

Comp 311
Principles of Programming Languages
Lecture 21
Semantics of OO Languages

Corky Cartwright
Mathias Ricken
October 20, 2010

Overview I

- In OO languages, data values (except for designated non-OO types) are special *records* [*structures*] (finite mappings from *names* to *values*, which are not necessarily objects). In OO parlance, the components of record are called *members*.
- Some *members* of an object may be functions called *methods*. Methods take **this** (the object in question) as an implicit parameter. Some OO languages like Java also support *static* methods that do not depend on **this**; these methods have no implicit parameters.
- Methods (*instance* methods in Java) can only be invoked on objects (the implicit parameter). Additional parameters are optional, depending on whether the method expects them.
- For a language to qualify as OO, it must support *inheritance*, an explicit taxonomy for classifying objects based on their members. In single inheritance, this taxonomy forms a tree; in multiple inheritance, it forms a rooted DAG (directed acyclic graph). Inheritance also provides a simple mechanism for defining some objects as extensions of others.
- Most OO languages are *class*-based (my preference because it supports a simple static type system). In *class*-based OO languages, every object is an instance of a class (an object template)

Overview II

- In *single inheritance* class-based languages, every class must declare its immediate superclass. In *multiple inheritance* class-based languages, every class must declare one or more immediate superclasses. Each superclass is either another declared class or a built-in universal (least common denominator) class [**Object** in Java]. Every class inherits all members of its immediate superclass(es); it has the option of overriding (replacing) the definitions of inherited methods.
- Java does not allow true multiple inheritance but it supports a cleaner alternative (*multiple interface inheritance*) using special classes called *interfaces*.
- The *superclass relation* is the transitive closure of the *immediate superclass relation*.
- A class cannot shadow a method defined in the parent class(es); it can only override it (replace its definition in the current class). The overriding method appears in class instances (objects) in place of the overridden one.
- A class can only shadow a field defined in the parent class; it cannot override it. Shadowing is simply the hiding of the parent field by the new fields exactly as in lexical scoping. The shadowed field still exists, but it can only be accessed by a **super** or by upcasting the type of the receiver (in a typed OO language).
- The method lookup process in OO languages is called *dynamic dispatch*. The meaning of a method call depends on the method code in **this**. In contrast, the meaning of a field reference is fixed for all subclasses of the class where the field is introduced. The field can only be shadowed but that does not affect the meaning of code that cannot see a shadowing definition.

Overview III

- Implications of overriding vs. shadowing: a *method invocation* always refers to the specified method in the receiver object even when the method has a definition in the class where the invocation appears. This mechanism is called *dynamic dispatch*; it is sometimes (misleadingly!) called *dynamic binding*.
- In contrast, *field references* refer to the field determined by lexical scoping rules (the corresponding binding occurrence of the field).
- A static type system can be used to restrict (discipline) class definitions and guarantee for all method lookups that a match will be found.
- OO languages that are *not* class-based are *prototype*-based. Any object can be a prototype (factory) for creating descendant objects.
- In prototype-based OO languages, objects literally contain methods. A lookup failure within a given object triggers a search in the ancestor object instead of the superclass.
- A prototype-based OO program can be written in a disciplined fashion (where a few factory objects function as class surrogates) so they have the same structure as a class-based program *but* type-checking is problematic. Complex static analysis is possible but it is not transparent and not very helpful (IMO) in locating and identifying program bugs.

Overview IV

Thumbnail history:

Simula (1967)

- Allows Algol blocks to be autonomous data values with independent lifetimes → objects
- Classes are special procedures that return blocks
- Allows autonomous blocks to be defined as extensions of other blocks; inheritance = lexical scoping + copying!
- Inheritance is single because it is block extension.
- No conventional overriding (?) but inner mechanism looks like inverse of overriding
- Good software engineering insights but no clear design methodology

Smalltalk (1972)

- Dynamically typed
- Class based
- Supports reflective access to the runtime (much like Java reflection).
- Single inheritance
- Dynamic extension of objects
- If dynamic features are exploited, software engineering is compromised

Self (1987)

- Dynamically typed
- Prototype based
- Activation records are objects (Simula in reverse!)
- Dynamic scoping except for explicit closures (Ugh!)

Aside; devise a statically typed dynamically scoped language. Impossible? AFAIK,

Pedagogic OO Extensions of ML culminating in OCaml; not truly OO because type system interferes.

Pragmatic OO extensions of C: Objective C, C++; truly OO except storage management is manual

Eiffel/Dylan (Scheme with classes and inheritance)/Java/C#/Scala

Important distinction: structural subtyping (ML and other pedagogic extensions) vs. nominal (C++, Java, etc.)

Java as a Real OO language

- Java is most important real OO design.
- Two views:
 - C++ augmented by GC
 - Dylan (Scheme + objects) with a C-like syntax
- I strongly support the latter view. Why? The semantic definitions of C++ and Java are completely different while the semantics of Dylan and Java are very similar. It is easy to transliterate Scheme (sans **call/cc**) into Java. It is essentially impossible to transliterate C++ into Java.

Java Implementation I

Why talk about implementation? In real world languages, implementation dictates semantic choices. Tradeoffs!

Part I: Data representation

- Java objects include a header specifying the object class and hash code. The remaining record components [slots, fields in C parlance] are simply the members of the object. The class “identifier” is a pointer to an object (belonging to class **Class**) describing the class. How can **super** be supported? (For fields, it is trivial. For methods, use **Class** object.)
- Method lookup: simply extract the method from the receiver object. Trivial! In the absence of the method table optimization, inherited methods are simply components [slots] of the object record. But space optimization important.
- Space optimization: move (pointers to) member methods to a separate method table *for each class* which can be shared by all instances of the class. This table is part of the **Class** object for the class where the method definition appears. Note that lookup is now much more complex because only “local” methods (those explicitly defined in the object's class) are defined in the local method table. “Non-local” method references must look at superclasses.

Java Implementation II

- Interfaces can be supported in a variety of ways. Perhaps the simplest is to create a separate interface method table for each class implementing the interface. These tables are linked from the class method table. How can you find the link? Internally support `getInterfaceXXX` methods (dynamic dispatch)
- Observation: interface method dispatch is slower than class method dispatch.
- Fast **instanceof**: naïve implementation requires search (which can be messy if subject type is an interface). Constant time implementations are possible. One simple approach: assign consecutive indices starting at 0 to types (classes/interfaces). Maintain a bit string for each class specifying which types it belongs to. Then `instanceof` simply indexes this bitstring.
- Multiple inheritance in C++ is supported by partitioning an object into sub-objects corresponding to each superclass and referring to objects using “internal” pointers so that a subclass object literally looks like the relevant superclass object. The bookkeeping (pointer adjustment) can get messy. Object pointers do not necessarily point to the base of the object! How can executing code find the base of an object (required by a cast to a different type!)? By embedding a head pointer in each sub-object representation.

Java Implementation III

Part II: Runtime

- The central (control) stack holds activation records for methods starting with `main`. There is *no* static link because Java only supports local variables.
- All objects are stored in the heap. All fields are slots in objects.
- Object values are represented by pointers (to the record representing the objects).
- Object in the heap must be reached (transitively) through local variables on the stack. In compiled code, computed values are often cached in registers.
- Instances of (dynamic) inner classes include a pointer to an enclosing parent object (static link!) so that inner class code can access fields in the enclosing object.

Java Implementation IV

- Classes are loaded dynamically by the class loader; it maps a byte stream in a file to a class object including code for all of the methods. The class loader performs *byte code verification* to ensure the loaded classes are well-formed and type-correct. In Java systems using “exact” GC, the class loader must build stack maps (indicating which words in the current activation record are pointers) for a sufficient set of program “safe points” also called “consistent regions”. There is not a single stack map for each method because local variable usage can vary during method execution! (Allowing this variance was a bad design decision in my opinion!)
- The Java libraries are simply an archive (typically in zip file format) containing a file tree of class files (byte streams).
- Java allows programs to use custom class loaders. Our NextGen compiler supporting first-class generics critically depends on this fact. Do does DrJava (for different reasons).
- Header optimization: use the pointer to the method table as the class “identifier”; method table must contain pointer to the **Class** object. Method table also includes a pointer to the superclass method table.

Java Implementation V

- In the presence of the method table optimization (which is essentially universal), objects only contain header information and fields (in OO parlance).
- Interfaces can be supported in a variety of ways. Perhaps the simplest is to create a separate interface method table for each class implementing the interface. These tables are linked from the class method table using dynamic dispatch (essentially a hidden `getInterfaceXXXTable()` method).
- Fast **instanceof**: naïve implementation requires search (which can be messy if subject type is an interface). Constant time implementations are possible. One simple approach: assign consecutive indices starting at 0 to types (classes/interfaces). Maintain a bit string for each class specifying which types it belongs to. Then `instanceof` simply indexes this bitstring. Must be updated when a new class is loaded (if class indices monotonically increase, can use demand-driven updates based on index out-of-bounds exception).

Java Criticisms

- Not truly OO:
 - Values of primitive types are not objects (should be hidden)
 - Static fields and methods (useful in practice just like mutation functional languages)
- Interfaces are not fully satisfactory as replacement for multiple inheritance.

Interfaces should be generalized to “traits” which are classes with no fields. The complexity of multiple inheritance is due to the fact that the same field can be inherited in multiple ways (the “diamond” relationship). This pathology cannot occur in multiple trait inheritance.
- Type system is too baroque and too restrictive.

Generic (parameterized) nominal type systems are still not fully understood. When Java was invented, nominal typing was still a radical idea (in Eiffel, C++, and Objective C which were not type safe! Eiffel subsequently added an ugly runtime check to salvage type safe).
- Excessive generalization of some constructs and mechanisms leads to baroque language specification.

Examples: `<receiver>.new <type>(...),`
`new <InnerClassType>(...)` outside of enclosing class
excessively general local type inference for polymorphic methods
unrestricted wildcard types (wildcards as bounds!)
- Erasure based implementation of generic types.

A huge (!) mistake IMO but note that I am clearly biased.
- The run-time check in array element updating is awkward.

The designers wanted co-variant subtyping for arrays ($u \leq v$ implies $u[] \leq v[]$) which is important in absence of generic types. Co-variant subtyping is difficult to support in an OO language (because method input types behave contra-variantly in subtyping relationships!). Java 5/6 uses wildcard types to support co-variance but did not get the details right. C# does not support any covariance.