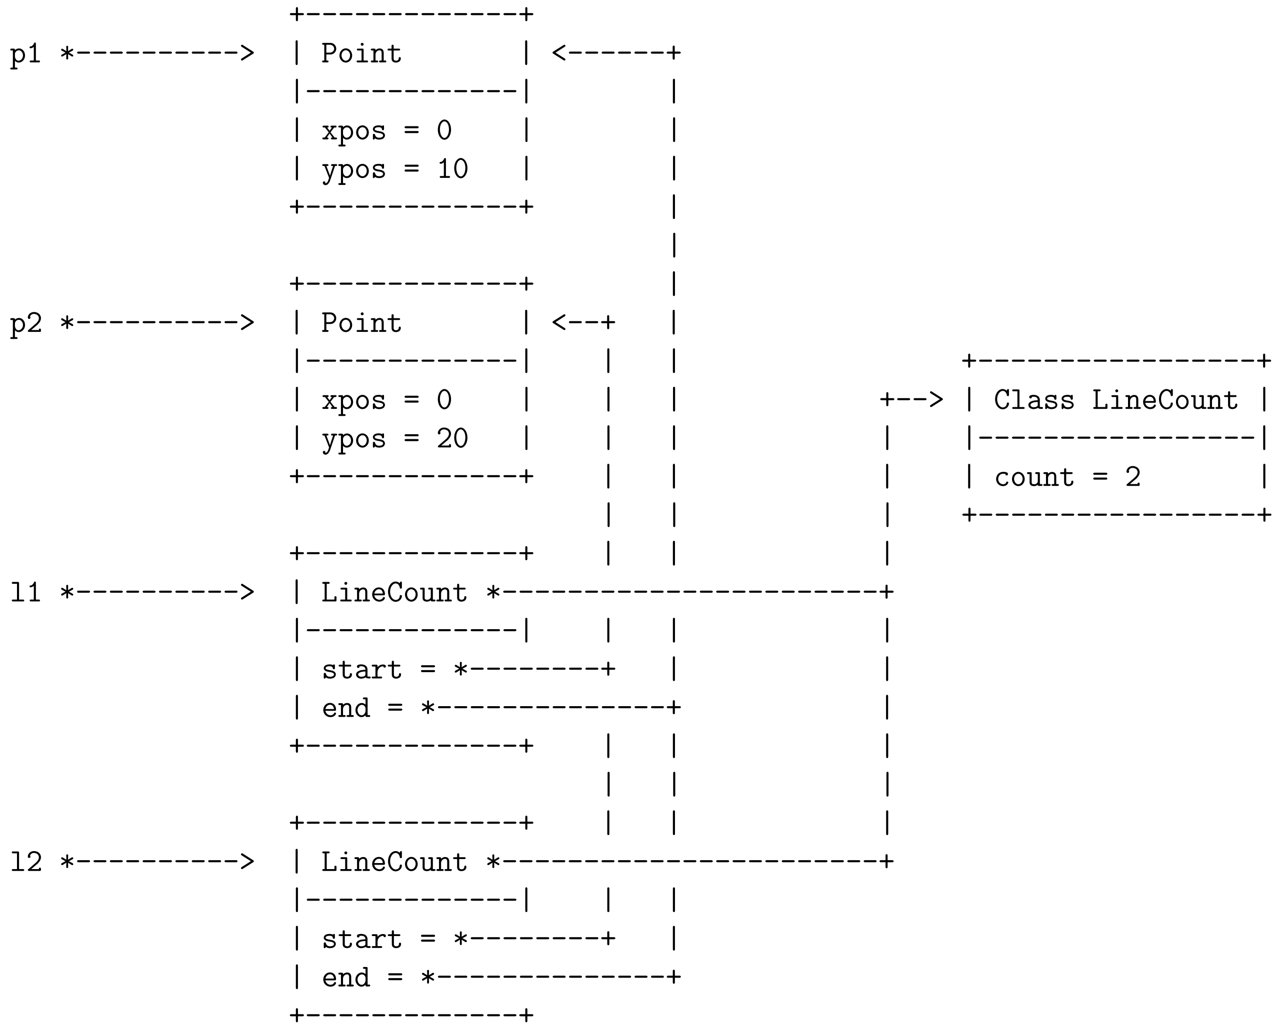

Note that both `l1` and `l2` have a link to the structure corresponding to class `LineCount`, which holds the static variable `count`. So when you have a `LineCount` object and you access its `count` field, Java sees it is not a field in the object itself, so it will follow the "static" pointer to the structure holding the static fields, and look for the field there.

18.4 Inheritance

So that takes care of static fields. What about inheritance? Well, when a class (say `A`) extends another class (say `B`), it allocates not only a chunk of memory to hold the new object of type `A`, but also allocates chunk of memory for an object of type `B` that will hold the fields that are inherited from `B`. Thus, in the presence of inheritance, creating an object actually ends up creating a chain of objects. Consider the following example:

```
public class CLine extends Line {
```

```
  private String color;

  protected CLine (Point s, Point e, String c) {
    super(s,e);
    color = c;
  }

  public static CLine create (Point s, Point e, String c) {
    return new CLine(s,e,c);
  }

  public String color () {
    return color;
}
```

Here's a sample execution with resulting diagram:

```
Point p1 = Point.create(0,10);
Point p2 = Point.create(0,20);
CLine cl = CLine.create(p2,p1,"red");
```

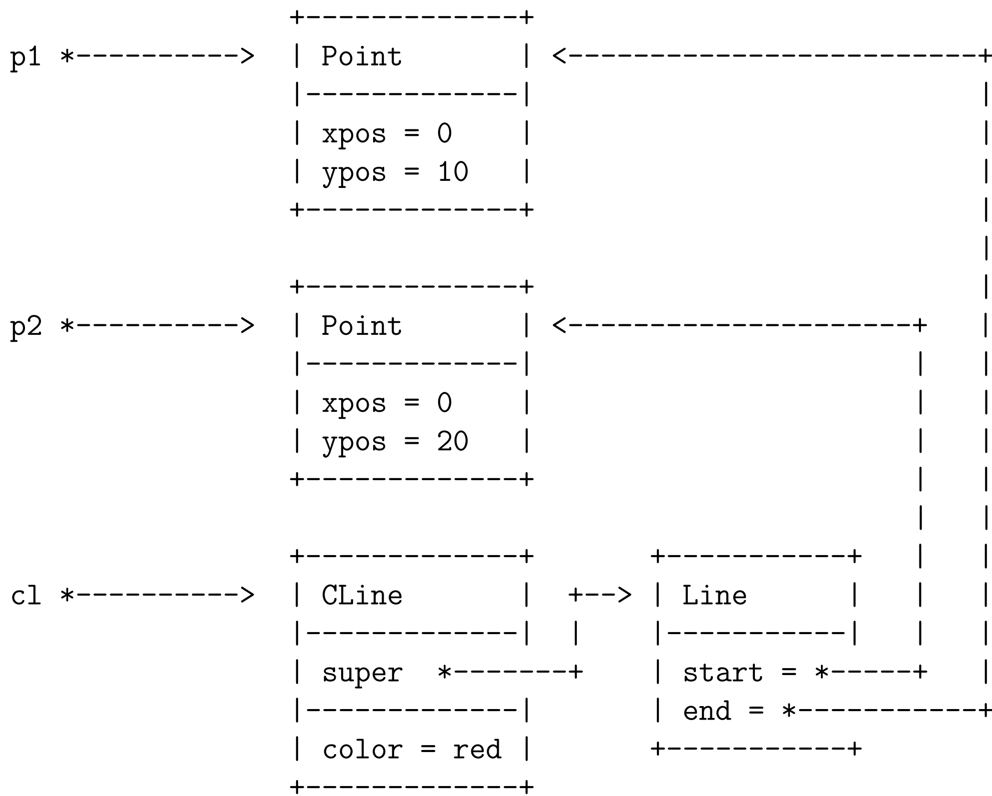

Thus, when we invoke `cl.end().yPos()`, we follow the arrow from `cl` to find the `CLine` object; it does not contain the `end` field, so we follow the `super` arrow that points to the inherited object, which does contain the `end` field, and to get the y-coordinate of that point, we follow the arrow to find the appropriate point object, and extract its `ypos` value.

The above diagram is grossly simplified, because, in particular, every object in Java extends at least the `Object` class. Thus, if you *really* wanted to be accurate, you would have to create inheritance arrows to objects of class `Object`, and so on, for every object. The above model is sufficient for quick back-of-the-envelope computations, though.

18.5 Shallow and Deep Copies

As we saw in the first example, if we pass an object to a method, we are really passing the address of the object to the method. And if the method just takes those values and store them somewhere, then we can get sharing, which may or may not be what we want.

An example of sharing was the `Line.create(p1,p1)` example earlier. The two fields `start` and `end` of the newly created `Line` object end up pointing to the same object, so that modifying one of course leads to modifying the other.

It makes sense sometimes to create a new object that is an exact copy of an object, but does not have any of its sharing. Of course, the best place to put this is as a method to the object itself: "Object, copy thyself!"

Let's write a method for `Point` that creates a new copy of a point:

```
public Point copy () {
  return new Point(xpos, ypos);
}
```

Simple enough. And indeed, if we draw what's happening when we write, say,

```
Point p1 = Point.create(0,10);
Point p2 = p1.copy();
Line l1 = Line.create(p2,p1);
```

we get what we expect, two distinct copies of the same `Point` object with the same field values:

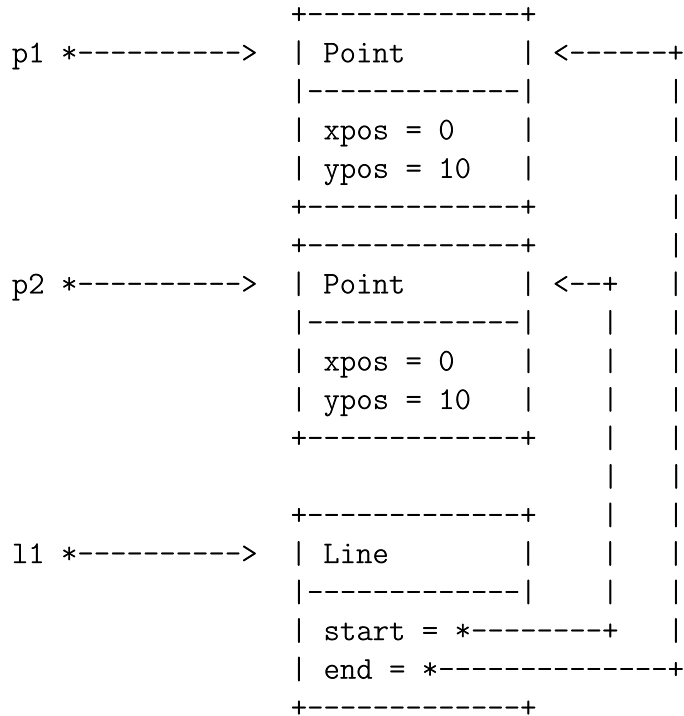

However, if we implement a similar method in `Line`, we get an interesting behavior:

```
public Line copy () {
  return new Line(start,end);
}
```

Let's execute the following code to see what happens:

```
Point p1 = Point.create(0,10);
Point p2 = Point.create(0,20);
Line l1 = Line.create(p2,p1);
Line l2 = l1.copy();
```

If you work through carefully what happens, you get the following:

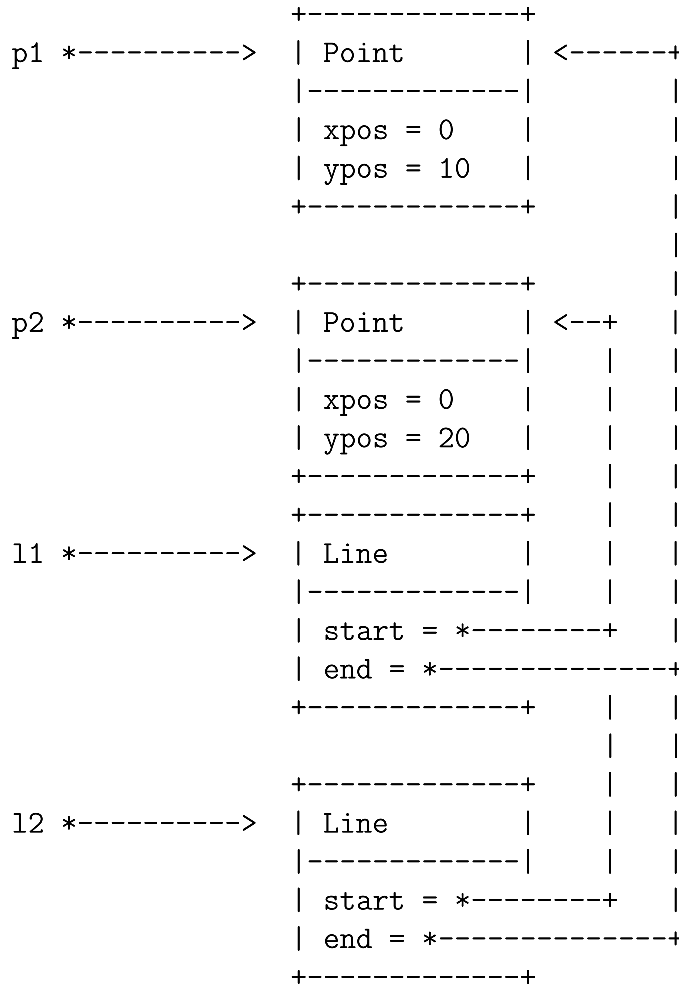

Which may not be *quite* what we want. It does create a fresh copy of `l1`, but because `start` and `end` are simply given the address of the object pointed to by `start` and by `end` in `l1`, the value of the fields end up the same in both `l1` and `l2`.

So while `l2` is a copy of `l1`, they are not fully disjoint. Rather, `l2` is what we call a *shallow copy* of `l1`. The "top level" of the objects are disjoint (in the sense that their fields live in different places), but any sharing within the values held in the fields is preserved.

If we wanted a truly disjoint new line, then we need to make what is called a *deep copy*, that is, a copy that recursively deep copies (creating new objects) for every object held in every variable, all the way down. Let's implement a method `deepCopy`, in `Line`:

```
public Line deepCopy () {
  Point newStart = start.deepCopy();
  Point newEnd = end.deepCopy();
  return new Line(newStart, newEnd);
}
```

So, you see, to create a deep copy of a line, we recursively deep copy all the values of all the relevant fields, and create a new line with those new values.

This also means that we need a `deepCopy` function in `Point`:

```
public Point deepCopy () {
  return new Point(xpos, ypos);
}
```

Pleasantly, this is exactly the same as a shallow copy. That's because copying an integer is the same as creating a new integer. (That's a nice property of primitive types.) So for some classes, a deep copy is going to look the same as a shallow copy. For the sake of clarity, let's keep the two methods distinct, just to make clear that they have different roles, even though they do the same thing.

Now, if we write:

```
Point p1 = Point.create(0,10);
Point p2 = Point.create(0,20);
Line l1 = Line.create(p2,p1);
Line l2 = l1.deepCopy();
```

You get that `l1` and `l2` are truly disjoint: each of their `start` and `end` fields point to *different* objects. I will let you draw the pictures, as a simple exercise.

Exercise: *Reimplement the Point ADT so that instead of storing the coordinates of the x and and y positions of points as integers, it stores them as elements of a new Coordinate ADT with a single creator* `Coordinate create(int)`*, an operation* `int coordinate ()` *to extract the integer coordinate value, and* `void setCoordinate (int)` *to set the coordinate value. Once you have implemented this and modified* `Point` *accordingly, define a* `deepCopy()` *operations for all of* `Coordinate`*,* `Point`*, and* `Line`*, so that deep copying a line gives rise to completely disjoint structures.*