

# CS4120/4121/5120/5121—Spring 2016

## oXi Language Specification

Cornell University  
Version of May 15, 2016

Thanks to the hard work and brilliant engineering of your compiler-building team, Xi has been gaining market share. In response to customer demands for object-oriented features, a new object-oriented version of the language, called oXi, has been partly designed. Your task is to extend your Xi implementation with the new object-oriented features, and to add some language extension of your own design.

oXi is backward compatible with Xi, so this language description focuses on the differences. The [original Xi spec](#) is still available.

## 0 Changes

- 5/11: spec change: interfaces can now begin with use statements.

## 1 Class definitions

An oXi program may contain class definitions in addition to function definitions. A class definition may contain instance variable (field) definitions, and method definitions. For example, the following code defines a `Point` class and an associated creator function `createPoint`:

```
1 class Point { // a mutable point
2     x,y: int
3
4     move(dx: int, dy: int) {
5         x = x + dx
6         y = y + dy
7     }
8     coords() : int, int {
9         return x, y
10    }
11    add(p: Point) : Point {
12        return createPoint(x + p.x, y + p.y)
13    }
14    initPoint(x0: int, y0: int): Point {
15        x = x0
16        y = y0
17        return this
18    }
19    clone(): Point { return createPoint(x, y) }
20    equals(p: Point): bool { return this == p }
21 }
22
23 createPoint(x: int, y: int): Point {
24     return new Point.initPoint(x, y)
25 }
```

As in Java, there is a special variable `this` that refers to the method receiver. The instance variables and methods of `this` are automatically in scope within methods.

Notice that the fields `x` and `y` can both be declared at once to have the type `int`. The same abbreviated syntax can now be used for local variable declarations, though no initializer expression may be provided in that case.

## 2 Class declarations

Classes do not have visibility modifiers. All class members, including instance variables, are visible everywhere inside their module (i.e., source file). To be used from a different module, the class must be declared in an interface that includes all the methods of the class<sup>1</sup>. Instance variables are always private and cannot be mentioned in the interface. For example, we might define an interface file for the Point class, hiding the x and y fields:

```
1 // A 2D Point with integer coordinates (x,y).
2 class Point {
3     move(dx: int, dy: int)
4     add(p: Point): Point
5     coords(): int, int
6     clone(): Point
7
8     // Initialize this to contain (x,y).
9     // Requires: p is uninitialized.
10    // Returns: this
11    initPoint(p: Point, x: int, y: int): Point
12 }
13
14 // Create the point (x,y).
15 createPoint(x: int, y: int): Point
```

A class may inherit from one other class, with an extends clause. For example, we might declare a subclass ColoredPoint that inherits from Point:

```
1 class Color {
2     r,g,b: int
3 }
4
5 class ColoredPoint extends Point {
6     col: Color
7     color(): Color { return col }
8
9     initColoredPoint(x0: int, y0: int, c: Color): ColoredPoint {
10        col = c
11        _ = initPoint(x0, y0)
12        return this
13    }
14 }
```

There are no class variables or class methods (Java's static) in oXi, because ordinary functions as in Xi, and the newly introduced global variables, can be used instead.

## 3 Modules and interfaces

The extension .xi continues to be used for files containing module definitions and the extension .ixi is used for interface files. Therefore, the statement use my\_module appearing in a module causes the compiler to look for the interface in the file my\_module.ixi.

The syntax of interfaces has also been extended so that, like modules, they can begin with "use" statements. If a module uses an interface, it also implicitly uses every interface that is used by that interface. It is legal for an interface to be used by a module along multiple "use" paths.

---

<sup>1</sup>The lack of private methods makes it possible to lay out dispatch tables at compile time

If a module is defined in a file `a_module.xi`, the compiler automatically looks for an interface in file `a_module.ixi`, exactly as though the module had contained a statement `use a_module.ixi`. However, such an interface need not exist.

If a module that defines a class `C` references (either explicitly or implicitly) an interface that declares class `C`, the class definition must match the interface declaration. It must implement all of the methods that the interface declares, and no additional methods. If such an interface exists, the order of the methods (in the object's dispatch table) is as declared in the interface, not as defined in the module. Similarly, if a module references an interface that declares a procedure or function that is defined in the module, the signature of the procedure or function declared in the interface must match the definition in the module.

A module does not need to have an interface. Further, even if there is an interface, not every class, procedure, or function in the module needs to be declared in that interface. Undeclared components are analogous to private classes or private methods in Java because they are not visible outside their own module.

Globals, including classes, functions, and global variables, are in scope throughout their defining module and throughout any module that uses an interface that declares them.

## 4 Subtyping and conformance

The subtyping relationship of `Xi` is extended to classes in a straightforward way. Every class name can be used as a type. A class is a subtype of the class it extends, if any. It is therefore also a subtype of any supertypes of that class. There is no top class in the subtype ordering (no `Object`).

Classes must conform to their declared superclasses. They may override methods, but if they do, the new method signature must match the signature in the superclass. They may not declare methods whose names shadow those in superclasses.

The subtyping rule on arrays is unchanged—if classes `B` and `C` extend class `A`, you cannot use a value of type `B[]` at type `A[]`, since then you could put a `C` into it and cause a run-time type error.

## 5 Initialization

Toward type safety, all storage is initialized before possible use in `oXi`. Local variables, array elements, and object fields are initialized with a default value determined by their types<sup>2</sup>. The language implementation may choose to elide this initialization if the storage is always overwritten with another value before use.

### 5.1 Null

Like C and Java, `oXi` incorporates Hoare's "[Billion-Dollar Mistake](#)". The special value `null` is a member of all class types and also a member of all array types. It is the default initialization value for all variables with class or array type, including global variables and array elements.

It is a run-time error to perform any operation on `null` that is specific to objects or arrays. The value `null` must be implemented as a pointer to memory location 0. A reference to this memory location will cause a page fault that will halt this program. This is an adequate implementation of the run-time checking for null values.

### 5.2 Other defaults

The default initialization values for `int` and `bool` are 0 and `false`, respectively.

---

<sup>2</sup>The memory returned by `_I_allloc_i` is automatically initialized to zero, as are global variables in section `.bss`

## 6 New operators

A new object is created with the syntax `new C`. This can only be done inside the module where `C` is defined—standalone creator operations must be implemented as functions like `createPoint` above. As the `Point` example shows, this operator has higher precedence than “.” does.

Inside the definition of the class `C`, the equality comparison `==` and the inequality comparison `!=` can be used on a value of type `C`. These serve as pointer comparisons, much like in Java, though ones that are private to `C`. Such comparisons are legal as long as one of two operands is of type `C`. It is a static error to compare an object and an array in this fashion. However, the special value `null` can always be compared against any value with object or array type, even outside a class definition.

## 7 Non-OO extensions

### 7.1 Global variables

The lack of global variables in Xi has been remedied in oXi. A module can contain global variable declarations, such as:

```
1 center: Point
2 corner: Point
3 len: int = 10
4 tenpoints: Point[len]
```

Integer variables may be initialized to an integer literal, and boolean variables may be initialized to a boolean literal. Global arrays with constant length can be declared, as shown, and are initialized with default element values. The length may be given either as an integer literal or as the name of a global variable. Global multidimensional array declarations are also supported.

Global variables cannot be declared in interfaces, so they are private to a given module. Two different global variables in different modules are different global variables even if they happen to be declared with the same name.

### 7.2 Assignment

To make type-checking convenient, we introduce a new judgment  $\Gamma \vdash e : \tau$  lvalue, which means that  $e$  is an expression that can be assigned values of type  $\tau$  (an *lvalue*). With this rule, we can rewrite the (ASSIGN) rule in a more generic form that encompasses the old rule and also the old (ARRASSIGN) rule:

$$\frac{\Gamma \vdash e_1 : \tau \text{ lvalue} \quad \Gamma \vdash_u e_2 : \tau}{\Gamma \vdash e_1 = e_2 : \Gamma} \text{ (ASSIGN)}$$

The lvalues are variables, array elements, and fields:

$$\frac{\Gamma(x) = \text{var } \tau}{\Gamma \vdash x : \tau \text{ lvalue}} \text{ (VARLVALUE)} \quad \frac{\Gamma \vdash e_1 : \tau[] \quad \Gamma \vdash_u e_2 : \text{int}}{\Gamma \vdash e_1[e_2] : \tau \text{ lvalue}} \text{ (ARRLVALUE)}$$

$$\frac{\Gamma \vdash_u e_1 : C \quad \text{class } C \{ \dots f : \tau \dots \} \in \Gamma}{\Gamma \vdash e_1.f : \tau \text{ lvalue}} \text{ (FIELDLVALUE)}$$

### 7.3 Break statement

To make it easier to end loops, a `break` statement similar to that in C and Java has been added. As in these languages, and its effect is to immediately terminate the closest lexically enclosing loop. As with the `return` statement, its type is `void`, but the `break` statement is legal only if lexically enclosed by a loop.


## 8 ABI

To give implementers flexibility, the ABI for oXi specifies as little as is required for different implementations to interoperate, with implementations otherwise free to lay out data as desired.

### 8.1 Method calls

The first common need is dynamic dispatch of method calls. For this, objects must specify pointers to their classes' dispatch tables at their first memory location, with the tables laid out as follows:

Starting with the class at the top of the inheritance hierarchy, and moving down towards the most concrete class, first allocate a private slot for the use of whatever compiler built that class, then a pointer for each method in the order they are specified in the interface declaration. For example, in the ColoredPoint example above, the layout must look like this:



When a method is invoked, the reference to the receiver object is passed as if it was the first argument, before the actual arguments, but after any hidden argument used to return multiple results. Object references are stored in arrays and are passed to and returned from functions in the same way as other scalar types like `int` and `bool`. Top-level functions that take or return object reference should encode their types into method signatures as follows:

1. `o`
2. A number giving the length of the unescaped type name.
3. The name, with underscores escaped in the usual fashion.

For example, the method `average(a: Point, b: Point): Point` would have its name encoded as `_Iaverage_o5Pointo5Pointo5Point`. The naming of symbols for methods is implementation-specific since they cannot be called from a different compilation unit by name, only via dispatch tables.

### 8.2 Fields

Fields of objects are laid out in the order they are declared, with each field taking up one 64-bit word.

### 8.3 Global variables

The symbol names for global variables should be encoded as follows:

1. `_I_g_`
2. The name of the variable, with underscores escaped.
3. `_`

#### 4. Encoding of the variable's type.

For example, a variable `points` of type `QPoint[]` will be encoded as `_I_g_points_a06QPoint`.

### 8.4 Initialization

Notice that in order to allocate objects of type `ColoredPoint`, the size of objects of type `Point` must be known, but it may not be available during `ColoredPoint`'s compilation. Similarly, if `ColoredPoint` does not override some methods from `Point`, it needs to copy pointers to them from `Point`'s dispatch vector into its own.

Because of this, object sizes and dispatch vectors are to be computed at application startup time. The size of an object of type `someClass`, including areas for superclasses and the dispatch vector pointer, should be stored in the `_I_size_someClass` variable, while its dispatch vector should be stored under the `_I_vt_someClass` symbol, with underscores in the class name escaped under usual rules.

The size variables are initially set to 0<sup>3</sup> to denote that the size information and the dispatch vector for the given class have not yet been computed. In that case, the function `_I_init_someClass()` is expected to fully compute the size and dispatch vector information. When initializing its own dispatch vector, a subclass must copy pointers for all the methods it does not override, as well as all of its superclass' private class information pointers into the appropriate slots in its dispatch vector. An implementation is expected to avoid computing object sizes and dispatch vector more than once.

You may arrange for initialization functions to be called at startup by placing their addresses into the `.ctors` section of the object file; see `examples/init.s` in the runtime distribution for an example. The order of invocation of these initializers is not specified, however; and therefore any superclasses must have their initialization functions called recursively if necessary.

---

<sup>3</sup>An implementation may set them to the correct value if it is able to compute both the size and the dispatch vector fully at compile time