

CSE341, Fall 2011, Lecture 26 Summary

Standard Disclaimer: This lecture summary is not necessarily a complete substitute for attending class, reading the associated code, etc. It is designed to be a useful resource for students who attended class and are later reviewing the material.

This lecture applies the general concepts of record and function subtyping from the last lecture to object-oriented languages, mostly Java. We will see that:

- Record and function subtyping are excellent guides to what should be allowed in statically typed object-oriented languages to prevent “field missing” and “method missing” errors.
- But Java has two key differences worth exploring:
 - It uses class and interface *names* for types, which is convenient and/but prevents some subtyping that would be sound.
 - It does not support contravariant arguments on methods. Instead it supports *static overloading* based on the types of method arguments.
- `self/this` is special with respect to subtyping: It is covariant in subclasses even though we sometimes think of it as an extra argument to methods.

Basic Subtyping for Objects and Java’s Restrictions Thereof

An object is basically a record holding fields (which we assume here are mutable) and methods. We assume the “slots” for methods are immutable: If an object’s method `m` is implemented with some code, then there is no way to mutate `m` to refer to different code. (An instance of a subclass could have different code for `m`, but that’s different than mutating a record field.)

With this perspective, sound subtyping for objects follows from sound subtyping for records and functions:

- A subtype can have extra fields.
- Because fields are mutable, a subtype cannot have a different type for a field.
- A subtype can have extra methods.
- Because methods are immutable, a subtype can have a subtype for a method, which means the method in the subtype can have contravariant argument types and a covariant result type.

That said, object types in Java and C# do not look like record types and function types. For example, we cannot write down a type that looks something like:

```
{fields : x:real, y:real, ...  
  methods: distToOrigin : () -> real, ...}
```

Instead, we reuse class names as types where if there is a class `Foo`, then the type `Foo` includes in it all fields and methods implied by the class definition (including superclasses). And, as discussed previously, we also have interfaces, which are more like record types except they do not include fields and we use the name of the interface as a type. Subtyping in Java and C# includes only the subtyping explicitly stated via the subclass relationship and the interfaces that classes explicitly indicate they implement (including interfaces implemented by superclasses).

All said, this approach is more restrictive than subtyping requires, but since it does not allow anything it should not, it soundly prevents “field missing” and “method missing” errors. In particular:

- A subclass can add fields but not remove them
- A subclass can add methods but not remove them
- A subclass can override a method with a covariant return type
- A class can implement more methods than an interface requires or implement a required method with a covariant return type

Shadowing Fields

Because fields are mutable, we know a subtype cannot have a different type for a field than a supertype. So it is surprising that Java actually allows a class to declare a field with the same name as a field already declared in a superclass. For example, this code is allowed in Java:

```
class Color { String colorName; }
class FancyColor extends Color { double shade; }
class Point { ... }
class ColorPoint extends Point {
    Color color;
    ...
}
class MyColorPoint extends ColorPoint {
    FancyColor color;
    ...
}
```

It would not be sound to allow the `MyColorPoint` class to change the `color` field to hold a `FancyColor` and have code assume this. After all, methods inherited from `ColorPoint` could assign `color` to hold instances of `Color` leading to errors if we then tried to read the `shade` field.

The conundrum is actually easily solved: In Java, the declaration `FancyColor color` adds a *different field* to instances of `MyColorPoint` that happens to have the same name as another field. Since it is a new, different field, it can have any type whatsoever, not necessarily related to the type of the other field named `color`.

Field look-ups do *not* use dynamic dispatch, so any uses of `color` in methods of `ColorPoint` will still use the `color` field of type `Color`. In `MyColorPoint`, the new field shadows the old one, so `color` will refer to the new field. To access the old field, code can use `super.color`.

Static Overloading

A similar but more common and convenient weirdness relates to Java's support for having multiple methods with the same name. Even though it would be sound, there is simply no way in Java to override a method with contravariant arguments. Instead, Java works like this:

If a class declares a method with name `m`, return type `t0` and argument types `t1, t2, ..., tn` (in that order), then:

- If the superclass already has a method name `m` with argument types `t1, t2, ..., tn` (in that order), then this is overriding. The type-checker requires `t0` to be the same as the return type in the superclass method or a subtype (covariant subtyping on result types).
- Else this is a new method added to the subclass, not overriding, even if the superclass (or the subclass) has other methods named `m`.

Having multiple methods with the same name can be convenient so that you do not have to think up different names for similar behavior. For example, a `Rational` class could have methods like:

```
void add(Rational r) {...}
void add(int i) {...}
void add(double d) {...}
```

But you must take care in Java to make sure you are overriding or not overriding when you intend to. It can be confusing to change or reorder argument types accidentally and, as a result, have different methods with the same name.

More importantly, we need to define which method gets called when there are multiple methods with the right name. That is, we must revise our definition of the most fundamental issue in object-oriented programming: What is the semantics of `e0.m(e1, ..., en)`? As before, we start as follows:

- Evaluate `e0, ..., en` to objects `v0, ..., vn`.
- Use the (run-time) class of `v0` to look up `m` (dynamic dispatch).

But now there may be multiple choices `m`. Java picks the “best” choice using the *static types* of `e1, ..., en`, not the classes `v0, ..., vn`. This semantics is called *static overloading* since we overload the name `m` to mean multiple things and disambiguate using the static types of the arguments. It is *not* dynamic dispatch on the arguments, which would be a different semantics called *multimethods*. There are programming languages with multimethods, but Java and C# do not have them.

The exact definition of “best” choice is remarkably complicated and we will not explain all the rules, but we will give the rough idea and some examples. Simple cases like the `add` methods for `Rational` are easy: Type-checking the argument should produce an `int`, `double`, or `Rational` and the “best” choice should be obvious. But even here we allow a subtype of `Rational` (using subtyping to pick the method that takes a `Rational`) or other numeric types (using a conversion to pick either the `int` version or the `double` version according to rules spelled out in the language definition). Also note that we always pick a method with the correct number of arguments. And if there is no choice that would type-check, the call, as usual, does not type-check.

Here is a more complicated example where we have 5 methods named `m` and the comments next to the methods and example method calls indicate which method will get used by each call:

```
class Color extends Object { String s; }
class FancyColor extends Color { double shade; }

class MyClass {
    void m(Object x)      {...} // A
    void m(Color x)       {...} // B
    void m(FancyColor x) {...} // C
    void m(Color x, FancyColor y) {...} // D
    void m(FancyColor x, Color y) {...} // E
}
MyClass obj = new MyClass(...);
Color c1 = new Color(...);
FancyColor c2 = new FancyColor(...);
Color c3 = new FancyColor(...); // subtyping!
obj.m(c1); // B
obj.m(c2); // C
```