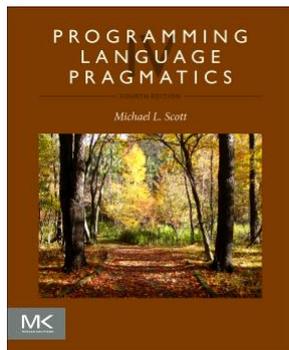


Chapter 10 :: Data Abstraction and Object Orientation

Programming Language Pragmatics, Fourth Edition

Michael L. Scott



ELSEVIER

Object-Oriented Programming

- Control or PROCESS abstraction is a very old idea (subroutines!), though few languages provide it in a truly general form (Scheme comes close)
- Data abstraction is somewhat newer, though its roots can be found in Simula67
 - An Abstract Data Type is one that is defined in terms of the operations that it supports (i.e., that can be performed upon it) rather than in terms of its structure or implementation



Object-Oriented Programming

- Why abstractions?
 - easier to think about - hide what doesn't matter to programmer
 - protection - prevent access to things you shouldn't see
 - plug compatibility
 - replacement of pieces, often without recompilation, definitely without rewriting libraries
 - division of labor in software projects



Object-Oriented Programming

- We talked about data abstraction some back in the unit on naming and scoping (Ch. 3)
- Recall that we traced the historical development of abstraction mechanisms
 - Static set of variables Basic
 - Locals Fortran
 - Statics Fortran, Algol 60, C
 - Modules Modula-2, Ada 83
 - Module types Euclid
 - Objects Smalltalk, C++, Eiffel, Java
Oberon, Modula-3, Ada 95



Object-Oriented Programming

- *Statics* allow a subroutine to retain values from one invocation to the next, while hiding the name in-between
- *Modules* allow a collection of subroutines to share some statics, still with hiding
 - If you want to build an abstract data type, though, you have to make the module a manager



Object-Oriented Programming

- *Module types* allow the module to *be* the abstract data type - you can declare a bunch of them
 - This is generally more intuitive
 - It avoids explicit object parameters to many operations
 - One minor drawback: If you have an operation that needs to look at the innards of two different types, you'd define both types in the same manager module in Modula-2
 - In C++ you need to make one of the classes (or some of its members) "friends" of the other class



Object-Oriented Programming

- Objects add inheritance and dynamic method binding
- Simula 67 introduced these, but didn't have data hiding
- The 3 key factors in OO programming
 - Encapsulation (data hiding)
 - Inheritance
 - Dynamic method binding



Encapsulation and Inheritance

- Visibility rules
 - Public and Private parts of an object declaration/definition
 - 2 reasons to put things in the declaration
 - so programmers can get at them
 - so the compiler can understand them
 - At the very least the compiler needs to know the size of an object, even though the programmer isn't allowed to get at many or most of the fields (members) that contribute to that size
 - That's why private fields have to be in declaration



Encapsulation and Inheritance Classes (C++)

- C++ distinguishes among
 - public class members
 - accessible to anybody
 - protected class members
 - accessible to members of this or derived classes
 - private
 - accessible just to members of this class
- A C++ structure (*struct*) is simply a class whose members are public by default
- C++ base classes can also be public, private, or protected



Encapsulation and Inheritance

Classes (C++)

- Example:

```
class circle : public shape { ...  
anybody can convert (assign) a circle* into a shape*
```

```
class circle : protected shape { ...  
only members and friends of circle or its derived classes  
can convert (assign) a circle* into a shape*
```

```
class circle : private shape { ...  
only members and friends of circle can convert (assign) a  
circle* into a shape*
```



Encapsulation and Inheritance

Classes (C++)

- Disadvantage of the module-as-manager approach: must include explicit create/initialize and destroy/finalize routines for every abstraction
 - Even w/o dynamic allocation inside module, users don't have necessary knowledge to do initialization
 - Ada 83 is a little better here: you can provide initializers for pieces of private types, but this is NOT a general approach
 - Object-oriented languages often give you constructors and maybe destructors
 - Destructors are important primarily in the absence of garbage collection



Encapsulation and Inheritance Classes (C++)

- A few C++ features you may not have learned:
 - classes as members

```
foo::foo (args0) : member1 (args1),  
member2 (args2) { ...
```

args1 and args2 need to be specified in terms of args0
 - The reason these things end up in the header of foo is that they get executed *before* foo's constructor does, and the designers consider it good style to make that clear in the header of foo::foo

Encapsulation and Inheritance Classes (C++)

- A few C++ features (2):

- initialization v. assignment

```
foo::operator= (&foo) v.
```

```
foo::foo (&foo)
```

```
    foo b;
```

```
    foo f = b;
```

```
        // calls constructor
```

```
    foo b, f;
```

```
        // calls no-argument
```

```
constructor
```

```
    f = b;
```

```
        // calls operator=
```



Encapsulation and Inheritance Classes (C++)

- A few C++ features (3):
 - virtual functions (see the next dynamic method binding section for details):

Key question: if child is derived from parent and I have a `parent* p` (or a `parent& p`) that points (refers) to an object that's actually a child, what member function do I get when I call `p->f` (`p . f`)?

 - Normally I get p's f, because p's type is `parent*`.
 - But if f is a virtual function, I get c's f.



Encapsulation and Inheritance Classes (C++)

- A few C++ features (4):
 - virtual functions (continued)
 - If a virtual function has a "0" body in the parent class, then the function is said to be a *pure* virtual function and the parent class is said to be *abstract*
 - You can't declare objects of an abstract class; you have to declare them to be of derived classes
 - Moreover any derived class *must* provide a body for the pure virtual function(s)
 - multiple inheritance in Standard C++ (see next)
 - friends
 - functions
 - classes



Initialization and Finalization

- In Section 3.2, we defined the lifetime of an object to be the interval during which it occupies space and can hold data
 - Most object-oriented languages provide some sort of special mechanism to *initialize* an object automatically at the beginning of its lifetime
 - When written in the form of a subroutine, this mechanism is known as a *constructor*
 - A constructor does not allocate space
 - A few languages provide a similar *destructor* mechanism to *finalize* an object automatically at the end of its lifetime



Initialization and Finalization

Issues

- choosing a constructor
- references and values
 - If variables are references, then every object must be created explicitly - appropriate constructor is called
 - If variables are values, then object creation can happen implicitly as a result of elaboration
- execution order
 - When an object of a derived class is created in C++, the constructors for any base classes will be executed before the constructor for the derived class
- garbage collection



Dynamic Method Binding

- Virtual functions in C++ are an example of *dynamic method binding*
 - you don't know at compile time what type the object referred to by a variable will be at run time
- Simula also has virtual functions (all of which are abstract)
- In Smalltalk, Eiffel, and Java *all* member functions are virtual



Dynamic Method Binding

- Note that inheritance does not obviate the need for generics
 - You might think: “hey, I can define an abstract list class and then derive `int_list`, `person_list`, etc. from it”, but the problem is you won't be able to talk about the elements because you won't know their types
 - That's what generics are for: abstracting over types
- Java doesn't have generics, but it does have (checked) dynamic casts



Dynamic Method Binding

- Data members of classes are implemented just like structures (records)
 - With (single) inheritance, derived classes have extra fields at the end
 - A pointer to the parent and a pointer to the child contain the same address - the child just knows that the struct goes farther than the parent does



Dynamic Method Binding

- Non-virtual functions require no space at run time; the compiler just calls the appropriate version, based on type of variable
 - Member functions are passed an extra, hidden, initial parameter: *this* (called *current* in Eiffel and *self* in Smalltalk)
- C++ philosophy is to avoid run-time overhead whenever possible (Sort of the legacy from C)
 - Languages like Smalltalk have (much) more run-time support



Dynamic Method Binding

- Virtual functions are the only thing that requires any trickiness (Figure 10.3)
 - They are implemented by creating a dispatch table (*vtable*) for the class and putting a pointer to that table in the data of the object
 - Objects of a derived class have a different dispatch table
 - In the dispatch table, functions defined in the parent come first, though some of the pointers point to overridden versions
 - You could put the whole dispatch table in the object itself, saving a little time, but potentially wasting a LOT of space



Dynamic Method Binding

```
class foo {  
    int a;  
    double b;  
    char c;  
public:  
    virtual void k( ...  
    virtual int l( ...  
    virtual void m();  
    virtual double n( ...  
    ...  
} F;
```

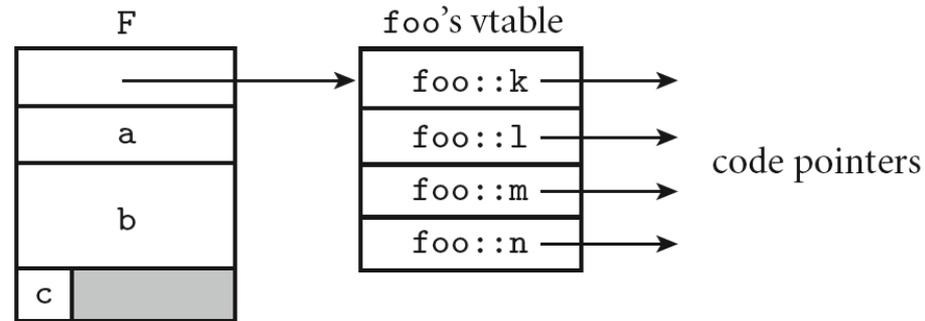


Figure 10.3 Implementation of virtual methods. The representation of object `F` begins with the address of the vtable for class `foo`. (All objects of this class will point to the same vtable.) The vtable itself consists of an array of addresses, one for the code of each virtual method of the class. The remainder of `F` consists of the representations of its fields.

Dynamic Method Binding

```
class bar : public foo {  
    int w;  
public:  
    void m() override;  
    virtual double s( ...  
    virtual char *t( ...  
    ...  
} B;
```

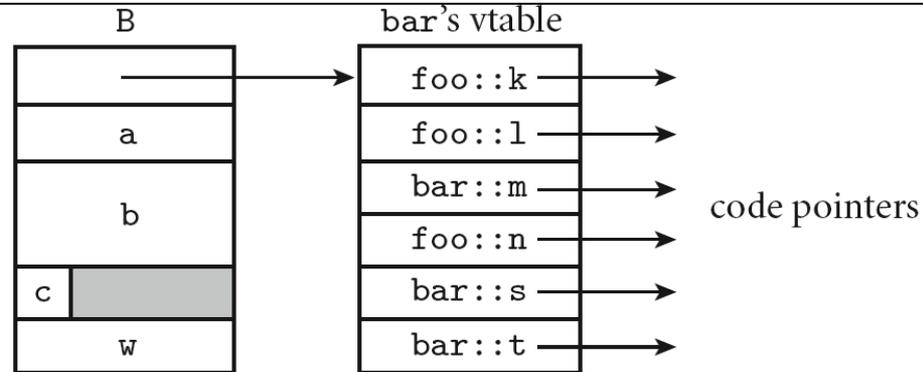


Figure 10.4 Implementation of single inheritance. As in Figure 10.3, the representation of object B begins with the address of its class's vtable. The first four entries in the table represent the same members as they do for `foo`, except that one—`m`—has been overridden and now contains the address of the code for a different subroutine. Additional fields of `bar` follow the ones inherited from `foo` in the representation of B; additional virtual methods follow the ones inherited from `foo` in the vtable of class `bar`.

Dynamic Method Binding

- Note that if you can query the type of an object, then you need to be able to get from the object to run-time type info
 - The standard implementation technique is to put a pointer to the type info at the beginning of the vtable
 - Of course you only have a vtable in C++ if your class has virtual functions
 - That's why you can't do a `dynamic_cast` on a pointer whose static type doesn't have virtual functions



Mix-In Inheritance

- Classes can inherit from only one “real” parent
- Can “mix in” any number of interfaces, simulating multiple inheritance
- Interfaces appear in Java, C#, Go, Ruby, etc.
 - contain only abstract methods, no method bodies or fields
- Has become dominant approach, superseding true multiple inheritance



Mix-In Inheritance

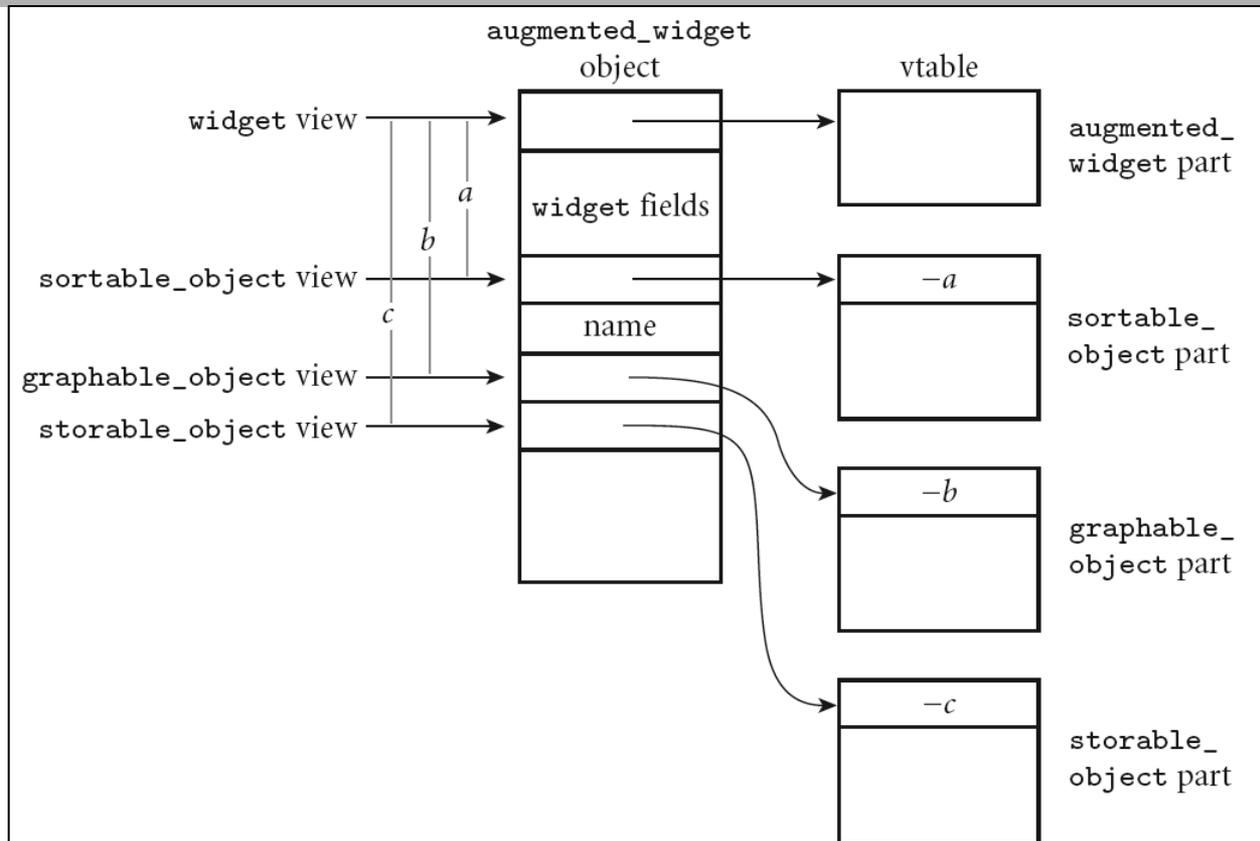


Figure 10.7 Implementation of mix-in inheritance. Objects of class `augmented_widget` contain four vtable addresses, one for the class itself (as in Figure 10.3), and three for the implemented interfaces. The view of the object that is passed to interface routines points directly at the relevant vtable pointer. The vtable then begins with a "this correction" offset, used to regenerate a pointer to the object itself.



True Multiple Inheritance

- In C++, you can say

```
class professor : public
teacher, public researcher {
    ...
}
```

Here you get all the members of teacher *and* all the members of researcher

- If there's anything that's in both (same name and argument types), then calls to the member are ambiguous; the compiler disallows them



True Multiple Inheritance

- You can of course create your own member in the merged class

```
    professor::print () {  
        teacher::print ();  
        researcher::print (); ...  
    }
```

Or you could get both:

```
    professor::tprint () {  
        teacher::print ();  
    }  
    professor::rprint () {  
        researcher::print ();  
    }
```



True Multiple Inheritance

- Virtual base classes: In the usual case if you inherit from two classes that are both derived from some other class B, your implementation includes two copies of B's data members
- That's often fine, but other times you want a *single* copy of B
 - For that you make B a virtual base class



Object-Oriented Programming

- Analogy to the real world is central to the OO paradigm - you think in terms of *real-world* objects that interact to get things done
- Many OO languages are strictly sequential, but the model adapts well to parallelism as well
- Strict interpretation of the term
 - uniform data abstraction - everything is an object
 - inheritance
 - dynamic method binding



Object-Oriented Programming

- Lots of conflicting uses of the term out there
object-oriented *style* available in many languages
 - data abstraction crucial
 - inheritance required by most users of the term OO
 - centrality of dynamic method binding a matter of dispute



Object-Oriented Programming

- SMALLTALK is, historically, the canonical object-oriented language
 - It has all three of the characteristics listed above
 - It's based on the thesis work of Alan Kay at Utah in the late 1960's
 - It went through 5 generations at Xerox PARC, where Kay worked after graduating
 - Smalltalk-80 is the current standard



Object-Oriented Programming

- Other languages are described in what follows:
- Modula-3
 - single inheritance
 - all methods virtual
 - no constructors or destructors
- Java
 - interfaces, *mix-in* inheritance
 - all methods virtual



Object-Oriented Programming

- Ada 95

- *tagged* types

- single inheritance

- no constructors or destructors

- *class-wide* parameters:

- methods static by default

- can define a parameter or pointer that grabs the object-specific version of all methods

- base class doesn't have to decide what will be virtual

- notion of child packages as an alternative to friends



Object-Oriented Programming

- Is C++ object-oriented?
 - Uses all the right buzzwords
 - Has (multiple) inheritance and generics (templates)
 - Allows creation of user-defined classes that look just like built-in ones
 - Has all the low-level C stuff to escape the paradigm
 - Has friends
 - Has static type checking



Object-Oriented Programming

- In the same category of questions:
 - Is Prolog a logic language?
 - Is Common Lisp functional?
- However, to be more precise:
 - Smalltalk is really pretty purely object-oriented
 - Prolog is primarily logic-based
 - Common Lisp is largely functional
 - C++ can be used in an object-oriented style

