Recursive object types in a logic of object-oriented programs

K. Rustan M. Leino

Digital Equipment Corporation Systems Research Center
130 Lytton Ave., Palo Alto, CA 94301, U.S.A.
http://www.research.digital.com/SRC/people/Rustan_Leino

Abstract. This paper formalizes a small object-oriented programming notation. The notation features imperative commands where objects can be shared (aliased), and is rich enough to allow subtypes and recursive object types. The syntax, type checking rules, axiomatic semantics, and operational semantics of the notation are given. A soundness theorem showing the consistency between the axiomatic and operational semantics is also given. A simple corollary of the soundness theorem demonstrates the soundness of the type system. Because of the way types, fields, and methods are declared, no extra effort is required to handle recursive object types.

0 Introduction

It is well known that C.A.R. Hoare's logic of the basic commands of imperative, procedural languages [9] has been useful in understanding imperative languages. Object-oriented programming languages being all the rage, one is surprised that the literature has not produced a corresponding logic for modern object-oriented programs. The control structures of object-oriented programs are similar to those treated by Hoare, but the data structures of object-oriented programs are more complicated, mainly because objects are (possibly shared) references to data fields.

This paper presents a logic for an object-oriented programming notation. In an early attempt at such a logic, Leavens gave an axiomatic semantics for an object-oriented language [11]. However, the language he used differs from popular object-oriented languages in that it is functional rather than imperative, so the values of the fields of objects cannot be changed. America and de Boer have given a logic for the parallel language POOL [4]. This logic applies to imperative programs with object sharing (sometimes called aliasing), but without subtyping and method overriding. In a logic that I will refer to as logic AL, Abadi and I defined an axiomatic semantics for an imperative, object-oriented language with object sharing [2], but it does not permit recursive object types. Poetzsch-Heffter and Müller have defined (but not proved sound) a Hoare-style logic for object-oriented programs that remove many of the previous limitations [18]. However, instead of following the standard methodological discipline of letting the designer of a method define its specification and then checking that implementations meet the specification, the specification of a method in the Poetzsch-Heffter and Müller logic is derived from the method's known implementations. The present logic deals with imperative features, subtyping, and recursive object types.

The literature has paid much attention to the type systems of object-oriented languages. Such papers tend to define some notion of types, the commands of some language, the type rules and operational semantics for the commands, and a soundness theorem linking the type system with the operational semantics. (Several examples of this are found in Abadi and Cardelli's book on objects [1].) But after all that effort, one still doesn't know how to *reason* about the programs that can be written with the provided commands, since no axiomatic semantics is given. In addition to giving a programming notation and its axiomatic semantics, this paper, like the paper describing logic AL, gives an operational semantics and a soundness theorem that links the operational semantics with the axiomatic semantics. The soundness theorem directly implies the soundness of the type system.

A complication with type systems is that types can be *recursive*, that is, an object type T may contain a field of type T or a method whose return type is T. The literature commonly treats recursive data types by introducing some sort of fix-point operator into the type system, good examples of which are a paper by Amadio and Cardelli on recursive types and subtypes [3] and the book by Abadi and Cardelli. By treating types in a dramatically different way, the present logic supports recursive object types without the need for any special mechanism like fix-points. The inclusion of recursive object types is one main advantage of the present logic over logic AL, which does not allow them. (The other main advantage over logic AL is that the present logic can be used with any first-order theory.) Because the given soundness theorem implies the soundness of the type system, the present work contributes also to the world of type systems.

In difference to the paper by Amadio and Cardelli, which considers unrestricted recursive types, the type system in the present paper uses a restriction along the lines of name matching. In particular, types are simply identifiers, and the subtype relation is simply a given partial order among those identifiers. This is much like the classes in Java [8] or the branded object types in Modula-3 [17]. But in contrast to languages like Java or Modula-3, fields and methods are declared separately from types in the language considered in this paper. (This is also done in Cecil [5] and Ecstatic [13].) Not only does this simplify the treatment without loss of applicability to languages like Java and Modula-3, but it also makes explicit the separation of concerns. For example, as the logic shows, having to know all the fields of a particular object type is necessary only for the allocation of a new object.

Furthermore, when a field or method is declared at some type T, each subtype of T automatically acquires, or *inherits*, that field or method. Consequently, one gets behavioral subtyping for free, something that can also be achieved by the inheritance discipline considered by Dhara and Leavens [6]. In contrast, subtype relations frequently found in the literature (including the subtype relation used in logic AL), involves the fields and methods of types. In such treatments of types, one often encounters words like "co-variant"; there will be no further occurrence of such words in this paper.

The rest of this paper is organized as follows. Section 1 relates the present logic to some work that has influenced it. Section 2 describes the declarations that can be used in program environments, and Section 3 describes the commands: their syntax, axiomatic semantics, and operational semantics. Section 4 discusses an example program. Then,

Section 5 states the soundness theorem. Section 6 discusses some limitations of the logic, and the paper concludes with a brief summary.

1 Sources of Influence

My work with Abadi has inculcated the present logic with its style and machinery. The present logic also draws from other sources with which I am quite familiar: my thesis [12], my work on an object logic with Nelson [15], and the Ecstatic language [13]. This section compares the features of these sources of influence with the features of the present logic.

My thesis includes a translation of common object-oriented language constructs into Dijkstra's *guarded commands*, an imperative language whose well-known axiomatic semantics is given in terms of *weakest preconditions* [7]. My attempt at an object logic with Nelson is also based on guarded commands, and Ecstatic is a richer object-oriented programming language defined directly in terms of weakest preconditions. Types, fields, and methods in these three sources are declared in roughly the same way as in the present logic. While these sources do provide a way to reason about object-oriented programs, they take for granted the existence of an operational semantics that implements the axiomatic semantics. The present paper includes an operational semantics for the given commands, and establishes the correctness of the operational semantics with respect to the axiomatic semantics by proving a soundness theorem.

Like logic AL, the present logic has few and simple commands. Each command in logic AL operates on an object store and produces a value placed in a special register called r. In the present logic, commands are allowed to refer to the initial value of that register, which simplifies many of the rules. (It also makes the commands "cute".) Another difference is that the present logic splits logic AL's *let* command into two commands: sequential composition and binding. The separation works well because the initial value of register r can be used. Perhaps surprisingly, another consequence of using the initial value of r is that the present logic manages fine without Abadi and Cardelli's ς binder that appears in logic AL to bind a method's self parameter.

2 Environments

This section starts defining the logic by describing program environments and the declarations that a program environment can contain.

A *program environment* is a list of declarations. A declaration introduces a type, a field, or a method.

An identifier is said to be declared in an environment if it is introduced by a type, field, or method declaration in the environment, or if it is one of the built-in types. I write $x \notin E$ to denote that identifier x is not declared in environment E.

The judgement $E \vdash \diamond$ says that E is a well-formed environment. The empty list, written $\emptyset$, is a valid environment.

Empty Environment

$$\frac{}{\emptyset \vdash \diamond}$$

The next three subsections describe types, fields, and methods, and give the remaining rules for well-formed environments.

2.0 Types

A *type* is an identifier. There are two built-in types, *Boolean* and *Object*. (Other types, like integers, can easily be added, but I omit them for brevity.) Types other than *Boolean* are called *object types*. A new object type is introduced by a *subtyping pair*, which has the form $T{<:}U$, where T is identifier that names the new type, and U is an object type. Like in Java, *Object* denotes the root of the class hierarchy. The analogue of a subtyping pair $T{<:}U$ in Java is a class T declared as a subclass of a class U: *class T extends U* { ... }.

A type is said to be declared in an environment if it is *Boolean*, *Object*, or if it occurs as the first component of a subtyping pair. To express this formally, the judgement $E \vdash_{type} T$ says that T is a type in environment E, and the judgement $E \vdash_{obj} T$ says that T is an object type in E. The rules for these judgements are as follows. Here and throughout this paper, I use T and U, possibly subscripted, to denote types.

Types in Environments

$$\frac{E \vdash_{obj} U \qquad T \notin E}{(E, T{<:}U) \vdash \diamond}$$

Declared Types ($\vdash_{type}$, $\vdash_{obj}$)

$$\frac{E \vdash \diamond}{E \vdash_{obj} \mathit{Object}} \qquad \frac{(E, T{<:}U, E') \vdash \diamond}{(E, T{<:}U, E') \vdash_{obj} T}$$

$$\frac{E \vdash \diamond}{E \vdash_{type} \mathit{Boolean}} \qquad \frac{E \vdash_{obj} T}{E \vdash_{type} T}$$

The reflexive, transitive closure of the subtyping pairs forms a partial order called the *subtyping order*. The judgement $E \vdash T <: U$ says that T and U are types in E that are ordered by the subtyping order. Type T is then said to be a *subtype* of U. The rules are:

Subtyping Order ($\vdash <:$)

$$\frac{E \vdash_{type} T}{E \vdash T <: T} \qquad \frac{(E, T{<:}U, E') \vdash \diamond}{(E, T{<:}U, E') \vdash T <: U} \qquad \frac{E \vdash T_0 <: T_1 \qquad E \vdash T_1 <: T_2}{E \vdash T_0 <: T_2}$$

2.1 Fields

A field is a map from an object type to another type. A field f is introduced by a *field triple*, written $\mathsf{f}\colon T \to U$, where f is the identifier that names the field, T is an object type called the *index type* of f, and U is a type called the *range type* of f. The analogue of a field triple $\mathsf{f}\colon T \to U$ in Java is an instance variable f of type U declared in a class T: *class T* { ... $U\ \mathsf{f}$; ... }.

An environment can contain field triples. A field f is said to be declared in an environment E if it occurs in some field triple $\mathsf{f}\colon T \to U$ in E. This is expressed by the judgement $E \vdash_{field} \mathsf{f}\colon T \to U$. The rules for these judgements are as follows. Here and throughout, I use f, possibly subscripted, to denote field names.

Fields in Environments

$$\frac{E \vdash_{obj} T \qquad E \vdash_{type} U \qquad \mathsf{f} \notin E}{(E, \mathsf{f}\colon T \to U) \vdash \diamond}$$

Declared Fields ($\vdash_{field}$)

$$\frac{(E, \mathsf{f}\colon T \to U, E') \vdash \diamond}{(E, \mathsf{f}\colon T \to U, E') \vdash_{field} \mathsf{f}\colon T \to U}$$

For a type T_0 declared in an environment E, the *set of fields of* T_0 in E, written $\mathit{Fields}(T_0, E)$, is the set of all field triples $\mathsf{f}\colon T \to U$ such that $E \vdash_{field} \mathsf{f}\colon T \to U$ and $E \vdash T_0 <: T$.

2.2 Methods

A *method quadruple* has the form $\mathsf{m}\colon T \to U : R$, where m is an identifier denoting a *method*, T is an object type, U is a type, and R is a *relation*. The analogue of a method quadruple $\mathsf{m}\colon T \to U : R$ in Java is a method m with return type U declared in a class T and given a specification R: *class* T { ... $U\ \mathsf{m}()$ {...} ... }. Note that the Java language does not have a place to write down the specification of a method. In the present language, the declaration of a method includes a specification, which specifies the effect of the method as a relation on the pre- and post-state of each method invocation. Note also that methods take no parameters (other than the object on which the method is invoked, an object commonly referred to as *self*). This simplifies the logic without losing theoretical expressiveness, since parameters can be passed through fields.

An environment can contain method quadruples. A method m is said to be declared in an environment E if it occurs in some method quadruple $\mathsf{m}\colon T \to U : R$ in E. This is expressed by the judgement $E \vdash_{method} \mathsf{m}\colon T \to U : R$. Formally, the rules are as follows. I use m, possibly subscripted, to denote methods.

Methods in Environments

$$\frac{\begin{array}{ll} E \vdash_{obj} T & E \vdash_{type} U \\ E, \emptyset \vdash_{rel} R & \mathsf{m} \notin E \end{array}}{(E, \mathsf{m}\colon T \to U : R) \vdash \diamond}$$

Declared Methods ($\vdash_{method}$)

$$\frac{(E, \mathsf{m}\colon T \to U : R, E') \vdash \diamond}{(E, \mathsf{m}\colon T \to U : R, E') \vdash_{method} \mathsf{m}\colon T \to U : R}$$

The judgement $E, \emptyset \vdash_{rel} R$, which will be described in more detail in Section 3.1, essentially says that R is a relation that may mention fields declared in E but doesn't mention any local program variables.

For a type T_0 declared in an environment E, the *set of methods of* T_0 in E, written $\mathit{Methods}(T_0, E)$, is the set of all method quadruples $\mathsf{m}\colon T \to U : R$ such that $E \vdash_{method} \mathsf{m}\colon T \to U : R$ and $E \vdash T_0 <: T$.

2.3 Relations

Methods are specified using *relations*. A relation is an untyped first-order predicate on a pre-state and a post-state. In order to make relations expressive, the present logic can be used with an *underlying logic*, which provides a set of function symbols and a set of first-order axioms about those function symbols. The example in Section 4 shows how an underlying logic may be used.

Syntactically, relations are made up only of:

– the constants *false*, *true*, and *nil*;
– constants for field names, and the special field alloc;
– the special variables $\grave{r}$, $\acute{r}$, $\grave{\sigma}$, $\acute{\sigma}$;
– other variables (I will write v to denote a typical variable);
– equality between terms;
– applications of the functions *select* and *store*;
– applications of the functions of the underlying logic;
– the usual logical connectives $\neg$, $\wedge$, and $\forall$.

The grammars for relations (R) and terms (e) are thus:

$$\begin{array}{l} R ::= e_0 = e_1 \mid \neg R \mid R_0 \wedge R_1 \mid \langle\, \forall x :: R \,\rangle \\ e ::= \mathit{false} \mid \mathit{true} \mid \mathit{nil} \mid \mathsf{f} \mid \grave{r} \mid \acute{r} \mid \grave{\sigma} \mid \acute{\sigma} \mid v \\ \quad \mid \mathit{select}(e_0, e_1, e_2) \mid \mathit{store}(e_0, e_1, e_2, e_3) \\ \quad \mid \mathit{pa}(e_0, \ldots, e_{ka-1}) \mid \cdots \mid \mathit{pz}(e_0, \ldots, e_{kz-1}) \end{array} ,$$

where $\mathit{pa}, \ldots, \mathit{pz}$ denote the function symbols of the underlying logic with arities $ka, \ldots, kz$, respectively. It will be convenient to also allow $\neq$, $\vee$, $\Rightarrow$, $\Leftarrow$, $\equiv$, and $\exists$ as the usual abbreviations of the operators above.

The semantics of a command (program statement) is defined in terms of a relation on a *register* and a *(data) store*, together called a *state*. The variables $\grave{r}$ and $\acute{r}$ denote the register in the pre- and post-states of the command, respectively, and $\grave{\sigma}$ and $\acute{\sigma}$ denote the store in those respective states. The value of a field f of an object e in a store σ is denoted $\mathit{select}(\sigma, e, \mathsf{f})$. The expression $\mathit{store}(\sigma, e_0, \mathsf{f}, e_1)$ represents the store that results from setting the f field of object e_0 in store σ to the value e_1. The relationship between *select* and *store* is defined as follows.

$$\begin{array}{l} \langle\, \forall \sigma, e_0, e_1, \mathsf{f}_0, \mathsf{f}_1, e :: \\ \quad \mathit{select}(\mathit{store}(\sigma, e_0, \mathsf{f}_0, e), e_0, \mathsf{f}_0) = e \;\wedge \\ \quad (e_0 \neq e_1 \vee \mathsf{f}_0 \neq \mathsf{f}_1 \;\Rightarrow \\ \qquad \mathit{select}(\mathit{store}(\sigma, e_0, \mathsf{f}_0, e), e_1, \mathsf{f}_1) = \mathit{select}(\sigma, e_1, \mathsf{f}_1)) \,\rangle \end{array} \tag{0}$$

The special field alloc is used to record which objects in the data store have been allocated; the alloc field of an object is *false* until the object is allocated, and is *true* from there on.

As we shall see, the logic allows relations to be rewritten. A rewriting uses the rules of logic and some axioms. In particular, a rewriting may use as axioms the definition of *select* and *store* (0), the distinctness of the boolean values and the distinctness of field name constants:

$$\mathit{false} \neq \mathit{true} \tag{1}$$

$$\text{all field name constants (including alloc) are distinct} \tag{2}$$

and the axioms of the underlying logic.

3 Commands

This section describes commands: their syntax, their axiomatic semantics, and their operational semantics.

3.0 Syntax

A *command* has a form dictated by the following grammar.

$a ::= c$	constant
$\mid v$	local variable
$\mid a_0 \diamond a_1$	conditional
$\mid a_0 \,;\, a_1$	composition
$\mid \mathit{with}\ v\colon T\ \mathit{do}\ a$	binding
$\mid [T\colon \mathsf{f}_i = c_i\ ^{i\in I},\ \mathsf{m}_j = a_j\ ^{j\in J}]$	object construction
$\mid \mathsf{f}$	field selection
$\mid \mathsf{f} := v$	field update
$\mid \mathsf{m}$	method invocation
$c ::= \mathit{false} \mid \mathit{true} \mid \mathit{nil}$	

Informally, the semantics of the language is as follows. (Recall from the previous section that commands operate on a register and a store.)

- The constants *false*, *true*, and *nil* evaluate to themselves. That is, they have the effect of setting the register to themselves.
- A local variable is an identifier introduced via a binding command. Every local variable is immutable: once bound (using *with*, see below), the value of a local variable cannot be changed. A local variable evaluates to its value.
- The conditional command evaluates a_0 if the register is initially *false*, and evaluates a_1 if the register is initially *true*. Note that the guard of the conditional is not shown explicitly in the command; rather, the initial value of the register is used as the guard.
- The sequential composition of a_0 and a_1 first evaluates a_0 and then evaluates a_1. The final values of the register and store in the evaluation of a_0 are used as the initial values of the register and store in the evaluation of a_1. Composition is usually written $a_0 \,;\, a_1$, but to keep the language looking like popular object-oriented languages, I also allow the alternative syntax $a_0 \,.\, a_1$ (see examples below).
- The binding command $\mathit{with}\ v\colon T\ \mathit{do}\ a$ introduces a local variable v for use in a. Its evaluation consists in evaluating a with v bound to the initial value of the register.
- The command $[T\colon \mathsf{f}_i = c_i\ ^{i\in I},\ \mathsf{m}_j = a_j\ ^{j\in J}]$ constructs a new object of type T, and sets the register to (a reference to the fields and methods of) the object. The command must list every field f_i from the set of fields of T. The initial value for field f_i is the given constant c_i. The command must also list every method m_j from the set of methods of T. The implementation of method m_j for the new object is given as the command a_j, which receives self as the initial value of the register and returns the method result value as the final value of the register. The command a_j cannot reference local variables other than those it declares.
- A field can be selected (f) and updated ($\mathsf{f} := v$). Both operate on the object referenced by the initial value of the register. Selection sets the register to the f field of the object. Update sets the f field of the object to the value of v, leaving the register unchanged.

– The method invocation m finds the implementation of method m for the object referenced by the initial value of the register, and then proceeds to evaluate that implementation. The evaluation of the implementation begins with the initial register and store values of the invocation, and the invocation ends with the final register and store values of the evaluation of the implementation. Other than the initial and final register values (which encode self and the result value, respectively), a method does not have explicit parameters; instead, parameters can be passed via the fields of the object.

Here are some examples that compare the present commands with programs written in other languages. The Modula-3 program statement *if b then S else T end* is written as the command $b\ ;\ (T \diamond S)$. The Modula-3 expression $\mathit{new}(T, \mathsf{f} := \mathit{true}).\mathsf{f}$, where T is an object type with one field f and no methods, is written as the command $[T : \mathsf{f} = \mathit{true}]\ ;\ \mathsf{f}$, or with the alternative syntax for composition, the command is written $[T : \mathsf{f} = \mathit{true}].\mathsf{f}$. The Modula-3 program $x.\mathsf{f} := \mathit{true}$ is written $\mathit{true}\ ;\ \mathit{with}\ v\colon \mathit{Boolean}\ \mathit{do}\ x.\mathsf{f} := v$.

As an example of object sharing, the command

$$[T\colon \mathsf{f} = c]\ ;\ \mathit{with}\ v\colon T\ \mathit{do}\ \mathit{with}\ w\colon T\ \mathit{do}\ (v.\mathsf{f} := y\ ;\ w.\mathsf{f})$$

allocates a new T object whose f field is set to c, creates two references to the object (v and w), updates the object's f field via v, and reads f back via w, returning y.

The following example shows the construction of a T object whose method or computes the disjunction of fields x and y:

$$[T\colon \mathsf{x} = \mathit{false},\ \mathsf{y} = \mathit{false},\ \mathsf{or} = \mathit{with}\ \mathit{self}\colon T\ \mathit{do}\ (\mathsf{x}\ ;\ (\mathit{self}.\mathsf{y} \diamond \mathit{true}))] \quad .$$

Note that although primitive, the programming notation is expressive enough to admit common object-oriented languages features like object construction, method invocation, and object sharing. The programming notation is kept minimal in order to simplify the associated rules.

3.1 Axiomatic Semantics

This subsection gives the axiomatic semantics of the commands. The judgement

$$E, V \vdash a : T \to U : R$$

says that command a in command environment (E, V) can be started in a state where the register contents has type T, and terminates in a state where the register contents has type U. The execution of a is such that its pre- and post-states satisfy the relation R. The rules of the axiomatic semantics double as type checking rules, because with a trivial R (such as $\grave{r} = \grave{r}$), the judgement expresses what it means for command a to be well-typed.

Before giving the axiomatic semantics, some other definitions and rules pertaining to constants, local variables, and command environments are in order.

There are three constants: *false*, *true*, and *nil*. The judgement $E \vdash_{\mathit{const}} c\colon T$ expresses that constant c has type T.

Type of Constants ($\vdash_{const}$)

$$\frac{E \vdash \diamond}{E \vdash_{const} false\colon Boolean} \qquad \frac{E \vdash \diamond}{E \vdash_{const} true\colon Boolean} \qquad \frac{E \vdash_{obj} T}{E \vdash_{const} nil\colon T}$$

A *local variable declaration* has the form $v\colon T$, where v is an identifier denoting a *local variable* and T is a type. A *command environment* is a pair (E, V), where E is a program environment and V is a list of local variable declarations. A local variable v is said to be declared in a command environment (E, V) if it occurs in some local variable declaration $v\colon T$ in V. This is expressed by the judgement $E, V \vdash_{var} v\colon T$. Thus, in a command environment (E, V), E contains declarations of types, fields, and methods, whereas V contains declarations of local variables. This separation allows a simple characterization of a command environment without local variable declarations: $(E, \emptyset)$. We saw this in the "Methods in Environments" rule in Section 2.2, and we will see it in the "Object Construction" rule below and in Theorems 0, 1, and 2 in Section 5.

The judgement

$$E, V \vdash_{rel} R$$

says that R is a relation whose free variables are fields or local variables declared in (E, V), or are among the special fields and variables alloc, $\grave{r}$, $\acute{r}$, $\grave{\sigma}$, and $\acute{\sigma}$. The obvious formal rules for this judgement are omitted. Thus, the judgement $E, \emptyset \vdash_{rel} R$ used in the hypothesis of the "Methods in Environments" rule in Section 2.2 implies that R does not mention local variables.

I write $x \notin (E, V)$ to denote that identifier x is not declared in command environment (E, V). The formal rules of the above are then:

Well-formed Command Environment **Declared Local Variables** ($\vdash_{var}$)

$$\frac{E \vdash \diamond}{E, \emptyset \vdash \diamond} \qquad \frac{E, V \vdash \diamond \quad v \notin (E, V) \quad E \vdash_{type} T}{E, (V, v\colon T) \vdash \diamond} \qquad \frac{E, (V, v\colon T, V') \vdash \diamond}{E, (V, v\colon T, V') \vdash_{var} v\colon T}$$

Now for the rules of the axiomatic semantics. There is one rule for each command, and one subsumption rule.

Subsumption

$$\frac{E, V \vdash a : T_1 \to T_2 : R \qquad E \vdash T_0 <: T_1 \qquad E \vdash T_2 <: T_3 \qquad \vdash_{fol} R \Rightarrow R' \qquad E, V \vdash_{rel} R'}{E, V \vdash a : T_0 \to T_3 : R'}$$

The judgement $\vdash_{fol} P$ represents provability in first-order logic, under axioms (0), (1), and (2) from Section 2.3 and the axioms of the underlying logic.

Constant **Local Variable**

$$\frac{E, V \vdash \diamond \qquad E \vdash_{const} c\colon T \qquad E \vdash_{type} U}{E, V \vdash c : U \to T : \acute{r} = c \ \wedge\ \grave{\sigma} = \acute{\sigma}} \qquad \frac{E, V \vdash_{var} v\colon T \qquad E \vdash_{type} U}{E, V \vdash v : U \to T : \acute{r} = v \ \wedge\ \grave{\sigma} = \acute{\sigma}}$$

Conditional

$$\frac{E, V \vdash a_0 : Boolean \to T : R_0 \qquad E, V \vdash a_1 : Boolean \to T : R_1}{E, V \vdash a_0 \diamond a_1 : Boolean \to T : (\grave{r} = false \ \Rightarrow\ R_0) \ \wedge\ (\grave{r} = true \ \Rightarrow\ R_1)}$$

Composition

$$\frac{E, V \vdash a_0 : T_0 \to T_1 : R_0 \qquad E, V \vdash a_1 : T_1 \to T_2 : R_1 \qquad \check{r} \text{ and } \check{\sigma} \text{ do not occur free in } R_0 \text{ or } R_1}{E, V \vdash a_0 \,;\; a_1 : T_0 \to T_2 : \langle\, \exists \check{r}, \check{\sigma} \;::\; R_0[\acute{r}, \acute{\sigma} := \check{r}, \check{\sigma}] \wedge R_1[\grave{r}, \grave{\sigma} := \check{r}, \check{\sigma}] \,\rangle}$$

Binding

$$\frac{E, (V, v\colon T) \vdash a : T \to U : R}{E, V \vdash \mathit{with}\ v\colon T\ \mathit{do}\ a : T \to U : R[v := \grave{r}]}$$

Object Construction

$$\frac{\begin{array}{l} E, V \vdash \diamond \qquad E \vdash_{\mathit{type}} U \qquad E \vdash_{\mathit{obj}} T \\ \mathsf{f}_i\colon T_i \to U_i\ ^{i \in I} \text{ are the elements of } \mathit{Fields}(T, E) \qquad E \vdash_{\mathit{const}} c_i\colon U_i\ ^{i \in I} \\ \mathsf{m}_j\colon T_j \to U_j : R_j\ ^{j \in J} \text{ are the elements of } \mathit{Methods}(T, E) \qquad E, \emptyset \vdash a_j : T \to U_j : R_j\ ^{j \in J} \end{array}}{\begin{array}{l} E, V \vdash [T\colon\ \mathsf{f}_i = c_i\ ^{i \in I},\ \mathsf{m}_j = a_j\ ^{j \in J}] : U \to T : \acute{r} \neq \mathit{nil} \wedge \\ \qquad \mathit{select}(\grave{\sigma}, \acute{r}, \mathsf{alloc}) = \mathit{false} \wedge \acute{\sigma} = \mathit{store}(\cdots(\mathit{store}(\grave{\sigma}, \acute{r}, \mathsf{alloc}, \mathit{true}),\ \ \acute{r}, \mathsf{f}_i, c_i)\ ^{i \in I} \end{array}}$$

Field Selection

$$\frac{E, V \vdash \diamond \qquad E \vdash_{\mathit{field}} \mathsf{f}\colon T \to U}{\begin{array}{l} E, V \vdash \mathsf{f} : T \to U : \\ \grave{r} \neq \mathit{nil} \Rightarrow \acute{r} = \mathit{select}(\grave{\sigma}, \grave{r}, \mathsf{f}) \wedge \grave{\sigma} = \acute{\sigma} \end{array}}$$

Field Update

$$\frac{\begin{array}{ll} E \vdash_{\mathit{field}} \mathsf{f}\colon T_0 \to U_0 & E \vdash T_1 <: T_0 \\ E, V \vdash_{\mathit{var}} v\colon U_1 & E \vdash U_1 <: U_0 \end{array}}{\begin{array}{l} E, V \vdash \mathsf{f} := v : T_1 \to T_1 : \\ \grave{r} \neq \mathit{nil} \Rightarrow \grave{r} = \acute{r} \wedge \acute{\sigma} = \mathit{store}(\grave{\sigma}, \grave{r}, \mathsf{f}, v) \end{array}}$$

Method Invocation

$$\frac{E, V \vdash \diamond \qquad E \vdash_{\mathit{method}} \mathsf{m}\colon T \to U : R}{E, V \vdash \mathsf{m} : T \to U : \grave{r} \neq \mathit{nil} \;\Rightarrow\; R}$$

3.2 Operational Semantics

The operational semantics is defined by the judgement

$$r, \sigma, \mu, S \vdash a \rightsquigarrow r', \sigma', \mu' \qquad .$$

It says that given an initial *operational state* (r, σ, μ) and *stack* S, executing command a terminates in operational state (r', σ', μ'). Operational states are triples whose first two components correspond to the register and data store components of states, as defined above. The third component is a *method store*. Let $\mathcal{H}$ denote a set of given *object names*. A *stack* is a partial function from local variables to $\mathcal{H} \cup \{\mathit{false}, \mathit{true}, \mathit{nil}\}$. A *method store* is a partial function μ from $\mathcal{H}$, such that

- $\mu(h)(\mathsf{type})$ is the allocated type of object h, and
- $\mu(h)(\mathsf{m})$, if defined, is the implementation of method m of object h.

A *store pair* is a pair (σ, μ) where σ is a data store and μ is a method store.

In addition to keeping the method implementations of objects, the method store keeps the allocated type of objects. The operational semantics records this information as it allocates a new object, but doesn't use it subsequently. The information is used only

to state and prove the soundness theorem. By conveniently recording this information in the operational semantics, where it causes no harm, one avoids the use of a *store type* (*cf.* [2]). The result is a simpler statement and proof of soundness.

To save space, I omit the rules for the operational semantics. They can be found in a SRC Technical Note [14].

4 Example

In this section, I show an example of a program that can be proved in the logic.

Let us consider a linked-list type with a method that appends a list to another. Reasoning about a program with such a type requires reasoning about reachability among linked-list nodes. To this end, we assume the underlying logic to contain a function symbol *Reach* (adapted from Greg Nelson's reachability predicate [16]). Informally, $Reach(e_0, e_1, \sigma, \mathsf{f}, e_2)$ is *true* whenever it is possible to reach from object e_0 to object e_1 via applications of f in σ, never going through object e_2.

The example in this section assumes that the underlying logic contains the following two axioms, which relate *Reach* to *select* and *store*, respectively.

$$\begin{array}{l} \langle\, \forall\, e_0, e_1, \sigma, \mathsf{f}, e_2 \;::\; Reach(e_0, e_1, \sigma, \mathsf{f}, e_2) = true \quad \equiv \\ \qquad e_0 = e_1 \vee (e_0 \neq e_2 \wedge Reach(select(\sigma, e_0, \mathsf{f}), e_1, \sigma, \mathsf{f}, e_2) = true) \,\rangle \end{array} \tag{3}$$

$$\begin{array}{l} \langle\, \forall\, e_0, e_1, \sigma, \mathsf{f}_0, e_2, \mathsf{f}_1, e_3, e_4 \;::\; \mathsf{f}_0 \neq \mathsf{f}_1 \;\Rightarrow \\ \qquad Reach(e_0, e_1, \sigma, \mathsf{f}_0, e_2) \;=\; Reach(e_0, e_1, store(\sigma, e_3, \mathsf{f}_1, e_4), \mathsf{f}_0, e_2) \,\rangle \end{array} \tag{4}$$

Axiom (3) resembles Nelson's axiom A1 and says that every object reaches itself, and that e_0 reaches e_1 if e_0 is not e_2 and $e_0.\mathsf{f}$ reaches e_1. Axiom (4) says that whether or not e_0 reaches e_1 via f_0 is independent of the values of another field f_1.

The example uses the following environment, which I shall refer to as E:

$$\begin{array}{cc} Node{<:}Object & \mathsf{next}\colon Node \to Node \\ \mathsf{appendArg}\colon Node \to Node & \mathsf{append}\colon Node \to Node : R \end{array}$$

where R is the relation

$$\begin{array}{rl} \grave{r} \neq nil \;\Rightarrow\; & Reach(\grave{r}, \acute{r}, \grave{\sigma}, \mathsf{next}, nil) \wedge select(\grave{\sigma}, \acute{r}, \mathsf{next}) = nil \wedge \\ & select(\acute{\sigma}, \acute{r}, \mathsf{next}) = select(\grave{\sigma}, \grave{r}, \mathsf{appendArg}) \wedge \\ & \langle\, \forall\, o, \mathsf{f} \;::\; select(\grave{\sigma}, o, \mathsf{f}) = select(\acute{\sigma}, o, \mathsf{f}) \vee \\ & \qquad\qquad (o = \acute{r} \wedge \mathsf{f} = \mathsf{next}) \vee \mathsf{f} = \mathsf{appendArg} \,\rangle \qquad . \end{array}$$

Informally, this relation specifies that the method store its argument (which is passed in via the $\mathsf{appendArg}$ field) at the end of the list linked via field next. More precisely, the relation specifies that the method find an object $\acute{r}$ reachable from $\grave{r}$ (self) via the next field such that $\acute{r}.\mathsf{next}$ is *nil*. The method is to set $\acute{r}.\mathsf{next}$ to the given argument node. The method can modify $\acute{r}.\mathsf{next}$ and can modify the $\mathsf{appendArg}$ field of any object (since this field is used only as a way to pass a parameter to append anyway), but it is not allowed to modify the store in any other way.

To present the example code, I introduce a new command, *isnil*, which tests whether or not the register is *nil*.

Nil Test

$$\frac{E, V \vdash \diamond \qquad E \vdash_{obj} T}{E, V \vdash \mathit{isnil} : T \to \mathit{Boolean} : \\ \quad (\grave{r} = \mathit{nil} \Rightarrow \acute{r} = \mathit{true}) \wedge (\grave{r} \neq \mathit{nil} \Rightarrow \acute{r} = \mathit{false}) \wedge \grave{\sigma} = \acute{\sigma}}$$

(Section 6 discusses expression commands such as *nil* tests.) To write the program text, I assume the following binding powers, from highest to lowest: := . ; <> *with* ... *do* . Now, consider the following command.

$$\begin{array}{l} [\mathit{Node}: \mathsf{next} = \mathit{nil},\ \mathsf{appendArg} = \mathit{nil}, \\ \quad \mathsf{append} = \mathit{with\ self}: \mathit{Node\ do}\ \ \mathsf{appendArg}\ ;\ \ \mathit{with\ n}: \mathit{Node\ do} \\ \qquad \mathit{self}.\mathsf{next}\ ;\ \ \mathit{isnil}\ ; \\ \qquad (\ \ \mathit{self}.\mathsf{next}.\mathsf{appendArg} := n\ ;\ \ \mathsf{append} \\ \qquad \diamond\ \mathit{self}.\mathsf{next} := n\ ;\ \ \mathit{self}\ \ )\ \ ] \end{array} \tag{5}$$

This command allocates and returns a *Node* object whose next and appendArg fields are initially *nil* . The implementation of append starts by giving names to the self object and the method's argument. Then it either calls append recursively on *self* .next or sets *self* .next to the given argument, depending on whether or not *self* .next is *nil* .

With axioms (3) and (4) in the underlying logic, one can prove the following judgement about the given allocation command.

$$E, \emptyset \vdash (5) : \mathit{Object} \to \mathit{Node} : \grave{r} = \grave{r} \tag{6}$$

Though the relation in this judgement ($\grave{r} = \grave{r}$) is trivial, establishing the judgement requires showing that the given implementation of append satisfies its declared specification. I omit the proof, which is straightforward.

I conclude the example with three remarks. First, remember that to reason about a call to a method, in particular the recursive call to append , one uses the specification of the method being called, not its implementation. This makes the reasoning independent of the actual implementation of the callee, which may in fact even be a different implementation than the one shown.

Second, remember that only partial correctness is proved. That is, judgement (6) says that *if* the method terminates, its pre- and post-states will satisfy the specified relation. Indeed, an invocation of method append on an object in a cyclic structure of *Node* objects will not terminate.

Third, the static type of the field *self* .next and the argument *self* .appendArg is *Node* , but the dynamic types of these objects in an execution may be any subtype of *Node* . Note, however, that judgement (6) is independent of the dynamic types of these objects. Indeed, having established judgement (6) means that the method works in every execution. This is because the logic is *sound*, as is shown in the next section.

5 Soundness

This section states a soundness theorem, which proves the correctness of the operational semantics with respect to the axiomatic semantics. I first motivate the soundness theorem, and then state it together with an informal explanation. Some additional formal definitions and the proof itself are found in a SRC Technical Note [14].

As a gentle step in presenting the full soundness theorem, consider the following theorem.

Theorem 0. If one can derive both $E, \emptyset \vdash a : \mathit{Object} \to \mathit{Boolean} : \acute{r} = \mathit{true}$ and $\mathit{nil}, \sigma_0, \emptyset, \emptyset \vdash a \rightsquigarrow r, \sigma, \mu$, then $r = \mathit{true}$.

Here and in the next two theorems, σ_0 denotes a data store that satisfies $\langle\, \forall\, h \in \mathcal{H} :: \mathit{select}(\sigma_0, h, \mathsf{alloc}) = \mathit{false} \,\rangle$, and $\emptyset$ denotes the partial function whose domain is empty. The theorem says that if in an environment E one can prove that a command a satisfies the transition relation $\acute{r} = \mathit{true}$, then any terminating execution of command a from a "reset" state ends with a register value of true.

A simple theorem about the result type of a command is the following.

Theorem 1. If one can derive $E, \emptyset \vdash a : \mathit{Object} \to T : R$ and $E \vdash_{obj} T$ and $\mathit{nil}, \sigma_0, \emptyset, \emptyset \vdash a \rightsquigarrow r, \sigma, \mu$, then the value r has type T, that is, either $r = \mathit{nil}$ or $E \vdash \mu(r)(\mathsf{type}) <: T$.

This theorem says that if one can prove, using the axiomatic semantics, that a command a has final type T, where T is an object type, and one can show that, operationally, the program terminates with a register value of r, then r is a value of type T (that is, it is nil or its allocated type is a subtype of T). This theorem shows the soundness of the type system's treatment of object types.

An interesting theorem that says something about the final object store of a program is the following.

Theorem 2. If one can derive both $E, \emptyset \vdash a : \mathit{Object} \to T : R$ and $\mathit{nil}, \sigma_0, \emptyset, \emptyset \vdash a \rightsquigarrow r, \sigma, \mu$, then $R[\grave{r}, \grave{\sigma}, \acute{r}, \acute{\sigma} := \mathit{nil}, \sigma_0, r, \sigma]$ holds as a first-order predicate.

This theorem says that if one can prove the two judgements about a, then relation R actually describes the relation between the initial and final states.

To prove the theorems above, one needs to prove something stronger. I call the stronger theorem, of which the theorems above are corollaries, the main theorem. The theorem is stated as follows.

Main Theorem. If (7) $E, V \vdash a : T \to U : R$, (8) $r, \sigma, \mu, S \vdash a \rightsquigarrow r', \sigma', \mu'$, (9) $E, \sigma, \mu \Vdash r : T$, (10) $E \Vdash \sigma, \mu$, and (11) $E, V, \sigma, \mu \Vdash S$; then (12) $r, \sigma, r', \sigma', S \Vdash R$, (13) $(\sigma, \mu) \preceq (\sigma', \mu')$, (14) $E, \sigma', \mu' \Vdash r' : U$, and (15) $E \Vdash \sigma', \mu'$.

In the antecedent of this theorem, (7) and (8) express the judgements that have been derived for some command a. One can hope to say something interesting in the conclusion of the theorem only if the execution under consideration is from a "reasonable" state (r, σ, μ) and uses a "reasonable" stack S. Therefore, judgement (9) states that r is a value of type T, judgement (10) says that store pair (σ, μ) *matches* the environment E, and judgement (11) says that S is a *well-typed stack*.

In the conclusion of the theorem, (12) expresses that R does indeed describe the relation between the initial and final states of the execution, and (14) expresses that r' has type U. In addition, to use the theorem as a sufficiently strong induction hypothesis in the proof, (13) says that (σ, μ) is *continued by* (σ', μ'). This property expresses a

kind of monotonicity that holds between two store pairs, the first of which precedes the other in some execution. Also, judgement (15) says that (σ', μ'), like the initial store pair, matches the environment.

By removing (12) from the conclusion of the main theorem, one gets a corollary that expresses that the type system is sound with respect to the operational semantics. Such a corollary follows directly from the main theorem, but could also be proved directly in the same way the main theorem is.

6 Limitations of the Logic

The object construction command is rather awkward. Because it lists method implementations, a method cannot directly construct objects whose type and method implementations are the same as for self. Instead, one can declare object types representing classes, as is done, for example, by Abadi and Cardelli [1] (see SRC TN 1997-025 for an example [14]). One can consider modifying the present logic to remove the limitation from the object construction command. For example, like in common class-based object-oriented languages, one can extend the program environment to include method implementations. One must then have a "link-time" check that ensures that every method that may be called by the program at run-time has an implementation. Or, like in common object-based languages, one can add a construct for cloning objects or their method implementations.

Another omission from the present logic is the ability to compare objects for equality. Just like one would expect to add primitive types like integers to the present logic, one would expect to add more general expressions, including comparison expressions.

A logic of programs provides a connection between programs and their specifications. In the present logic, method declarations contain specifications that are given simply as transition relations. Transition relations are not practically suited for writing down method specifications, because they are painfully explicit. Specification features like modifies clauses and abstract fields would remedy the situation, but lie outside the scope of this paper. To mention some work in this area, Lano and Haughton [10] have surveyed object-oriented specifications, and my thesis [12] shows how to deal with modifies clauses and data abstraction in modular, object-oriented programs. The logic for POOL [4] includes some specification features that can be used to state properties of recursive data structures.

7 Summary

I have presented a sound logic for object-oriented programs whose commands are imperative and whose objects are references to data fields. The programming notation requires that types, fields, and methods be declared in the environment before they can be used in a program. The main contributions of the paper are the logic itself, the soundness theorem, and the way that types are handled, which makes the subtype relation and the admission of recursive object types trivial.

Acknowledgements. I am grateful to Martín Abadi, Luca Cardelli, Greg Nelson, and Raymie Stata for helpful comments on the logic and the presentation thereof.

References

1. Martín Abadi and Luca Cardelli. *A Theory of Objects*. Springer-Verlag, 1996.
2. Martín Abadi and K. Rustan M. Leino. A logic of object-oriented programs. In *Theory and Practice of Software Development: Proceedings / TAPSOFT '97, 7th International Joint Conference CAAP/FASE*, volume 1214 of *Lecture Notes in Computer Science*, pages 682–696. Springer, April 1997.
3. Roberto M. Amadio and Luca Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, September 1993.
4. Pierre America and Frank de Boer. Reasoning about dynamically evolving process structures. *Formal Aspects of Computing*, 6(3):269–316, 1994.
5. Craig Chambers. The Cecil language: Specification & rationale, version 2.1, March 7, 1997. Available from `http://www.cs.washington.edu/research/projects/cecil/www/Papers/cecil-spec.html`, 1997.
6. Krishna Kishore Dhara and Gary T. Leavens. Forcing behavioral subtyping through specification inheritance. Technical Report TR #95-20c, Iowa State University, Department of Computer Science, 1997.
7. Edsger W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, NJ, 1976.
8. James Gosling, Bill Joy, and Guy Steele. *The Java™ Language Specification*. Addison-Wesley, 1996.
9. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580,583, October 1969.
10. Kevin Lano and Howard Haughton. *Object-Oriented Specification Case Studies*. Prentice Hall, New York, 1994.
11. Gary Todd Leavens. *Verifying Object-Oriented Programs that Use Subtypes*. PhD thesis, MIT Laboratory for Computer Science, February 1989. Available as Technical Report MIT/LCS/TR-439.
12. K. Rustan M. Leino. *Toward Reliable Modular Programs*. PhD thesis, California Institute of Technology, January 1995. Available as Technical Report Caltech-CS-TR-95-03.
13. K. Rustan M. Leino. Ecstatic: An object-oriented programming language with an axiomatic semantics. In *The Fourth International Workshop on Foundations of Object-Oriented Languages*, January 1997. Proceedings available from `http://www.cs.indiana.edu/hyplan/pierce/fool/`.
14. K. Rustan M. Leino. Recursive object types in a logic of oject-oriented programs. Technical Note 1997-025a, Digital Equipment Corporation Systems Research Center, January 1998.
15. K. Rustan M. Leino and Greg Nelson. Object-oriented guarded commands. Internal manuscript KRML 50, Digital Equipment Corporation Systems Research Center, March 1995.
16. Greg Nelson. Verifying reachability invariants of linked structures. *Conference Record of the Tenth Annual ACM Symposium on Principles of Programming Languages*, pages 38–47, January 1983.
17. Greg Nelson, editor. *Systems Programming with Modula-3*. Series in Innovative Technology. Prentice-Hall, Englewood Cliffs, NJ, 1991.
18. Arnd Poetzsch-Heffter and Peter Müller. A logic for the verification of object-oriented programs. In R. Berghammer and F. Simon, editors, *Programming Languages and Fundamentals of Programming*. Christian Albrechts-Universität Kiel, 1997.