

# Self: The Power of Simplicity\*

David Ungar<sup>†</sup>  
Randall B. Smith<sup>‡</sup>

SMLI-TR-94-30


December 1994

## Abstract:

Self is an object-oriented language for exploratory programming based on a small number of simple and concrete ideas: prototypes, slots, and behavior. Prototypes combine inheritance and instantiation to provide a framework that is simpler and more flexible than most object-oriented languages. Slots unite variables and procedures into a single construct. This permits the inheritance hierarchy to take over the function of lexical scoping in conventional languages. Finally, because Self does not distinguish state from behavior, it narrows the gaps between ordinary objects, procedures, and closures. Self's simplicity and expressiveness offer new insights into object-oriented computation.

\* This work was partially supported by Xerox Corporation, and partially by National Science Foundation Presidential Young Investigator Grant #CCR-8657631, Sun Microsystems, Inc., The Powell Foundation, International Business Machines Corporation, Apple Computer, Inc., Digital Equipment Corporation, NCR Corporation, Texas Instruments, Inc., and Cray Laboratories.

† At the time of original publication, Ungar was employed at the Computer Systems Laboratory of Stanford University, Stanford, CA, and Smith was employed at Xerox Palo Alto Research Center, Palo Alto, CA.

 *Sun Microsystems*  
*Laboratories, Inc.*  
A Sun Microsystems, Inc. Business  
M/S 29-01  
2550 Garcia Avenue  
Mountain View, CA 94043

**email addresses:**  
david.ungar@eng.sun.com  
randall.smith@eng.sun.com

# Self: The Power of Simplicity

*David Ungar*

*Randall B. Smith*

Sun Microsystems Laboratories  
2550 Garcia Avenue  
Mountain View, CA 94043

To thine own self be true....  
—William Shakespeare

## 1 Introduction

Object-oriented programming languages are gaining acceptance, partly because they offer a useful perspective for designing computer programs. However, they do not all offer exactly the same perspective; there are many different ideas about the nature of object-oriented computation. In this paper, we present Self, a programming language with a new perspective on objects and message passing. Like the Smalltalk-80<sup>TM</sup> language [6], Self is designed to support exploratory programming [13], and therefore includes runtime typing (i.e., no type declarations) and automatic storage reclamation.

But unlike Smalltalk, Self includes neither classes nor variables. Instead, Self has adopted a prototype metaphor for object creation [2, 3, 4, 8, 10]. Furthermore, while Smalltalk and most other object-oriented languages support variable access as well as message passing, Self objects access their state information by sending messages to “self,” the receiver of the current message. Naturally, this results in many messages sent to “self,” and the language is named in honor of these messages.

---

<sup>1</sup>Smalltalk-80 is a trademark of ParcPlace Systems. In this paper, the term “Smalltalk” will be used to refer to the Smalltalk-80 programming language.

One of the strengths of object-oriented programming lies in the uniform access to different kinds of stored and computed data, and the ideas in Self result in even more uniformity, which results in greater expressive power. We believe that these ideas offer a new and useful view of object-oriented computation.

Several principles have guided the design of Self:

**Messages-at-the-bottom.**

Self features message passing as the fundamental operation, providing access to stored state solely via messages. There are no variables, merely slots containing objects that return themselves. Since Self objects access state solely by sending messages, message passing is more fundamental to Self than to languages with variables.

**Occam's razor.**

Throughout the design, we have aimed for conceptual economy:

- As described above, Self's design omits classes and variables. Any object can perform the role of an instance or serve as a repository for shared information.
- There is no distinction between accessing a variable and sending a message.
- As in Smalltalk, the language kernel has no control structures. Instead, closures and polymorphism support arbitrary control structures within the language.
- Unlike Smalltalk, Self objects and procedures are woven from the same yarn by representing procedures as prototypes of activation records. This technique allows activation records to be created in the same way as other objects, by cloning prototypes. In addition to sharing the same model of creation, procedures also store their variables and maintain their environment information the same way as ordinary objects, as described in Section 4.

**Concreteness.**

Our tastes have led us to a metaphor whose elements are as concrete as possible [14, 15]. So, in the matter of classes versus prototypes, we have chosen to try prototypes. This makes a basic difference in the way that new objects are created. In a class-based language, an object would be created by instantiating a plan in its class. In a prototype-based language like Self, an object would be created by cloning (copying) a prototype. In Self, any object can be cloned.

The remainder of the paper describes Self in more detail, and concludes with two examples. We use Smalltalk as our yardstick, as it is the most widely known language in which everything is an object. Familiarity with Smalltalk will therefore be helpful to the reader.

## **2 Prototypes: Blending Classes and Instances**

In Smalltalk, unlike C++, Simula, Loops, or Ada<sup>®</sup>, everything is an object and every object contains a pointer to its class, an object that describes its format and

	<b>class-based systems</b>	<b>Self: no classes</b>
inheritance relationships	instance of subclass of	inherits from
creation metaphor	build according to plan	clone an object
initialization	executing a "plan"	cloning an example
one-of-a-kind	need extra object for class	no extra object needed
infinite regress	class of class of class of ...	none required

holds its behavior (see Figure 1). In Self, too, everything is an object. But, instead of a class pointer, a Self object contains named slots which may store either state or behavior. If an object receives a message and it has no matching slot, the search continues via a parent pointer. This is how Self implements inheritance. Inheritance in Self allows objects to share behavior, which in turn allows the programmer to alter the behavior of many objects with a single change. For instance, a point object<sup>2</sup> would have slots for its non-shared characteristics:  $x$  and  $y$  (see Figure 1). Its parent would be an object that held the behavior shared among all points:  $+$ ,  $-$ , etc.

## 2.1 Comparing Prototypes and Classes

One of Self's most interesting aspects is the way it combines inheritance, prototypes, and object creation, eliminating the need for classes (see above table).

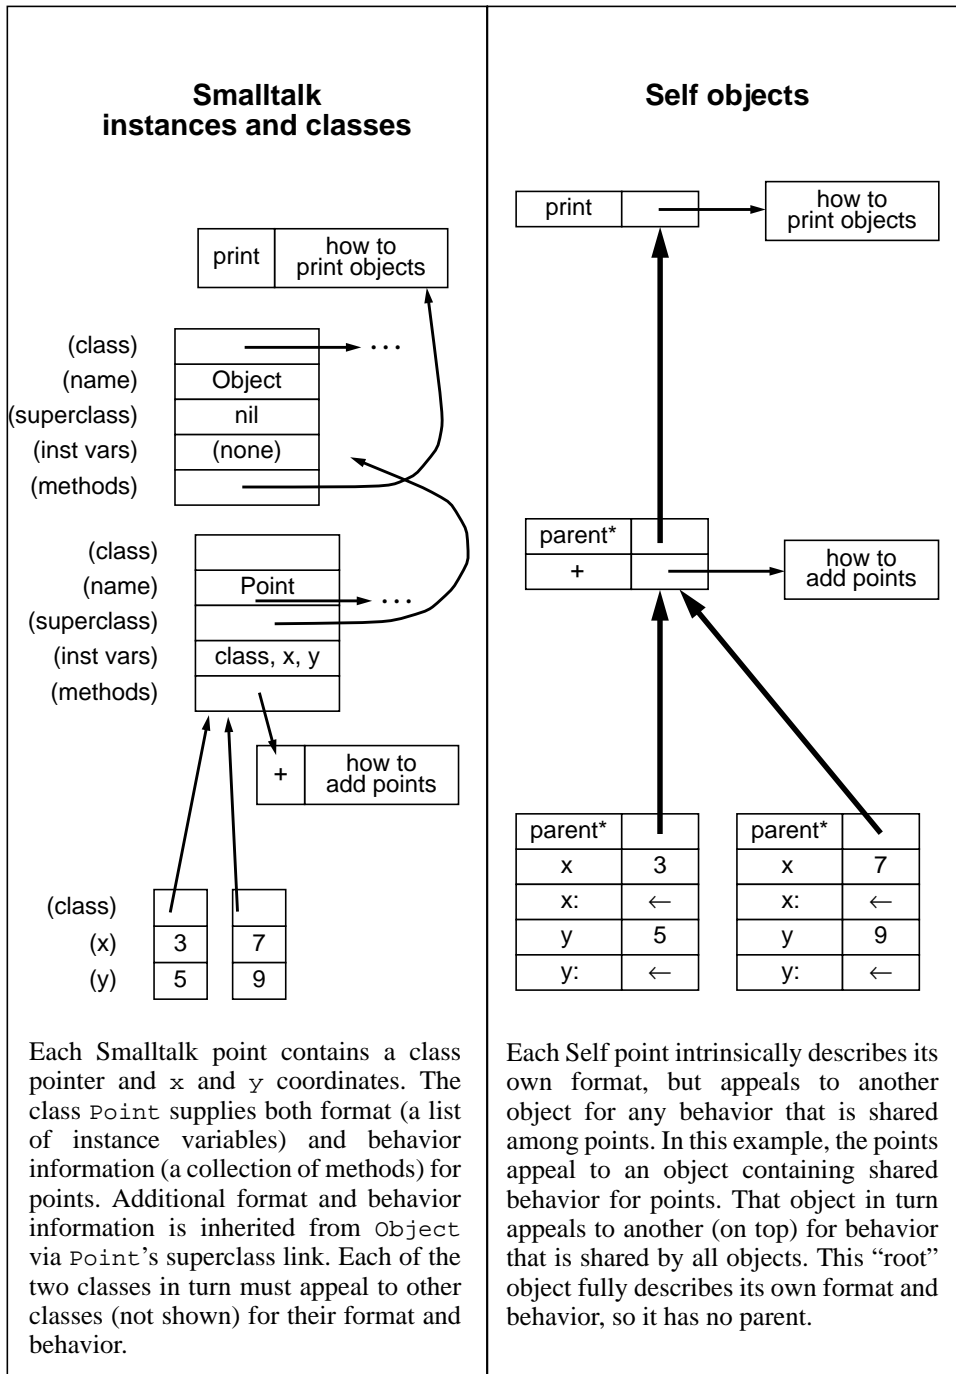
### Simpler relationships.

Prototypes can simplify the relationships between objects. To visualize the way objects behave in a class-based language, one must grasp two relationships: the "is a" relationship, that indicates that an object is an instance of some class, and the "kind of" relationship, that indicates that an object's class is a subclass of some other object's class. In a system such as Self with prototypes instead of classes, there is only one relationship, "inherits from," that describes how objects share behavior and state. This structural simplification makes it easier to understand the language and easier to formulate an inheritance hierarchy.

A working system would provide the chance to discover whether class-like objects would be so useful that programmers would create them without encouragement from the language. The absence of the class-instance distinction may make it too hard to understand which objects exist solely to provide shared information for other objects. Perhaps Self programmers will create entirely new organizational structures. In any case, Self's flexibility poses a challenge to the programming environment; it will have to include navigational and descriptive aids.<sup>3</sup>

<sup>2</sup>Throughout this paper, we appeal to point objects in examples. A Smalltalk point represents a point in two-dimensional Cartesian coordinates. It contains two instance variables: the  $x$  and  $y$  coordinates.

<sup>3</sup>See [19] for a description of the structures that have evolved since this was first written.



**Figure 1.** A comparison of Smalltalk instances and classes with Self objects. At the bottom of each figure are two point objects that have been created by a user program.

**Creation by copying.**

Creating new objects from prototypes is accomplished by a simple operation, copying, with a simple biological metaphor, cloning. Creating new objects from classes is accomplished by instantiation, which includes the interpretation of format information in a class (see Figure 2). Instantiation is similar to building a house from a plan. Copying appeals to us as a simpler metaphor than instantiation.

**Examples of pre-existing modules.**

Prototypes are more concrete than classes because they are examples of objects rather than descriptions of format and initialization. These examples may help users to reuse modules by making them easier to understand. A prototype-based system allows users to examine a typical representative rather than requiring them to make sense out of its description.

**Support for one-of-a-kind objects.**

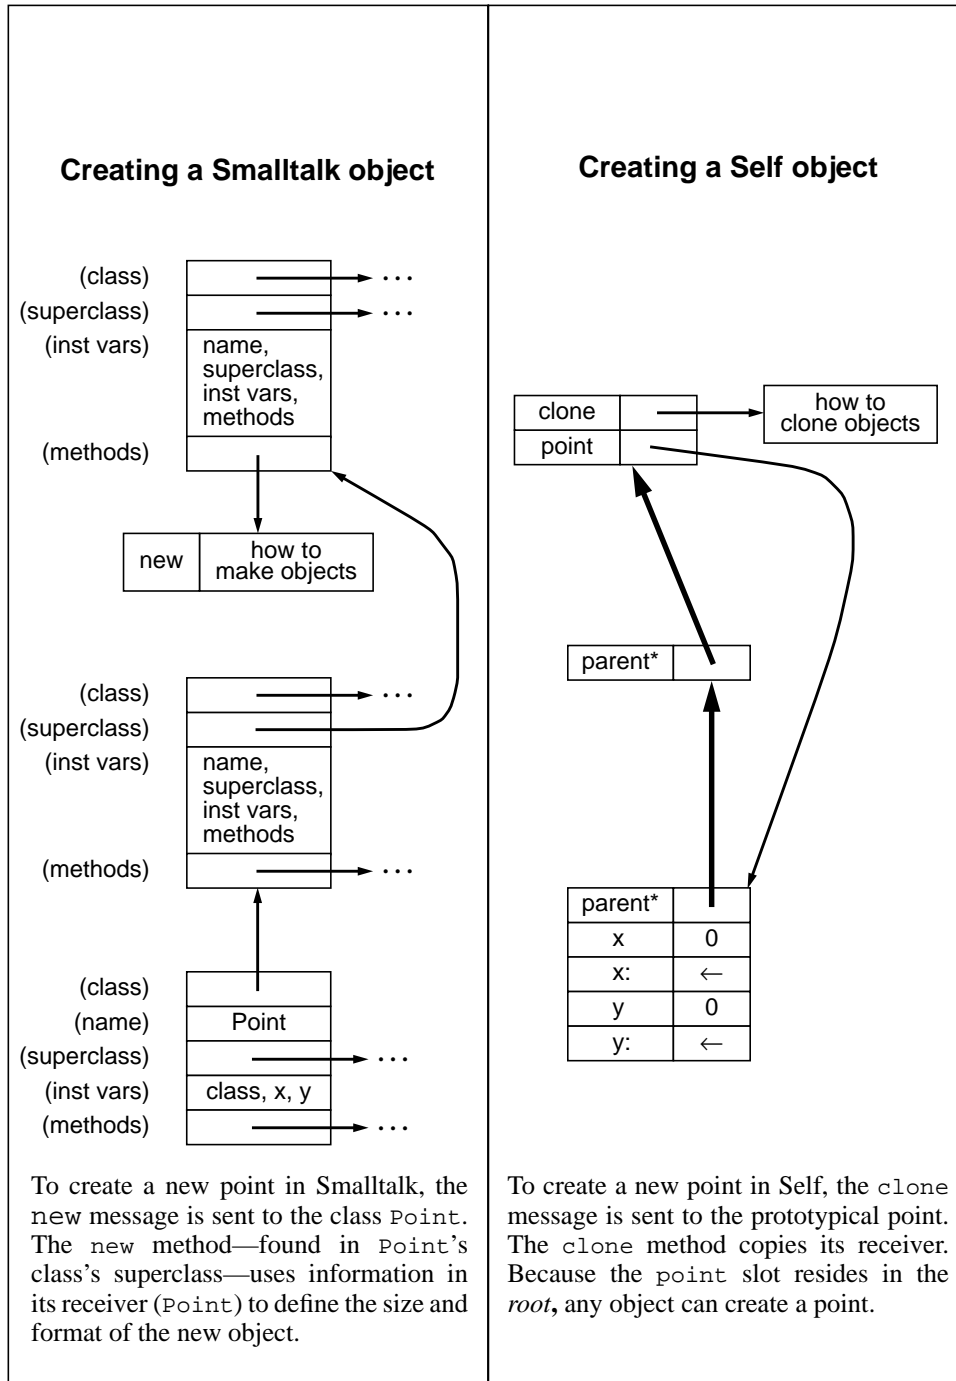
Self provides a framework that can easily include one-of-a-kind objects with their own behavior. Since each object has named slots, and slots can hold state or behavior, any object can have unique slots or behavior (see Figure 3). Class-based systems are designed for situations where there are many objects with the same behavior. There is no linguistic support for an object to possess its own unique behavior, and it is awkward to create a class that is guaranteed to have only one instance. Self suffers from neither of these disadvantages. Any object can be customized with its own behavior. A unique object can hold the unique behavior, and a separate “instance” is not needed.

**Elimination of meta-regress.**

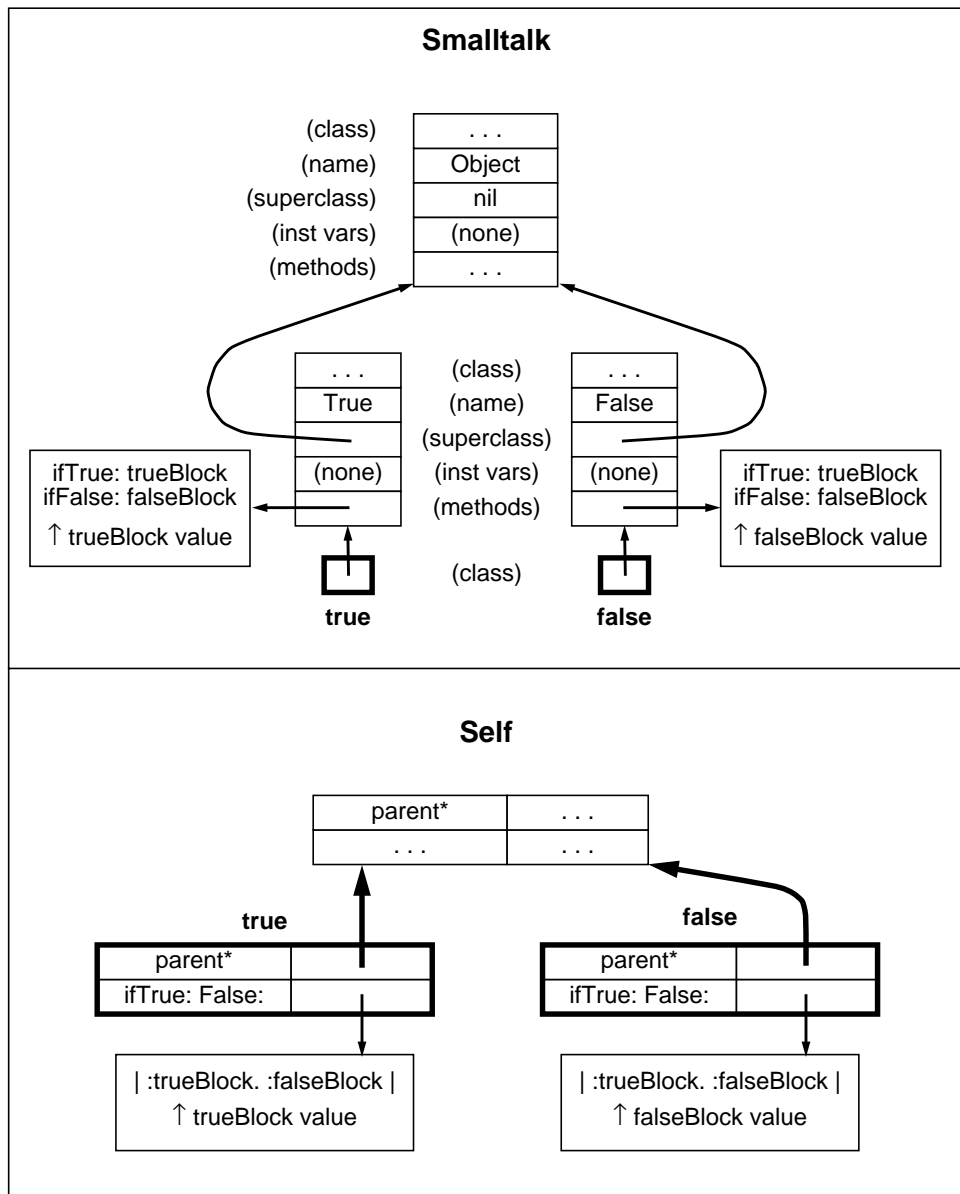
No object in a class-based system can be self-sufficient; another object (its class) is needed to express its structure and behavior. This leads to a conceptually infinite meta-regress: a point is an instance of class `Point`, which is an instance of meta-class `Point`, which is an instance of metaclass `Point`, *ad infinitum*. On the other hand, in prototype-based systems, an object can include its own behavior; no other object is needed to breathe life into it. Prototypes eliminate meta-regress.

The discussion of prototypes in this paper naturally applies to them as realized in Self. Prototype-based systems without inheritance would have a problem: each object would include all of its own behavior—just like the real world—and these systems would surrender one of the most pleasant differences between computers and the real world, the ability to make sweeping changes by changing shared behavior.

Once inheritance is introduced into the language, the natural tendency is to make the prototype the same object that contains the behavior for that kind of object. For instance, the behavior of all points could be changed by changing the behavior of the prototypical point. Unfortunately, such a system must supply two ways to create objects: one to make an object that is the offspring of a prototype, and another to copy an object that is not a prototype. The ultimate result is that prototypes would become special and not prototypical at all.



**Figure 2.** Object creation in Smalltalk and in Self.



**Figure 3.** It is easier to define unique objects in Self than in a class-based system like Smalltalk. Consider the objects that represent the true and false boolean values. A system needs only one instance of each object, but in Smalltalk, there must be a class for each. In Self, since any object can contain behavior, it is straightforward to create specialized objects for true and false.



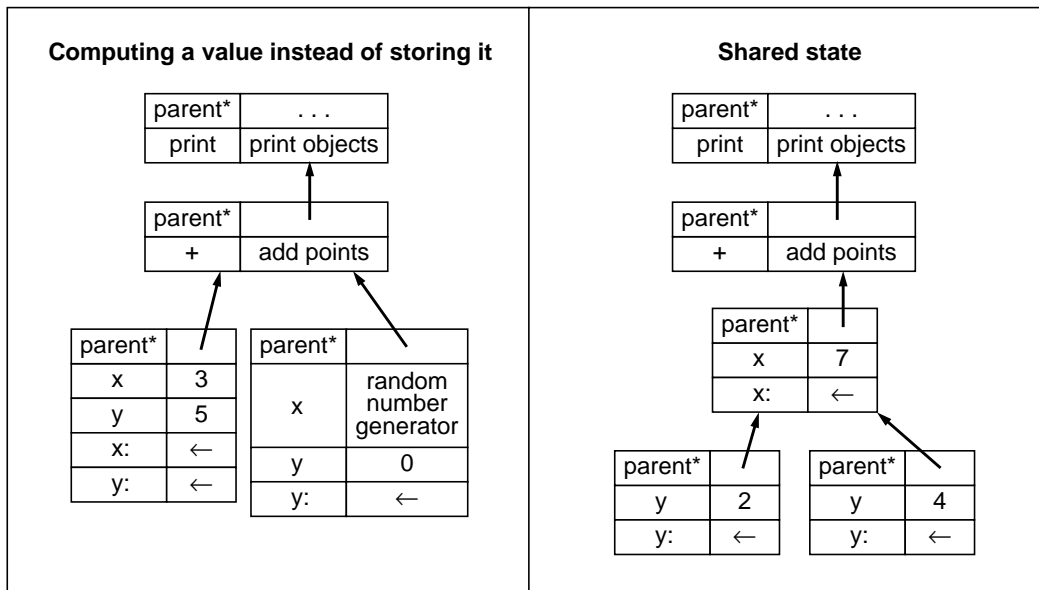
Our solution is to put the shared behavior for a family of objects in a separate object that is the parent of all of them, even the prototype. That way the prototype is absolutely identical to every other member of the family. The object containing the shared behavior plays a role akin to a class, except that it contains no information about representation; it merely holds some shared behavior. So to add some behavior to all points in Self, one would add that behavior to the parent of the points.

### 3 Blending State and Behavior

In Self, there is no direct way to access a variable; instead, objects send messages to access data residing in named slots. So, to access its “x” value, a point sends itself the “x” message. The message finds the “x” slot, and evaluates the object found therein. Since the slot contains a number, the result of the evaluation is just the number itself. In order to change contents of the “x” slot to, say, 17, instead of performing an assignment like “x←17,” the point must send itself the “x:” message with 17 as the argument. The point object must contain a slot named “x:” containing the assignment primitive. Of course, all these messages sent to “self” would make for verbose programs, so our syntax allows the “self” to be elided. The result is that accessing state via messages in Self becomes as easy to write as accessing variables directly in Smalltalk; “x” accesses the slot by the same name, and “x: 17” stores seventeen in the slot.

Accessing state via messages makes inheritance more powerful. Suppose we wish to create a new kind of point, whose “x” coordinate is a random number instead of a stored value. We copy the standard point, remove the “x:” slot (so that “x” cannot be changed) and replace the contents of the “x” slot with the code to generate a random number (see Figure 4). If, instead of modifying the “x” slot, we had replaced the “x:” slot with a “halt” method, we would obtain a breakpoint on write. Thus, Self can express the idioms associated with *active variables* and *dæmons*. Accessing state via messages also makes it easier to share state. To create two points that share the same “x” coordinate, the “x” and “x:” slots can be put in a separate object that is a parent of each of the two points (see Figure 4).

In most object-oriented languages, accessing a variable is a different operation than sending a message. This dilutes the message passing model of computation with assignment and access. As a result, message passing becomes less powerful. For instance, the inclusion of variables makes it harder for a specialization (subclass) to replace a variable with a computed result, because there may be code in a superclass that directly accesses the variable. Also, class-based languages typically store the names and order of instance variables in an object’s class (as shown in Figure 1). This further limits the power of inheritance; the specification within a class unnecessarily restricts an instance’s format. Finally, variable access requires scoping rules, yet a further complication. For instance, Smalltalk has five kinds of variables: local variables (temporaries), instance variables, class variables, pool variables, and global variables, whose scopes roughly correspond to rungs on a ladder of instantiation.



**Figure 4.** Two examples of flexibility in Self. On the left is a point whose  $x$  coordinate is computed by a random number generator. Since all the code for point sends messages to get the  $x$  value, the random  $x$  point can reuse all existing point code. On the right are two points that share the same  $x$  variable.

## 4 Closures and Methods

The Lisp community has obtained excellent results with closures (or lambda-expressions) as a basis for control structures [1, 17]. Experience with Smalltalk blocks supports this; closures provide a powerful, yet easy-to-use metaphor for users to exploit and define their own control structures. Furthermore, this ability is crucial to any language that supports user-defined abstract data types. In Self, a closure is represented by an object containing an environment link and a method named “value,” “value:,” “value:With:,” and so forth, depending on the number of arguments.

### Local variables.

Methods require storage for local variables, and in Self, their slots fulfill this function. In Smalltalk, a method creates an activation record whose initial contents are *described* by the method. For example, the number of temporary variables listed in the method describes the number of fields set aside in the activation record to hold variables. This is similar to the way a class contains a structural description used to instantiate its instances. But in Self, objects that play the role of methods are *prototypes* of activation records; they are copied and invoked to run the subroutine. So, local variables are allocated by reserving slots for them in the prototype activation

record. One advantage is that the prototype's slots may be initialized to any value; they may even contain private methods and closures (blocks).

### **Environment link.**

In general, a method must contain a link to its enclosing closure or scope. This link is used to resolve references to variables not found in the method itself. In Self, instead of having separate scope information, a method's *parent* link performs this function. If a slot is not found in the current scope, `lookup` proceeds to the next outer scope by following the parent link.

Some interesting mechanisms are needed to make the parent links handle lexical scoping. First, the parent link must be set to the appropriate object. This is simple for an ordinary object; the parent link is automatically set to its prototype's parent as a result of cloning. For methods, the object created by the compiler serves as a prototype activation, and when invoked, is cloned. The clone's parent is then set to the message's receiver. In this fashion, the method's scope is embedded in the receiver's. For Self blocks, an environment link is set to the block's enclosing activation record when the block is created. Later, when the block is activated, its method's parent link is set from the block's environment link.

Second, in order to allow the slots containing local variables to be accessed in the same way as everything else, the implicit “`self`” operand must take on an unusual meaning: start the *message lookup* with the current activation record, but set the *receiver* of the message to be the same as the current receiver. In a way, this is the opposite of the “*super*” construct in Smalltalk, which starts the lookup with the receiver's superclass (see Figure 5).

## **5 Speculation: Where is Self Headed?**

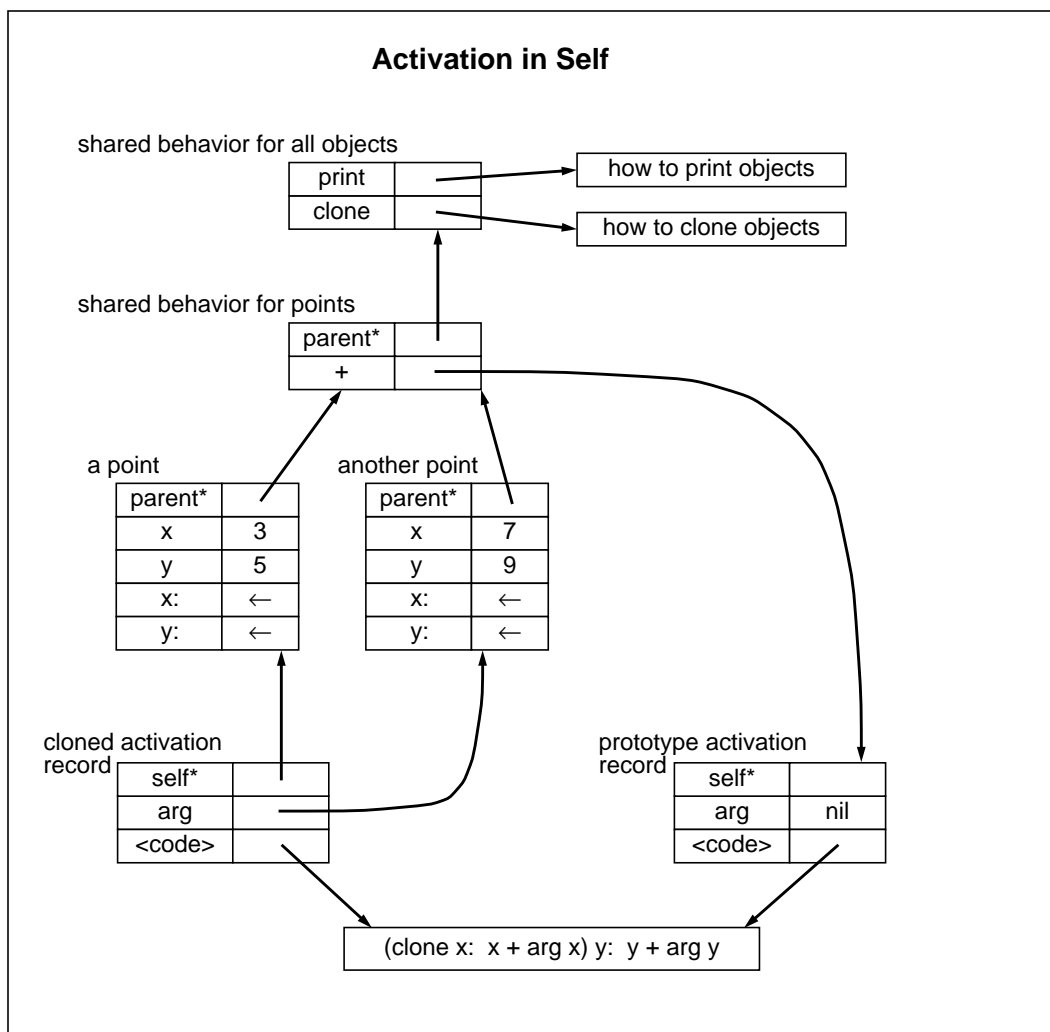
In designing Self, we have been led to some rather strange recurring themes. We present them here for the reader to ponder.

### **Behaviorism.**

In many object languages, objects are passive; an object is what it is. In Self, an object is what it does. Since variable access is the same as message passing, ordinary passive objects can be regarded merely as methods that always return themselves. For example, consider the number 17. In Smalltalk, the number 17 represents a particular (immutable) state. In Self, the number 17 is just an object that returns itself and behaves a certain way with respect to arithmetic. The only way to know an object is by its actions.

### **Computation viewed as refinement.**

In Smalltalk, the number 17 is a number with some particular state, and the state information is used by the arithmetic primitives—addition, for example. In Self, 17 can be viewed as a *refinement of shared behavior* for numbers that responds to addition by returning 17 more than its argument. Since an activation record's parent is



**Figure 5.** This figure shows what happens when the point (3, 5) is sent the “+” message with argument (7, 9). Lookup for plus starts at (3, 5) and finds a matching slot in the object holding shared behavior for points. Since the contents of the slot is a method object, it is cloned, the clone’s argument slot is set to the argument of the message, and its parent is set to the receiver. When the code for “+” executes, the lookup for `x` will find the receiver’s `x` slot by following the inheritance chain from the current activation record. It will also find the contents of the “arg” slot in the same way. It is this technique of having the lookup for the implicit “self” receiver start at the current activation that allows local variables, instance variables, and method lookup to be unified in Self.

set to the receiver of the message in Self, method activation can be viewed as the creation of a short-lived *refinement* of the receiver. Likewise, block (closure) activation can be viewed as the creation of a *refinement* of the activation record for the enclosing context scope.

In our examples, we render the shared behavior object for points as an ordinary *passive* object. Another twist would be to build class-like objects out of *methods*. In Self, the shared behavior object for points could be a method with code that simply returned a clone of the prototypical point. This method could then be installed in the “point” slot of the root object. One object would then be serving two roles: its code would create new points, and its slots (locals) would hold the shared behavior for points. At this writing, we do not believe that this is the best way to construct a system, but the use of methods to hold shared behavior for a group of objects is an example of the flexibility afforded by Self.

### Parents viewed as shared parts.

Finally, one can view the parents of an object as shared parts of the object. From this perspective, a Self point contains a private part with *x* and *y* slots, a part shared with other points containing *+*, *-*, etc. slots, and a part shared with all other objects containing behavior common to all objects. Viewing parents as shared parts broadens the applicability of inheritance.

## 6 Syntax

In this section, we outline the syntax for a textual representation of Self objects. Where possible, we have followed Smalltalk syntax to avoid confusion. We have added slot list syntax for creating objects inline. In general, Self objects are written enclosed in parentheses, and include a list of slots and, in the case of methods, also include code. Blocks are written as methods enclosed in square brackets instead of parentheses.

The code follows Smalltalk syntax, with a few notable exceptions: The receiver is omitted when it is “self.” The return value of a method is always the result of the last expression. Keyword messages associate from right to left. Case is used to make keyword message-sends easier to read: the first keyword must be lowercase, and subsequent keywords in the same selector must be uppercase. These changes reduce the number of parentheses and other lexical noise in Self code.

The slot list syntax has little precedent in Smalltalk. The slot list, if present, must be nestled in a pair of vertical bars. Each item in the slot list must be separated from the next by a period. (A trailing period is optional.) Finally, there are several forms for slots:

- A selector by itself denotes *two* slots: a slot initialized to *nil*, and a slot named with a trailing colon initialized to the assignment primitive (denoted by  $\leftarrow$ ). For example, the object

```
(|x.|)
```

contains two slots: one called `x` containing *nil*, and another one called `x:` containing `←`. This has the same effect as declaring a Smalltalk variable.

- A selector followed by a left arrow (“`<-`”) and an expression also denotes two slots: a slot initialized to the value of the expression, and a corresponding assignment slot. The effect is similar to an initialized variable. For example, the method object

```
(
  | tally <- 0 |
  10 do: [tally: tally + random].
  tally
)
```

(where “`random`” is a slot in the root that returns random numbers) returns the sum of 10 random numbers. It contains a slot named “`tally`” initialized to zero, and a slot named “`tally:`” containing the assignment primitive.

- A selector followed by an equals sign (“`=`”) and an expression denotes only one slot, initialized to the value of the expression. The effect is identical to that of the left-arrow form, except that the variable is read-only (no assignment slot is created).
- Finally, a unary selector (i.e., identifier) preceded by a colon defines a slot bound to the corresponding argument of the message. For example: “`[|:a.:b|a<b]`” defines a block with two arguments, “`a`” and “`b`.”

The arguments for a method may also be moved into the selector as in Smalltalk:

```
display:At: = (
  | aForm. aPoint |
  (bitblt copy
   destination: self At: aPoint Source: aForm)
  copybits)
```

and

```
display: aForm At: aPoint = (
  (bitblt copy
   destination: self At: aPoint Source: aForm)
  copybits)
```

are equivalent.

## 7 Examples

The first example shows a simple point. It assumes that the system is organized into two kinds of objects: objects that hold shared traits, and prototypes. (In the examples, the primitive `_AddSlotsIfAbsent:` adds slots to its receiver if the slots do not already exist; the primitive `_Define:` redefines its receiver in place.)

```

_AddSlotsIfAbsent: ( |
  traits = ().
  prototypes* = ().
| )

```

*Create an object for holding traits objects.  
Create an object for holding prototypical objects. It is a parent to make it easier to refer to its contents.*

```

traits _AddSlotsIfAbsent:(| clonable=()|)
traits clonable _Define:(|
  copy=(_Clone).
| )

```

*Define object holding traits of clonable objects.  
Define a copy method that invokes the “clone” primitive.*

```

traits _AddSlotsIfAbsent: ( | point = () | )
traits point _Define: ( |
  parent* = traits clonable.
  printString = (
    x printString, '@',
    y printString ).
+ aPoint = (
  | newPoint |
  newPoint: copy.
  newPoint x: x + aPoint x.
  newPoint y: y + aPoint y.
  newPoint ).
- aPoint = (
  (copy x: x - aPoint x)
  y: y - aPoint y ).
| )

```

*Define the object holding traits of points.  
Inherit copying methods.  
Define a method to construct a printable string.  
Concatenate the strings for x and y, separated by an '@'.  
Define method for addition.  
Uses a local slot (newPoint).  
Copy the receiver to serve as the result.  
Set its x coordinate.  
Set its y coordinate.  
Return the new point.  
Define subtraction, exploiting the fact that assignment (e.g., x: y:) returns self.  
(By convention, methods return self if at all possible.)*

```

prototypes _AddSlotsIfAbsent: ( | point = () | )
point _Define: ( |
  parent* = traits point.
  x <- 0.
  y <- 0.
| )

```

*Define the prototypical point.  
Inherit shared traits via a read-only parent slot.  
Define read/write slots x and y initialized to 0. (i.e., also define x: and y: with the assignment primitive.)*

```

traits integer _AddSlots: ( |
  @ y = ( (point copy x: self) y: y ).
| )

```

*Add behavior for creating points to existing integer traits object.*

The following example is more interesting; it shows how to exploit Self’s unique features to build a data structure that holds a binary tree of objects.

```

traits _AddSlotsIfAbsent: ( | emptyTree = () | )
traits emptyTree _Define: ( |
  parent* = traits clonable.
  includes: x = ( false ).
  insert: x = (
    parent: treeNode copy
    contents: x ).
  size = 0.
  do: aBlock = ( self ).
| )

```

*Define an object holding emptyTree traits.  
Inherit copying methods.  
Empty trees never include anything.  
Create a new treeNode, set its contents, and switch my parents. Uses dynamic inheritance.*

```

prototypes _AddSlotsIfAbsent: ( | tree = () | )
tree _Define: ( |
  parent* <- traits emptyTree.
| )
Define the prototypical tree as empty. It has
an assignable parent slot set to
traits emptyTree.

traits _AddSlotsIfAbsent: ( | treeNode = () | )
traits treeNode _Define: ( |
  parent* = traits clonable.
includes: x = ( | subT |
  x = contents ifTrue: [^true].
  subT: x < contents
    ifTrue: [left] False: [right].
  subT includes: x ).
insert: x = ( | subT |
  x = contents ifTrue: [^self].
  subT: x < contents
    ifTrue: [left] False: [right].
  subT insert: x.
  self ).
size = ( left size + 1 + right size ).
do: aBlock = (
  left do: aBlock.
  aBlock value: contents.
  right do: aBlock.
  self ).
copy = (
  (resend.copy left: left copy)
  right: right copy ).
| )
Define the object holding tree node traits.
Inherit copying methods.
This method uses a local variable, subT.
Send contents to self and compare. Uses Self's
ability to inherit state (the contents slot).
Return true if found here.
Select subtree.
Recurse and return result of recursive call.
Insert x into the receiver; same recursion as
includes:.
Copy subtrees when copying a tree node. This
provides new empty trees when a tree node
is copied for insertion into a tree.

prototypes _AddSlotsIfAbsent: ( | treeNode = () | )
treeNode _Define: ( |
  parent* = traits treeNode.
  left <- tree.
  right <- tree.
  contents.
| )
Define the prototypical tree node.
A constant parent slot.
Assignable subtree slots, initialized to the
prototypical (empty) tree.
Assignable slot for contents.

```

## 8 Related Work

We would like to express our deep appreciation to the past and present members of the System Concepts Laboratory at Xerox PARC for blazing the trail with Smalltalk [6]. The way Self accesses state via message passing owes much to conversations with Peter Deutsch, and is reminiscent of an earlier unpublished language of his called “O.” Some Smalltalk programmers have already adopted this style of variable accessing [11]. Trellis/Owl is an independently designed object-oriented



language that includes syntactic sugar to make message invocation look like element access and assignment [12]. This is the reverse of our approach. We stuck with message-passing syntax in Self to emphasize behavioral connotations. Strobe is a frame-based language for AI that also mixes data and behavior in slots [16]. Loops, an extension of InterLisp with objects, also includes active variables [18].

We would like to thank Henry Lieberman for calling our attention to prototypes [10]. Unlike Self, Lieberman's prototypes include shared information as well. His clones inherit from their prototype, adding private slots on demand. Although this approach obviates the need for traits objects, its prototypes are heavier-weight objects than Self's. Exemplars is the name given to prototypes in a project that added a prototype-based object hierarchy to Smalltalk [8]. Like our design for Self, objects are created by *cloning* exemplars, and multiple representations of the same data type are permitted. Unlike Self, this system also includes classes as an abstract type hierarchy, and two forms of multiple inheritance. One interesting contribution is the exemplar system's support for or-inheritance. Self seems to be more unorthodox than exemplars in two respects: it eliminates variable accessing from the language, and it unifies objects, classes, and methods.

The Alternate Reality Kit [14, 15] is a direct manipulation simulation environment based on prototypes and active objects, and it has given us much insight into the world of prototypes. Alan Borning's experience with prototype-based environments, especially ThingLab [2, 3, 4] made him a wonderful sounding board when we were struggling to grasp the implications of prototypes.

The DeltaTalk proposal [5] included several ideas for merging Smalltalk methods and blocks, which helped us to understand the problems in this area. The Actors [7] system has active objects, but these are processes, unlike Self's procedural model. Actors also rejects classes, replacing inheritance with delegation.

Oaklisp [9] is a version of Scheme with message passing at the bottom. However, Oaklisp is class-based and maintains the inheritance hierarchy separately from the lexical nesting; it does not seem to integrate lambdas and objects.

## 9 Conclusions

Self offers a new paradigm for object-oriented languages that combines both simplicity and expressiveness. Its simplicity arises from realizing that classes and variables are not needed. Their elimination banishes the metaclass regress, dispels the illusory distinction between instantiation and subclassing, and allows for the blurring of the differences between objects, procedures, and closures. Reducing the number of basic concepts in a language can make the language easier to explain, understand, and use. However, there is a tension between making the language simpler and making the organization of a system manifest. As the variety of constructs decreases, so does the variety of linguistic clues to a system's structure.

Making Self simpler made it powerful. Self can express idioms from traditional object-oriented languages such as classes and instances, but can go beyond them to express one-of-a-kind objects, active values, inline objects and classes, and over-riding instance variables. We believe that contemplation of Self provides insights into the nature of object-oriented computation.

## 10 Acknowledgments

We would like to thank Daniel Weise and Mark Miller for listening patiently and tutoring us on Scheme. Craig Chambers, Martin Rinard, and Elgin Lee have helped distill and refine the language. Finally, we would like to thank all the readers and reviewers for many helpful comments and criticisms, especially Dave Robson, who helped separate the wheat from the chaff.

## 11 References

- [1] Abelson, H., G. J. Sussman, and J. Sussman. *Structure and Interpretation of Computer Programs*. Massachusetts: MIT Press, 1984.
- [2] Borning, A. H. “ThingLab—A Constraint-Oriented Simulation Laboratory.” Ph.D. thesis, Stanford University, Palo Alto, California, 1979.
- [3] Borning, A. H. “The Programming Language Aspects of ThingLab, A Constraint-Oriented Simulation Laboratory.” *ACM Transactions on Programming Languages and Systems* 3, no. 4 (1981): 353–387.
- [4] Borning, A. H. “Classes Versus Prototypes in Object-Oriented Languages.” *Proceedings of the ACM/IEEE Fall Joint Computer Conference* (1986): 36–40.
- [5] Borning, A. H. and T. O’Shea. “DeltaTalk: An Empirically and Aesthetically Motivated Simplification of the Smalltalk-80<sup>TM</sup> Language.” 1986.
- [6] Goldberg, A. and D. Robson. *SmallTalk-80: The Language and Its Implementation*. Massachusetts: Addison-Wesley, 1983.
- [7] Hewitt, C. and G. Agha. MIT Artificial Intelligence Laboratory. “ACTORS: A Conceptual Foundation For Concurrent Object-Oriented Programming.” 1987.
- [8] LaLonde, W. R., D. A. Thomas, and J. R. Pugh. “An Exemplar Based Smalltalk.” *OOPSLA '86 Conference Proceedings* (1986). Also *SIGPLAN Notices* 21, no. 11 (1986): 322–330.
- [9] Lang, K. J. and B. A. Pearlmutter. “Oaklisp: An Object-Oriented Scheme with First Class Types.” *OOPSLA '86 Conference Proceedings* (1986). Also *SIGPLAN Notices* 21, no. 11 (1986): 30–37.
- [10] Lieberman, H. “Using Prototypical Objects to Implement Shared Behavior in Object-Oriented Systems.” *OOPSLA '86 Conference Proceedings* (1986). Also *SIGPLAN Notices* 21, no. 11 (1986): 214–223.
- [11] Rochat, R. “In Search of Good SmallTalk Programming Style.” *Tektronix Laboratories Computer Research Laboratory Technical Report CR-86-19* (1986).

- [12] Schaffert, C., T. Cooper, B. Bullis, M. Kilian, and C. Wilpolt. “An Introduction to Trellis/Owl.” *OOPSLA '86 Conference Proceedings* (1986). Also *SIGPLAN Notices* 21, no. 11 (1986): 9–16.
- [13] Sheil, B. “Power Tools for Programmers.” *Datamation* 29, no. 2 (1983): 131–144.
- [14] Smith, R. B. “The Alternate Reality Kit: An Animated Environment for Creating Interactive Simulations.” *Proceedings of 1986 IEEE Computer Society Workshop on Visual Languages* (1986): 99–106.
- [15] Smith, R. B. “Experiences with the Alternate Reality Kit: An Example of the Tension Between Literalism and Magic.” *Proceedings of the CHI+GI '87 Conference* (1987): 61–67.
- [16] Smith, R. G. “Strobe: Support for Structured Object Knowledge Representation.” *Proceedings of the 1983 International Joint Conference on Artificial Intelligence* (1983): 855–858.
- [17] Steele, G. L., Jr. “Lambda, the Ultimate Imperative.” AI Memo 353. MIT Artificial Intelligence Laboratory (1976).
- [18] Stefik, M., D. Bobrow, and K. Kahn. “Integrating Access-Oriented Programming into a Multiprogramming Environment.” *IEEE Software Magazine* 3, no. 1 (1986): 10–18.
- [19] Ungar, D., C. Chambers, B. Chang, and U. Höelzle. “Organizing Programs Without Classes.” *Lisp and Symbolic Computation* 4, no. 3 (1991).
- [20] Ungar, D. and R. B. Smith. “SELF: The Power of Simplicity.” *OOPSLA '87 Conference Proceedings* (1987). Also *SIGPLAN Notices* 22, no. 11 (1987): 227–241.

© Copyright 1994 Sun Microsystems, Inc. The SMLI Technical Report Series is published by Sun Microsystems Laboratories, a division of Sun Microsystems, Inc. Printed in U.S.A. This paper was originally published in *LISP AND SYMBOLIC COMPUTATION: An International Journal* 4, no. 3. The Netherlands: Kluwer Academic Publishers, 1991. It is a substantial revision of "Self: The Power of Simplicity" in *OOPSLA '87 Conference Proceedings* (1987) and *SIGPLAN Notices* 22, no. 12 (1987): 227–241.

Unlimited copying without fee is permitted provided that the copies are not made nor distributed for direct commercial advantage, and credit to the source is given. Otherwise, no part of this work covered by copyright hereon may be reproduced in any form or by any means graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an information retrieval system, without the prior written permission of the copyright owner.

#### TRADEMARKS

Sun, Sun Microsystems, and the Sun logo are trademarks or registered trademarks of Sun Microsystems, Inc. UNIX and OPEN LOOK are registered trademarks of UNIX System Laboratories, Inc. All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCworks, and SPARCcompiler are licensed exclusively to Sun Microsystems, Inc. Ada is a registered trademark of U. S. Department of Defense A. J. P. O. Smalltalk-80 is a trademark of ParcPlace Systems. All other product names mentioned herein are the trademarks of their respective owners.

For information regarding the SML Technical Report Series, contact Jeanie Treichel, Editor-in-Chief <jeanie.treichel@eng.sun.com>.

For distribution issues, contact Amy Tashbook, Assistant Editor <amy.tashbook@eng.sun.com>.