

A Logic of Object-Oriented Programs

Martín Abadi¹ and K. Rustan M. Leino²

¹ Computer Science Department, University of California at Santa Cruz, CA, USA

² Microsoft Research, Redmond, WA, USA

Abstract. We develop a logic for reasoning about object-oriented programs. The logic is for a language with an imperative semantics and aliasing, and accounts for self-reference in objects. It is much like a type system for objects with subtyping, but our specifications go further than types in detailing pre- and postconditions. We intend the logic as an analogue of Hoare logic for object-oriented programs. Our main technical result is a soundness theorem that relates the logic to a standard operational semantics.

1 Introduction

In the realm of procedural programming, Floyd and Hoare defined two of the first logics of programs [9, 11]; many later formalisms and systems built on their ideas, and addressed difficult questions of concurrency and data abstraction, for example. An analogous development has not taken place in object-oriented programming. Although there is much formal work on objects (see Section 6), the previous literature on objects does not seem to contain an analogue for Floyd’s logic or Hoare’s logic. In our opinion, this is an important gap in the understanding of object-oriented programming languages.

Roughly imitating Hoare, we develop a logic for the specification and verification of object-oriented programs. We focus on elementary goals: we are interested in logical reasoning about pre- and postconditions of programs written in a basic object-oriented programming language (a variant of the calculi of Abadi and Cardelli [1]). Like Hoare, we deal with partial correctness, not with termination.

The programming language presents many interesting and challenging features of common object-oriented languages. In particular, the operational semantics of the language is imperative and allows aliasing. Objects have fields and methods, and the self variable permits self-reference. At the type level, the type of an object lists the types of its fields and the result types of its methods; a subtyping relation supports subsumption and inheritance. However, the language lacks many class-based constructs common in practice. It also lacks “advanced” features, like concurrency; some of these features have been studied in the literature (e.g., see [6, 14, 42]).

Much like Hoare logic, our logic includes one rule for reasoning about pre- and postconditions for each of the constructs of the programming language. In order to formulate these rules, we introduce object specifications. An object specification is a generalization of an object type: it lists the specifications of fields,

the specifications of the methods' results, and also gives the pre/postcondition descriptions of the methods.

Some of the main advantages of Hoare logic are its formal precision and its simplicity. These advantages make it possible to study Hoare logic, and for example to prove its soundness and completeness; they also make it easier to extend and to implement Hoare logic. We aim to develop a logic with some of those same advantages. Our rules are not quite as simple as Hoare's, in part because of aliasing, and in part because objects are more expressive than first-order procedures and give some facilities for higher-order programming (see [5, 3]). However, our rules are precise; in particular, we are able to state and to prove a soundness theorem. We do not know of any previous, equivalent soundness theorem in the object-oriented literature.

In the next section we describe the programming language. In Section 3 we develop a logic for this language, and in Section 4 we give some examples of the use of this logic in verification. In Section 5, we discuss soundness and completeness with respect to the operational semantics of Section 2. Finally, in Sections 6 and 7, we review some related work, discuss possible extensions of our own work, and conclude. A preliminary version of this work has been presented in a conference [2]. The present version incorporates several improvements; in particular, it includes proofs.

2 The Language

In this section we define a small object-oriented language similar to the calculi of Abadi and Cardelli. Those calculi have few syntactic forms, but are quite expressive. They are object-based; they do not include primitives for classes and inheritance, which can be simulated using simpler constructs.

We give the syntax of our language, its operational semantics, and a set of type rules. These aspects of the language are (intentionally) not particularly novel or exotic; we describe them only as background for the rest of the paper.

2.1 Syntax and Operational Semantics

We assume we are given a set \mathcal{V} of program variables (written x , y , z , and w possibly with subscripts), a set \mathcal{F} of field names (written f and g , possibly with subscripts), and a set \mathcal{M} of method names (written m , possibly with subscripts). These sets are disjoint.

The grammar of the language is:

$a, b ::= x$	variables
$false \mid true$	constants
$if\ x\ then\ a_0\ else\ a_1$	conditional
$let\ x = a\ in\ b$	let
$[f_i = x_i^{i \in 1..n},\ m_j = \zeta(y_j)b_j^{j \in 1..m}]$	object construction
$x.f$	field selection
$x.m$	method invocation
$x.f := y$	field update

Throughout, we assume that the names f_i and m_j are all distinct in the construct $[f_i = x_i \text{ }^{i \in 1..n}, m_j = \varsigma(y_j)b_j \text{ }^{j \in 1..m}]$, and we allow the renaming of bound variables in all expressions.

Informally, the semantics of the language is as follows:

- Variables are identifiers; they are not mutable: $x := a$ is not a legal statement. This restriction is convenient but not fundamental. (We can simulate assignment by binding a variable to an object with a single field and updating that field.)
- *false* and *true* evaluate to themselves.
- *if x then a₀ else a₁* evaluates a_0 if x is *true* and evaluates a_1 if x is *false*.
- *let x = a in b* evaluates a and then evaluates b with x bound to the result of a . We define $a ; b$ as a shorthand for *let x = a in b* where x does not occur free in b .
- $[f_i = x_i \text{ }^{i \in 1..n}, m_j = \varsigma(y_j)b_j \text{ }^{j \in 1..m}]$ creates and returns a new object with fields f_i and methods m_j . The initial value for the field f_i is the value of x_i . The method m_j is set to $\varsigma(y_j)b_j$, where ς is a binder, y_j is a variable (the self parameter of the method), and b_j is a program (the body of the method).
- Fields can be both selected and updated. In the case of selection ($x.f$), the value of the field is returned; in the case of update ($x.f := y$), the value of the object is returned.
- When a method of an object is invoked ($x.m$), its self variable is bound to the object itself and the body of the method is executed. The method does not have any explicit parameters besides the self variable; however, additional parameters can be passed via the fields of the object.

Objects are references (rather than records), and the semantics allows aliasing. For example, the program fragment

$$\text{let } x = [f = z_0] \text{ in let } y = x \text{ in } (x.f := z_1 ; y.f)$$

allocates some storage, creates two references to it (x and y), updates the storage through x , and then reads it through y , returning z_1 .

In order to formalize the operational semantics, we use some notations for partial functions. We write $A \rightarrow B$ for the set of partial functions from A to B . We write \emptyset for the totally undefined partial function. When $f \in A \rightarrow B$, $a \in A$, and $b \in B$, we write $f.(a \mapsto b)$ for the function that coincides with f except possibly at a , and that maps a to b . When $a_i \in A \text{ }^{i \in 1..n}$ are distinct and $b_i \in B \text{ }^{i \in 1..n}$, we write $(a_i \mapsto b_i \text{ }^{i \in 1..n})$ for the function in $A \rightarrow B$ that maps a_i to b_i for $i \in 1..n$ and is otherwise undefined.

The formal operational semantics is in terms of stacks and stores. A stack maps variables to booleans or references. A store maps object fields to booleans or references and maps object methods to closures. We write $\sigma, S \vdash b \rightsquigarrow v, \sigma'$ to mean that, given the initial store σ and the stack S , executing the program b leads to the result v and to the final store σ' .

We define the notions of store, stack, and result as follows:

- We assume we are given a set of *object names* \mathcal{H} . The set of *results* \mathcal{R} is $\mathcal{H} \cup \{false, true\}$.
- A *stack* is a function in $\mathcal{V} \rightarrow \mathcal{R}$.
- A *method closure* is a triple of a variable $x \in \mathcal{V}$ (standing for self), a program b , and a stack S ; we write it $\langle \varsigma(x)b, S \rangle$. The set of method closures is \mathcal{C} .
- A *store* is a function σ in $\mathcal{H} \rightarrow ((\mathcal{F} \cup \mathcal{M}) \rightarrow (\mathcal{R} \cup \mathcal{C}))$. There is a condition on σ : if $h \in \mathcal{H}$, $f \in \mathcal{F}$, and $\sigma(h)(f)$ is defined, then $\sigma(h)(f) \in \mathcal{R}$; if $h \in \mathcal{H}$, $m \in \mathcal{M}$, and $\sigma(h)(m)$ is defined, then $\sigma(h)(m) \in \mathcal{C}$. In other words, field names are mapped to results and method names to closures.

The operational semantics is represented with a set of rules, given below. According to these rules, a variable x reduces to its value in the stack, without change in the store. The constants *false* and *true* reduce to themselves, without change in the store. The execution of a conditional expression consists in evaluating the guard and, depending on the outcome of this evaluation, evaluating one of the branches. The *let* construct evaluates an expression, binds a local variable to the result of that evaluation, and then evaluates another expression. The execution of an object construction requires evaluating the fields, constructing method closures, picking a new location, and mapping that location to an appropriate suite of fields and methods. The execution of a field selection on an object requires evaluating the object and then extracting the value of the appropriate field from the store. The execution of a method invocation is similar, but there the value returned is the result of evaluating the appropriate method body with an extended stack that maps the self variable to the value of the object. Finally, the execution of a field update modifies the store and returns the value of the object being affected.

Operational semantics

Variables

$$\frac{S(x) = v}{\sigma, S \vdash x \rightsquigarrow v, \sigma}$$

Constants

$$\frac{}{\sigma, S \vdash false \rightsquigarrow false, \sigma} \quad \frac{}{\sigma, S \vdash true \rightsquigarrow true, \sigma}$$

Conditional

$$\frac{S(x) = false \quad \sigma, S \vdash a_1 \rightsquigarrow v, \sigma'}{\sigma, S \vdash \text{if } x \text{ then } a_0 \text{ else } a_1 \rightsquigarrow v, \sigma'}$$

$$\frac{S(x) = true \quad \sigma, S \vdash a_0 \rightsquigarrow v, \sigma'}{\sigma, S \vdash \text{if } x \text{ then } a_0 \text{ else } a_1 \rightsquigarrow v, \sigma'}$$

Let

$$\frac{\sigma, S \vdash a \rightsquigarrow v, \sigma' \quad \sigma', S.(x \mapsto v) \vdash b \rightsquigarrow v', \sigma''}{\sigma, S \vdash \text{let } x = a \text{ in } b \rightsquigarrow v', \sigma''}$$

Object construction

$$\frac{S(x_i) = v_i \quad i \in 1..n \quad h \notin \text{dom}(\sigma) \quad h \in \mathcal{H} \quad \sigma' = \sigma.(h \mapsto (f_i \mapsto v_i \quad i \in 1..n, m_j \mapsto \langle \zeta(y_j) b_j, S \rangle \quad j \in 1..m))}{\sigma, S \vdash [f_i = x_i \quad i \in 1..n, m_j = \zeta(y_j) b_j \quad j \in 1..m] \rightsquigarrow h, \sigma'}$$

Field selection

$$\frac{S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(f) = v}{\sigma, S \vdash x.f \rightsquigarrow v, \sigma}$$

Method invocation

$$\frac{S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(m) = \langle \zeta(y) b, S' \rangle \quad \sigma, S'.(y \mapsto h) \vdash b \rightsquigarrow v, \sigma'}{\sigma, S \vdash x.m \rightsquigarrow v, \sigma'}$$

Field update

$$\frac{S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(f) \text{ is defined} \quad S(y) = v \quad \sigma' = \sigma.(h \mapsto \sigma(h).(f \mapsto v))}{\sigma, S \vdash x.f := y \rightsquigarrow h, \sigma'}$$

The judgment $\sigma, S \vdash b \rightsquigarrow v, \sigma'$ represents only computations that terminate with a result, not computations that do not terminate or that terminate with an error. For example, intuitively, the execution of *let* $x = [m = \zeta(y) \text{ true}]$ *in* $x.m$ terminates with the output *true*. Formally, we can derive

$$\sigma, S \vdash \text{let } x = [m = \zeta(y) \text{ true}] \text{ in } x.m \rightsquigarrow \text{true}, \sigma'$$

for all σ and S and for some σ' . On the other hand, intuitively, the execution of *let* $x = \text{true}$ *in* $x.m$ yields an error, while the execution of *let* $x = [m = \zeta(x) x.m]$ *in* $x.m$ does not terminate. Formally,

$$\sigma, S \vdash \text{let } x = \text{true} \text{ in } x.m \rightsquigarrow v, \sigma'$$

and

$$\sigma, S \vdash \text{let } x = [m = \zeta(x) x.m] \text{ in } x.m \rightsquigarrow v, \sigma'$$

cannot be derived for any σ , S , v , and σ' . The search for a derivation of the former judgment “gets stuck”, while the search for a derivation of the latter judgment diverges.

We have defined a small language in order to simplify the presentation of our rules. In examples, we sometimes extend the syntax with additional, standard constructs, such as integers. The rules for such constructs are straightforward.

2.2 Types

We present a first-order type system for our language. The types are *Bool* and object types, which have the form:

$$[f_i: A_i \ i \in 1..n, \ m_j: B_j \ j \in 1..m]$$

This is the type of objects with a field f_i of type A_i , for $i \in 1..n$, and with a method m_j with result type B_j , for $j \in 1..m$. The order of the components does not matter.

The type system includes a reflexive and transitive subtyping relation. A longer object type is a subtype of a shorter one, and in addition object types are covariant in the result types of methods. More precisely, the type $[f_i: A_i \ i \in 1..n+p, \ m_j: B_j \ j \in 1..m+q]$ is a subtype of $[f_i: A_i \ i \in 1..n, \ m_j: B'_j \ j \in 1..m]$ provided B_j is a subtype of B'_j , for $j \in 1..m$. Thus, object types are invariant in the types of fields; this invariance is essential for soundness [1].

Formally, we write $\vdash A$ to express that A is a well-formed type, and $\vdash A <: A'$ to express that A is a subtype of A' . We have the rules:

Well-formed types

$$\frac{}{\vdash Bool} \quad \frac{\vdash A_i \ i \in 1..n \quad \vdash B_j \ j \in 1..m}{\vdash [f_i: A_i \ i \in 1..n, \ m_j: B_j \ j \in 1..m]}$$

Subtypes

$$\frac{}{\vdash Bool <: Bool} \quad \frac{\vdash A_i \ i \in 1..n+p \quad \vdash B_j <: B'_j \ j \in 1..m \quad \vdash B_j \ j \in m+1..m+q}{\vdash [f_i: A_i \ i \in 1..n+p, \ m_j: B_j \ j \in 1..m+q] <: [f_i: A_i \ i \in 1..n, \ m_j: B'_j \ j \in 1..m]}$$

A typing environment is a (possibly empty) list of pairs $x:A$, where x is a variable and A is a type. The variables of each environment are distinct. We write \emptyset for the empty environment, and say that x is in E when it appears in some pair $x:A$ in E . We write $E \vdash \diamond$ to express that E is a well-formed typing environment. We have two rules for forming typing environments:

Well-formed typing environments

$$\frac{}{\emptyset \vdash \diamond} \quad \frac{E \vdash \diamond \quad \vdash A \quad x \text{ not in } E}{E, x: A \vdash \diamond}$$

We write $E \vdash a : A$ to express that, in environment E , program a has type A . There is one typing rule for each construct, and an additional rule for subsumption. We write $\stackrel{sym}{=}$ for the relation of syntactic equality (up to reordering of object components).

Well-typed programs

Subsumption

$$\frac{\vdash A <: A' \quad E \vdash a : A}{E \vdash a : A'}$$

Variables

$$\frac{E, x: A, E' \vdash \diamond}{E, x: A, E' \vdash x : A}$$

Constants

$$\frac{E \vdash \diamond}{E \vdash \text{false} : \text{Bool}} \quad \frac{E \vdash \diamond}{E \vdash \text{true} : \text{Bool}}$$

Conditional

$$\frac{E \vdash x : \text{Bool} \quad E \vdash a_0 : A \quad E \vdash a_1 : A}{E \vdash \text{if } x \text{ then } a_0 \text{ else } a_1 : A}$$

Let

$$\frac{E \vdash a : A \quad E, x: A \vdash b : B}{E \vdash \text{let } x = a \text{ in } b : B}$$

Object construction for $A \stackrel{\text{syn}}{=} [f_i: A_i \ i \in 1..n, \ m_j: B_j \ j \in 1..m]$

$$\frac{E \vdash \diamond \quad E \vdash x_i : A_i \ i \in 1..n \quad E, y_j: A \vdash b_j : B_j \ j \in 1..m}{E \vdash [f_i = x_i \ i \in 1..n, \ m_j = \varsigma(y_j)b_j \ j \in 1..m] : A}$$

Field selection

$$\frac{E \vdash x : [f: A]}{E \vdash x.f : A}$$

Method invocation

$$\frac{E \vdash x : [m: B]}{E \vdash x.m : B}$$

Field update for $A \stackrel{\text{syn}}{=} [f_i: A_i \ i \in 1..n, \ m_j: B_j \ j \in 1..m]$

$$\frac{E \vdash x : A \quad k \in 1..n \quad E \vdash y : A_k}{E \vdash x.f_k := y : A}$$

This type system is much like those of common programming languages in that it is independent of verification rules. In particular, types are not automatically associated with specifications, and subtyping does not impose any “behavioral” constraints (as in the work of Liskov and Wing [23], for example). However, as the next section explains, specifications are a generalization of types.

3 Verification

In this section, which is the core of the paper, we give rules for verifying object-oriented programs written in the language of Section 2. We start with an informal explanation of our approach.

3.1 Transition Relations

The purpose of our verification rules is to allow reasoning about pre- and postconditions. These pre- and postconditions concern the initial and final stores, the stack, and the result of the execution of a given program.

In our rules, we express pre- and postconditions in formulas of standard, untyped first-order logic that we call *transition relations*. These formulas mention the unary predicates alloc and alloc , two binary functions δ and σ , and the special variable r (which is not in the set \mathcal{V} of program variables). Intuitively, $\delta(x, f)$ is the value of field f of object x before the execution, and $\sigma(x, f)$ is its value after the execution. Similarly, $\mathit{alloc}(x)$ and $\mathit{alloc}(x)$ indicate whether x has been allocated before and after the execution. Finally, the variable r represents the result of the execution.

For example, we may want to prove that, after any execution of the program $x.f := y$, the result is x and the field f of x equals y . We can express this with the transition relation $r = x \wedge \sigma(x, f) = y$. As a second example, we may want to prove that, after any execution of $x.f$, the result equals the initial value of the field f of x , and that the store is not changed by the execution. This statement is captured by the transition relation $r = \delta(x, f) \wedge (\forall y, z. \delta(y, z) = \sigma(y, z) \wedge (\mathit{alloc}(y) \equiv \mathit{alloc}(y)))$.

We work in standard first-order logic, so the functions δ and σ are total. Hence, $\delta(x, f)$ is defined even if $\mathit{alloc}(x)$ does not hold, and $\sigma(x, f)$ is defined even if $\mathit{alloc}(x)$ does not hold. In those cases, the values of $\delta(x, f)$ and $\sigma(x, f)$ are not important. Similarly, the values of expressions such as $\delta(f, x)$ and $\mathit{alloc}(f)$, which are intuitively meaningless, are not important.

Given a program, a transition relation is much like a Hoare triple from the point of view of expressiveness. For example, a transition relation such as $(\delta(x, f) = \delta(x, g)) \Rightarrow (\sigma(x, f) = \sigma(x, g))$ can be understood as assuming a precondition $(\delta(x, f) = \delta(x, g))$ and asserting a postcondition $(\sigma(x, f) = \sigma(x, g))$. However, the precondition and postcondition are given by separate formulas in a Hoare triple, while there is no such formal separation in a transition relation. This difference is largely a matter of convenience.

Formally, we write that T is a transition relation to mean that T is a well-formed formula of the standard, untyped first-order logic, made up only of:

- the constants *false* and *true*;
- the variable r , the binary functions δ and σ , and the unary predicates alloc and alloc ;
- constants for field names (such as f);
- other variables (such as x);
- the usual logical connectives \neg , \wedge , and \forall , and the equality predicate $=$ (from which \vee , \Rightarrow , \equiv , \exists , and \neq can be defined as abbreviations).

The grammar for transition relations is thus:

$$\begin{aligned} T &::= e_0 = e_1 \mid \mathit{alloc}(e) \mid \mathit{alloc}(e) \mid \neg T \mid T_0 \wedge T_1 \mid (\forall x. T) \\ e &::= \mathit{false} \mid \mathit{true} \mid r \mid x \mid f \mid \delta(e_0, e_1) \mid \sigma(e_0, e_1) \end{aligned}$$

3.2 Specifications and Subspecifications

In order to permit reasoning about pre- and postconditions, our verification rules also deal with *specifications*, which generalize types. A specification can be either *Bool* or an *object specification*, of the form:

$$[f_i: A_i \quad i \in 1..n, \quad m_j: \varsigma(y_j)B_j :: T_j \quad j \in 1..m]$$

where each A_i and B_j is a specification, and each T_j is a transition relation. The variable y_j is bound in B_j and T_j . Informally, an object satisfies the specification $[f_i: A_i \quad i \in 1..n, \quad m_j: \varsigma(y_j)B_j :: T_j \quad j \in 1..m]$ if, for $i \in 1..n$, it has a field f_i that satisfies specification A_i , and, for $j \in 1..m$, it has a method m_j with a result that satisfies B_j and whose execution satisfies T_j when y_j equals self. We may think of B_j as a predicate on the result, and then we may read $B_j :: T_j$ as the conjunction of that predicate and T_j . As for object types, the order of the components of object specifications does not matter.

Just like there is a subtyping relation on types, there is a *subspecification* relation on specifications. This relation is reflexive and transitive. A longer object specification is a subspecification of a shorter one, and in addition object specifications are covariant in the result specifications and in the transition relations for methods. Intuitively, when A and A' are object specifications, A is a subspecification of A' only if any object that satisfies A also satisfies A' .

3.3 Rules for Specifications

In our rules for specifications, we use several judgments analogous to those introduced for types in Section 2.2, and in those cases we use similar notations but with a \Vdash instead of a \vdash . In particular, we write $\Vdash A$ to express that A is a well-formed specification, and $\Vdash A <: A'$ to express that A is a subspecification of A' . The following rules for specifications generalize the corresponding rules for types:

Well-formed specifications

$$\frac{\Vdash A_i \quad i \in 1..n \quad \Vdash B_j \quad j \in 1..m \quad T_j \text{ is a transition relation } j \in 1..m}{\Vdash \text{Bool}} \quad \frac{}{\Vdash [f_i: A_i \quad i \in 1..n, \quad m_j: \varsigma(y_j)B_j :: T_j \quad j \in 1..m]}$$

Subspecifications

$$\frac{\Vdash \text{Bool} <: \text{Bool}}{\Vdash \text{Bool} <: \text{Bool}}$$

$$\frac{\Vdash A_i \quad i \in 1..n+p \quad \Vdash B_j <: B'_j \quad j \in 1..m \quad \Vdash B_j \quad j \in m+1..m+q \quad \Vdash_{\text{fol}} T_j \Rightarrow T'_j \quad j \in 1..m \quad T_j \text{ is a transition relation } j \in 1..m+q \quad T'_j \text{ is a transition relation } j \in 1..m}{\Vdash [f_i: A_i \quad i \in 1..n+p, \quad m_j: \varsigma(y_j)B_j :: T_j \quad j \in 1..m+q] <: [f_i: A_i \quad i \in 1..n, \quad m_j: \varsigma(y_j)B'_j :: T'_j \quad j \in 1..m]}$$

In this last rule, \Vdash_{fol} represents provability in first-order logic with the standard axioms for $=$ and the axioms $false \neq true$ and $f \neq g$ for every pair of different field names f and g .

3.4 Specification Environments

A *specification environment* is much like a typing environment, except that it contains specifications instead of types. We write $E \Vdash \diamond$ to mean that E is a well-formed specification environment. We have the rules:

Well-formed specification environments

$$\frac{}{\emptyset \Vdash \diamond} \quad \frac{E \Vdash \diamond \quad E \Vdash A \quad x \text{ not in } E}{E, x : A \Vdash \diamond}$$

Here, given a well-formed specification environment E , we write $E \Vdash A$ to mean $\Vdash A$ and that all the free program variables of A are in E . We omit the obvious rule for this judgment. Similarly, when all the free program variables of a transition relation T are in E , we write:

$$E \Vdash T \text{ is a transition relation}$$

In order to formulate the verification rules, we introduce the judgment:

$$E \Vdash a : A :: T$$

This judgment states that, in specification environment E , the execution of a satisfies the transition relation T , and its result satisfies the specification A .

For this judgment, there is one rule per construct plus a subsumption rule; the rules are all given below. The rules guarantee that, whenever $E \Vdash a : A :: T$ is provable, all the free program variables of a , A , and T are in E . The rules have interesting similarities both with the rules of the operational semantics (Section 2.1) and with the typing rules (Section 2.2). The treatment of transition relations reiterates parts of the operational semantics, while the treatment of specifications generalizes that of types.

The subsumption rule enables us to weaken a specification and a transition relation; it generalizes both the subsumption rule for typing and the standard rule of consequence from Hoare logic. The rule for *if-then-else* allows the replacement of the boolean guard with its value in reasoning about each of the alternatives. The rule for *let* achieves sequencing by representing an intermediate state with the auxiliary binary function $\tilde{\sigma}$ and unary predicate *alloc*. (Note that the grammar for transition relations does not allow the auxiliary symbols $\tilde{\sigma}$ and *alloc*; however, \Vdash_{fol} applies to any first-order formula, even one that is not a transition relation.) The variable x bound by *let* cannot escape because of the hypotheses that $E \Vdash B$ and that $E \Vdash T''$ is a transition relation. The rule for object construction has a complicated transition relation, but this transition relation directly reflects the operational semantics; the introduction of an object specification requires the verification of the methods of the new object. The rule

for method invocation takes advantage of an object specification for yielding a specification and a transition relation; in these, the formal self is replaced with the actual self. The remaining rules are mostly straightforward.

In several rules, we use transition relations of the form $Res(e)$, where e is a term; $Res(e)$ is defined by:

$$Res(e) \triangleq r = e \wedge (\forall x, y. \delta(x, y) = \acute{\sigma}(x, y) \wedge (alloc(x) \equiv alloc(x)))$$

and it means that the result is e and that the store does not change. We also write $u_1[u_2/u_3]$ for the result of substituting u_2 for u_3 in u_1 .

Well-specified programs

Subsumption

$$\frac{\begin{array}{l} \Vdash A <: A' \quad \Vdash_{fol} T \Rightarrow T' \quad E \Vdash a : A :: T \\ E \Vdash A' \quad E \Vdash T' \text{ is a transition relation} \end{array}}{E \Vdash a : A' :: T'}$$

Variables

$$\frac{E, x : A, E' \Vdash \diamond}{E, x : A, E' \Vdash x : A :: Res(x)}$$

Constants

$$\frac{E \Vdash \diamond}{E \Vdash false : Bool :: Res(false)} \quad \frac{E \Vdash \diamond}{E \Vdash true : Bool :: Res(true)}$$

Conditional

$$\frac{\begin{array}{l} E \Vdash x : Bool :: Res(x) \\ E \Vdash a_0 : A_0 :: T_0 \quad A_0[true/x] \stackrel{syn}{=} A[true/x] \quad T_0[true/x] \stackrel{syn}{=} T[true/x] \\ E \Vdash a_1 : A_1 :: T_1 \quad A_1[false/x] \stackrel{syn}{=} A[false/x] \quad T_1[false/x] \stackrel{syn}{=} T[false/x] \end{array}}{E \Vdash \text{if } x \text{ then } a_0 \text{ else } a_1 : A :: T}$$

Let

$$\frac{\begin{array}{l} E \Vdash a : A :: T \quad E, x : A \Vdash b : B :: T' \\ E \Vdash B \quad E \Vdash T'' \text{ is a transition relation} \\ \Vdash_{fol} T[\delta/\acute{\sigma}, alloc/alloc, x/r] \wedge T'[\delta/\acute{\sigma}, alloc/alloc] \Rightarrow T'' \end{array}}{E \Vdash \text{let } x = a \text{ in } b : B :: T''}$$

Object construction for $A \stackrel{syn}{=} [f_i : A_i \text{ }^{i \in 1..n}, m_j : \zeta(y_j) B_j :: T_j \text{ }^{j \in 1..m}]$

$$\frac{E \Vdash \diamond \quad E \Vdash x_i : A_i :: Res(x_i) \text{ }^{i \in 1..n} \quad E, y_j : A \Vdash b_j : B_j :: T_j \text{ }^{j \in 1..m}}{E \Vdash [f_i = x_i \text{ }^{i \in 1..n}, m_j = \zeta(y_j) b_j \text{ }^{j \in 1..m}] : A :: \neg alloc(r) \wedge alloc(r) \wedge (\forall z. z \neq r \Rightarrow (alloc(z) \equiv alloc(z))) \wedge \acute{\sigma}(r, f_1) = x_1 \wedge \dots \wedge \acute{\sigma}(r, f_n) = x_n \wedge (\forall z, w. z \neq r \Rightarrow \delta(z, w) = \acute{\sigma}(z, w))}$$

Field selection

$$\frac{E \Vdash x : [f: A] :: Res(x)}{E \Vdash x.f : A :: Res(\dot{\sigma}(x, f))}$$

Method invocation

$$\frac{E \Vdash x : [m: \varsigma(y)B :: T] :: Res(x)}{E \Vdash x.m : B[x/y] :: T[x/y]}$$

Field update for $A \stackrel{syn}{=} [f_i: A_i \text{ }^{i \in 1..n}, m_j: \varsigma(z_j)B_j :: T_j \text{ }^{j \in 1..m}]$

$$\frac{E \Vdash x : A :: Res(x) \quad k \in 1..n \quad E \Vdash y : A_k :: Res(y)}{E \Vdash x.f_k := y : A :: \begin{aligned} &r = x \wedge \dot{\sigma}(x, f_k) = y \wedge \\ &(\forall z, w. \neg(z = x \wedge w = f_k) \Rightarrow \dot{\sigma}(z, w) = \dot{\sigma}(z, w)) \wedge \\ &(\forall z. alloc(z) \equiv alloc(z)) \end{aligned}}$$

4 Examples

We discuss a few instructive examples, some of them with derivations. From now on, we use some abbreviations, allowing general expressions to appear where the grammar requires a variable. In case a , $a_i \text{ }^{i \in 1..n}$, and b are not variables, we define:

$$\begin{aligned} & \text{if } b \text{ then } a_0 \text{ else } a_1 \triangleq \text{let } x = b \text{ in if } x \text{ then } a_0 \text{ else } a_1 \\ [f_i = a_i \text{ }^{i \in 1..n}, m_j = \varsigma(y_j)b_j \text{ }^{j \in 1..m}] \triangleq & \text{let } x_1 = a_1 \text{ in } \dots \text{let } x_n = a_n \text{ in} \\ & [f_i = x_i \text{ }^{i \in 1..n}, m_j = \varsigma(y_j)b_j \text{ }^{j \in 1..m}] \\ & a.f \triangleq \text{let } x = a \text{ in } x.f \\ & a.m \triangleq \text{let } x = a \text{ in } x.m \\ & a.f := b \triangleq \text{let } x = a \text{ in} \\ & (x.f ; \text{let } y = b \text{ in } x.f := y) \end{aligned}$$

where the variables x and $x_i \text{ }^{i \in 1..n}$ are fresh. Rules for these abbreviations can be derived directly from the rules for the language proper. For example, for field selection, we may use the rule:

$$\frac{E \Vdash a : [f: A] :: T}{E \Vdash a.f : A :: (\exists x. T[x/r] \wedge Res(\dot{\sigma}(x, f)))}$$

4.1 Field Update and Selection

Our first example concerns the program:

$$([f = false].f := true).f$$

This program constructs an object with one field, f , whose initial value is *false*. It then updates the value of the field to *true*. Finally, a field selection retrieves the new value of the field.

Using our rules, we can prove that $r = \text{true}$ holds upon termination of this program. Formally, we can derive the judgment:

$$\emptyset \Vdash ([f = \text{false}].f := \text{true}).f : \text{Bool} :: (r = \text{true})$$

4.2 Aliasing

The following three programs exhibit the rôle of aliasing:

let $x = [f = \text{false}]$ *in* *let* $y = [g = \text{false}]$ *in* ($y.g := \text{true} ; x.f$)

let $x = [f = \text{false}]$ *in* *let* $y = [f = \text{false}]$ *in* ($y.f := \text{true} ; x.f$)

let $x = [f = \text{true}]$ *in* *let* $y = x$ *in* ($y.f := \text{false} ; x.f$)

For each of these programs we can verify that $r = \text{false}$. The first program shows that an update of a field g has no effect on another field f . The second program shows that separately constructed objects have different fields, even if those fields have the same name. The third program shows that an update of a field of an aliased object can be seen through all the aliases.

4.3 Method Invocations and Recursion

The next example illustrates the use of method invocation; it shows how object specifications play the rôle of loop invariants for recursive method invocations.

We consider an object-oriented implementation of Euclid's algorithm for computing greatest common divisors. This implementation uses an object with two fields, f and g , and a method m :

$$\begin{aligned} & [f = 1, g = 1, \\ & m = \varsigma(y) \text{ if } y.f < y.g \text{ then } (y.g := y.g - y.f ; y.m) \\ & \quad \text{else if } y.g < y.f \text{ then } (y.f := y.f - y.g ; y.m) \\ & \quad \text{else } y.f] \end{aligned}$$

Setting f and g to two positive integer values and then invoking the method m has the effect of reducing both f and g to the greatest common divisor of those two values.

We can prove that this object satisfies the following specification:

$$\begin{aligned} & [f : \text{Nat}, g : \text{Nat}, \\ & m : \varsigma(y) \text{ Nat} :: 1 \leq \delta(y, f) \wedge 1 \leq \delta(y, g) \Rightarrow \\ & \quad r = \delta(y, f) \wedge r = \delta(y, g) \wedge r = \text{gcd}(\delta(y, f), \delta(y, g))] \end{aligned}$$

(The proof relies on the addition of standard axioms about integers to the underlying first-order logic.) In verifying the body of the method m , we can use the specification of m , recursively. The derivation included in the next example demonstrates how a method specification can be used recursively in a formal proof.

4.4 Nontermination

As we mentioned initially, our rules are for partial correctness, not for termination. Nontermination can easily arise because of recursive method invocations. Consider, for example, the nonterminating program:

$$[m = \zeta(x) x.m].m$$

Using our rules, we can prove that anything holds upon termination of this program, vacuously. Formally, we can derive the judgment:

$$\emptyset \Vdash [m = \zeta(x) x.m].m : A :: T$$

for any specification A and transition relation T without free variables.

We show the proof as a sequence of judgments, from the desired conclusion back to “true”, indicating between braces the rules applied and other justifications. We write *False* and *True* as abbreviations for the formulas ($false = true$) and ($false = false$), respectively.

$$\begin{aligned}
& \emptyset \Vdash [m = \zeta(x) x.m].m : A :: T \\
\Leftarrow & \{ \text{subsumption, using the assumptions and } \Vdash_{fol} \text{False} \Rightarrow T \} \\
& \emptyset \Vdash [m = \zeta(x) x.m].m : A :: \text{False} \\
\equiv & \{ \text{shorthand} \} \\
& \emptyset \Vdash \text{let } x = [m = \zeta(x) x.m] \text{ in } x.m : A :: \text{False} \\
\Leftarrow & \{ \text{let} \} \\
& \emptyset \Vdash [m = \zeta(x) x.m] : [m : \zeta(x) A :: \text{False}] :: \text{True} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x.m : A :: \text{False} \\
& \emptyset \Vdash A \\
& E \Vdash \text{False is a transition relation} \\
& \Vdash_{fol} \text{True} \wedge \text{False} \Rightarrow \text{False} \\
\Leftarrow & \{ \text{assumptions and first-order logic} \} \\
& \emptyset \Vdash [m = \zeta(x) x.m] : [m : \zeta(x) A :: \text{False}] :: \text{True} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x.m : A :: \text{False} \\
\Leftarrow & \{ \text{subsumption; object construction} \} \\
& \emptyset \Vdash \diamond \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x.m : A :: \text{False} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x.m : A :: \text{False} \\
\Leftarrow & \{ \text{well-formed specification environments; simplification} \} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x.m : A :: \text{False} \\
\Leftarrow & \{ \text{method invocation} \} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash x : [m : \zeta(x) A :: \text{False}] :: \text{Res}(x) \\
\Leftarrow & \{ \text{variables} \} \\
& \emptyset, x : [m : \zeta(x) A :: \text{False}] \Vdash \diamond \\
\Leftarrow & \{ \text{well-formed specification environments} \} \\
& \emptyset \Vdash \diamond \\
& \Vdash A \\
& \text{False is a transition relation} \\
\Leftarrow & \{ \text{well-formed specification environments; assumptions} \} \\
& \text{true}
\end{aligned}$$

Note that in the step that uses the rule for object construction we derive

$$\emptyset \Vdash [m = \zeta(x) x.m] : [m: \zeta(x)A :: False] :: True$$

from

$$\emptyset, x: [m: \zeta(x)A :: False] \Vdash x.m : A :: False$$

The assumption $x: [m: \zeta(x)A :: False]$ then enables us to check that the method body $x.m$ implements the specification for m given in $[m: \zeta(x)A :: False]$. Thus, recursively, we use the specification of an object in verifying its implementation.

4.5 Putting It All Together

Let a_0 be the following object:

$$[f = false, m = \zeta(s) \text{ if } s.f \text{ then } false \text{ else } true]$$

Below we give a detailed proof that this object satisfies the following specification:

$$[f: Bool, \\ m: \zeta(s) Bool :: (\dot{\sigma}(s, f) = false \Rightarrow r = true) \wedge \\ (\dot{\sigma}(s, f) = true \Rightarrow r = false)]$$

The specification says that the object's m method returns the negation of the object's f field. Below we also give a detailed proof of the following judgment, illustrating the use of a_0 :

$$\emptyset \Vdash \text{let } x = a_0 \text{ in } (x.f := true ; x.m) : Bool :: r = false$$

The proofs include applications of all our verification rules.

For convenience in the proofs, we define a transition relation T , as follows:

$$T \triangleq (\dot{\sigma}(s, f) = false \Rightarrow r = true) \wedge (\dot{\sigma}(s, f) = true \Rightarrow r = false)$$

We also define A_0 to be the specification of a_0 :

$$A_0 \triangleq [f: Bool, m: \zeta(s) Bool :: T]$$

We begin with a proof of the specification of a_0 , more precisely, with a proof of the judgment $E \Vdash a_0 : A_0 :: r = r$. We omit some easy steps, and use our standard abbreviations. For any well-formed environment E not containing s , we calculate:

$$\begin{aligned} & E \Vdash a_0 : A_0 :: r = r \\ \Leftarrow & \quad \{ \text{subsumption} \} \\ & E \Vdash a_0 : A_0 :: \\ & \quad \neg \text{alloc}(r) \wedge \text{alloc}(r) \wedge (\forall z. z \neq r \Rightarrow (\text{alloc}(z) \equiv \text{alloc}(z))) \wedge \\ & \quad \dot{\sigma}(r, f) = false \wedge (\forall z, w. z \neq r \Rightarrow \dot{\sigma}(z, w) = \dot{\sigma}(z, w)) \\ \Leftarrow & \quad \{ \text{object construction} \} \end{aligned}$$

$$\begin{aligned}
& E \Vdash \diamond \\
& E \Vdash \text{false} : \text{Bool} :: \text{Res}(\text{false}) \\
& E, s : A_0 \Vdash \text{if } s.f \text{ then } \text{false} \text{ else } \text{true} : \text{Bool} :: T \\
\Leftarrow & \quad \{ E \text{ is well-formed; constants} \} \\
& E, s : A_0 \Vdash \text{if } s.f \text{ then } \text{false} \text{ else } \text{true} : \text{Bool} :: T \\
\Leftarrow & \quad \{ \text{conditional} \} \\
& E, s : A_0 \Vdash s.f : \text{Bool} :: \text{Res}(\check{\sigma}(s, f)) \\
& E, s : A_0 \Vdash \text{false} : \text{Bool} :: \\
& \quad (\text{true} = \text{false} \Rightarrow r = \text{true}) \wedge (\text{true} = \text{true} \Rightarrow r = \text{false}) \\
& E, s : A_0 \Vdash \text{true} : \text{Bool} :: \\
& \quad (\text{false} = \text{false} \Rightarrow r = \text{true}) \wedge (\text{false} = \text{true} \Rightarrow r = \text{false}) \\
\Leftarrow & \quad \{ \text{field selection; subsumption, using } \neg(\text{false} = \text{true}) \} \\
& E, s : A_0 \Vdash s : A_0 :: \text{Res}(s) \\
& E, s : A_0 \Vdash \text{false} : \text{Bool} :: \text{Res}(\text{false}) \\
& E, s : A_0 \Vdash \text{true} : \text{Bool} :: \text{Res}(\text{true}) \\
\Leftarrow & \quad \{ \text{variables; constants} \} \\
& E, s : A_0 \Vdash \diamond \\
\Leftarrow & \quad \{ \text{well-formed specification environments} \} \\
& E \Vdash \diamond \\
& E \Vdash A_0 \\
& s \text{ not in } E \\
\Leftarrow & \quad \{ E \text{ is well-formed and does not contain } s; \text{ misc. rules} \} \\
& \text{true}
\end{aligned}$$

Next, we prove that the program

$$\text{let } x = a_0 \text{ in } (x.f := \text{true} ; x.m)$$

which contains a_0 , yields the result *false*:

$$\begin{aligned}
& \emptyset \Vdash \text{let } x = a_0 \text{ in } (x.f := \text{true} ; x.m) : \text{Bool} :: r = \text{false} \\
\Leftarrow & \quad \{ \text{let} \} \\
& \emptyset \Vdash a_0 : A_0 :: r = r \\
& \emptyset, x : A_0 \Vdash x.f := \text{true} ; x.m : \text{Bool} :: r = \text{false} \\
& \emptyset \Vdash \text{Bool} \\
& \emptyset \Vdash (r = \text{false}) \text{ is a transition relation} \\
& \Vdash_{\text{fol}} x = x \wedge r = \text{false} \Rightarrow r = \text{false} \\
\Leftarrow & \quad \{ \text{previous proof about } a_0; \text{ misc. rules} \} \\
& \emptyset, x : A_0 \Vdash x.f := \text{true} ; x.m : \text{Bool} :: r = \text{false} \\
\Leftarrow & \quad \{ \text{definition of } ; \} \\
& \emptyset, x : A_0 \Vdash \text{let } z = x.f := \text{true} \text{ in } x.m : \text{Bool} :: r = \text{false} \\
\Leftarrow & \quad \{ \text{let} \} \\
& \emptyset, x : A_0 \Vdash x.f := \text{true} : A_0 :: \check{\sigma}(x, f) = \text{true} \\
& \emptyset, x : A_0, z : A_0 \Vdash x.m : \text{Bool} :: \check{\sigma}(x, f) = \text{true} \Rightarrow r = \text{false} \\
& \emptyset, x : A_0 \Vdash \text{Bool} \\
& \emptyset, x : A_0 \Vdash (r = \text{false}) \text{ is a transition relation} \\
& \Vdash_{\text{fol}} \check{\sigma}(x, f) = \text{true} \wedge (\check{\sigma}(x, f) = \text{true} \Rightarrow r = \text{false}) \Rightarrow r = \text{false} \\
\Leftarrow & \quad \{ \text{misc. rules} \}
\end{aligned}$$

$$\begin{aligned}
& \emptyset, x: A_0 \Vdash x.f := true : A_0 :: \acute{o}(x, f) = true \\
& \emptyset, x: A_0, z: A_0 \Vdash x.m : Bool :: \grave{o}(x, f) = true \Rightarrow r = false \\
\Leftarrow & \quad \{ \text{subsumption; field update; method invocation} \} \\
& \emptyset, x: A_0 \Vdash x : A_0 :: Res(x) \\
& \emptyset, x: A_0, z: A_0 \Vdash x : A_0 :: Res(x) \\
\Leftarrow & \quad \{ \text{variables; misc. rules} \} \\
& true
\end{aligned}$$

5 Soundness and Related Properties

In this section we discuss the relation between verification and typing, obtaining two simple results. We then discuss the relation between verification and operational semantics, proving in particular a soundness theorem. The soundness theorem is the main technical result of this paper. Finally, we comment on completeness.

5.1 Typing versus Verification

Our first result establishes a correspondence between typing rules and verification rules: it says that only well-typed programs can be verified.

Proposition 1. *If $E \Vdash a : A :: T$ then $E' \vdash a : A'$ for some E' and A' (obtained from E and A by deleting transition relations).*

This result provides a first formal sanity check for the verification rules. It also highlights a limitation of the verification rules: for example, it implies that the verification rules do not enable us to derive that the program

$$if\ true\ then\ true\ else\ (true.f)$$

yields $r = true$, because this program is not well-typed. We do not view this limitation as a serious one because we are primarily interested in well-typed programs.

Conversely, every well-typed program can be verified, at least in a trivial sense:

Proposition 2. *If $E' \vdash a : A'$ then $E \Vdash a : A :: (r = r)$ for some E and A (obtained from E' and A' by inserting trivial transition relations).*

5.2 Soundness

We have both an axiomatic semantics (the verification rules) and an operational semantics. As we prove next, the two semantics agree in the sense that all that can be derived with the verification rules is true operationally. For example, if a program yields a result according to the operational semantics, and the axiomatic semantics says that the result is *true*, then indeed the result is *true*. We call this property soundness.

A special case of our soundness theorem is:

Theorem 1. *Assume that the operational semantics says that program b yields result v when run with an empty stack and an empty initial store (that is, $\emptyset, \emptyset \vdash b \rightsquigarrow v, \sigma'$ is provable with the rules of Section 2.1 for some σ'). If $\emptyset \Vdash b : \text{Bool} :: (r = \text{true})$ is provable then v is the boolean true. Similarly, if $\emptyset \Vdash b : \text{Bool} :: (r = \text{false})$ is provable then v is the boolean false.*

The statement of this theorem clearly expresses the consistency of the verification rules with the operational semantics. It is however unsatisfactory in at least two respects: (i) it does not apply to programs with free variables, to programs that return objects, or to programs that start running with a nonempty store; (ii) it cannot be proved directly anyway.

Going a little beyond the first special case, we can show that if b yields result v when run with an empty stack and an empty initial store, and if $\emptyset \Vdash b : A :: T$ is provable, then v “satisfies” A , and T holds when interpreted as a predicate on the initial and the final stores, with v as the value of r .

The full statement of our soundness theorem is as follows.

Theorem 2. *Assume that $\sigma, S \vdash a \rightsquigarrow v, \sigma'$ is provable and that $\Sigma \models \sigma$. If $E \Vdash a : A :: T$ is provable and $\Sigma \models S : E$, then $(S, \sigma, \sigma', v) \models T$ and there exists Σ' such that $\Sigma' \succeq \Sigma$, $\Sigma' \models \sigma'$, and $\Sigma' \models v : (A, S)$.*

The notations used in this statement are defined precisely in the appendix, where we also prove the theorem. Here we explain these notations only informally. The hypothesis $\Sigma \models \sigma$ means that the store σ meets the store specification Σ . The hypothesis $\Sigma \models S : E$ means that the variables in the stack S meet the specifications given in E in the context of Σ . The store specification Σ is needed because the initial store σ may be nonempty; similarly, the use of E as a specification for the stack S is needed because the program b may have free variables. The conclusion $(S, \sigma, \sigma', v) \models T$ means that the transition relation holds when we interpret its symbols using S, σ, σ' , and v (for example, taking v as the value of r). The conclusions $\Sigma' \succeq \Sigma$ and $\Sigma' \models \sigma'$ imply that Σ' agrees with Σ but possibly specifies additional objects in the store (those allocated in going from σ to σ'); these two conclusions appear in order to permit a direct inductive proof. Finally, the conclusion $\Sigma' \models v : (A, S)$ means that the output v meets the specification A in the context of Σ' and S .

Theorem 1 is a corollary of Theorem 2. As another corollary, we obtain a soundness theorem for the type system of Section 2.2. Therefore, as might be expected, our soundness proof is no less intricate than proofs of type soundness for imperative languages. In fact, Theorem 2 generalizes concepts developed for sophisticated proofs of type soundness [1, 10, 22, 36, 41]. New techniques are required because specifications, unlike ordinary (non-dependent) types, may contain occurrences of program variables.

5.3 Completeness Issues

While we have soundness, we do not have its converse, completeness. Unfortunately, our rules do not seem to be complete even for well-typed programs.

Careful examination of the following three similar programs reveals a first difficulty:

$$\begin{aligned}
b_1 &\triangleq \text{let } x = (\text{let } y = \text{true} \text{ in } [\text{m} = \zeta(z) y]) \text{ in } x.\text{m} \\
b_2 &\triangleq \text{let } y = \text{true} \text{ in } (\text{let } x = [\text{m} = \zeta(z) y] \text{ in } x.\text{m}) \\
b_3 &\triangleq \text{let } x = (\text{let } y = \text{true} \text{ in } [\text{f} = y, \text{m} = \zeta(z) z.\text{f}]) \text{ in } x.\text{m}
\end{aligned}$$

All three programs are well-typed and yield the result *true*. Using our rules, we can prove $\emptyset \Vdash b_2 : \text{Bool} :: (r = \text{true})$ and $\emptyset \Vdash b_3 : \text{Bool} :: (r = \text{true})$ but not $\emptyset \Vdash b_1 : \text{Bool} :: (r = \text{true})$. A reasonable diagnosis is that the judgment $E \Vdash a : A :: T$ does not allow sufficient interaction between A and T (particularly in the rule for *let*). One remedy is transforming b_1 into b_2 (by let-floating [28]) or into b_3 (by adding an auxiliary field). We have considered other remedies, but do not yet know which is the “right” one.

A deeper difficulty arises because the verification rules rely on a “global store” model. As Meyer and Sieber have explained [24], the use of this model is a source of incompleteness for procedural languages with local variables. Some of their remarks apply to our language as well. For example, the following program is reminiscent of their Example 2: *let* $x = [\text{f} = \text{true}] \text{ in } (y.\text{m} ; x.\text{f})$. This program will always return *true* because the method invocation $y.\text{m}$ cannot affect the field f of the newly allocated object x . We can prove this, but only by adopting a strong specification for y , for example requiring that $y.\text{m}$ not modify the field f of any object. Since the work of Meyer and Sieber, there has been considerable progress in the semantics of procedural languages with local variables (e.g., see [26, 29]). Some of the insights gained in that area should be applicable to reasoning about objects.

6 Related Work

As we mentioned in the introduction, there has been much research on specification and verification for object-oriented languages. The words “object” and “logic” are frequently used together in the literature, but with many different meanings (e.g., [33]). We do not know of any previous Hoare logic for a language like ours.

This section mentions several pieces of related work, emphasizing the most closely related ones that preceded our research, but also mentioning some relevant subsequent work. Since the initial presentation of our results [2], there has been much activity around formal techniques for programs in Java and related languages. In particular, this activity has resulted in semantics [13], in verification systems such as LOOP [38], in program-checking tools such as ESC/Java (which does not aim to be sound or complete) [8], and also in new Hoare logics [25, 39, 40]. A detailed survey of all that recent work is beyond the scope of the present paper. The interested reader may for example consult the programs of the workshops titled “Formal Techniques for Java-like Programs” and the papers presented there (e.g., [7]).

Our work is most similar to that of Leavens [17], who developed a Hoare logic for a small language with objects. The language is statically typed and includes a subtyping relation, but does not permit side-effects or aliasing. In another related study, de Boer [6] gave verification rules for the parallel language POOL. These rules apply to programs with side-effects and aliasing, but without subtyping or recursive methods, and with only one global implementation for each method (rather than one implementation per object). Both Leavens and de Boer obtained soundness results. More recently, in his dissertation [30], Poetzsch-Heffter considered how to integrate Larch-style interface specifications with Hoare-style verification techniques for object-oriented programs. Further, Poetzsch-Heffter and Müller [31] provided a proof system for a class-based language, representing properties of the type system of the language in axioms.

Much of the emphasis of the previous research has been on issues of refinement and inheritance. Lano and Haughton [15], Leavens [17, 18], and Liskov and Wing [23] all studied notions of subtyping and of refinement of specifications (similar to our subspecification relation, though in some respects more sophisticated). Stata and Guttag [34] studied the notion of subclassing, and presented a pre-formal approach for reasoning about inheritance. Utting [37] studied modular reasoning, allowing data refinement between the implementation and specification of an object. Lano and Haughton [16] have collected other research on object-oriented specification.

In some existing formalisms (e.g., Leavens's), specifications can be written in terms of abstract variables. Specifications at different levels of abstraction can be related by simulation relations or abstraction functions. Some results on abstraction appear in Leino's dissertation [19], which also gives a guarded-command semantics for objects and uses this semantics for reasoning about programs. Leino and Nelson [21] and Müller [25] have developed this approach to abstraction further.

Recursive types and recursive specifications can be helpful in dealing with programs that manipulate unbounded object data structures, which our logic treats only in a limited way. Two continuations of this work consider recursion [20, 32]. The most recent, due to Reus and Streicher, introduces a semantics of object specifications.

In another direction, the work of Hofmann and Tang shows how our logic can be embedded into a theorem prover for higher-order logic, and develops a verification-condition generator [12, 35]. It also treats examples beyond those shown in this paper.

Several other extensions of our work may be worthwhile. Some of those extensions appear straightforward; for example, it would be trivial to account for a construct that compares the addresses of two objects, or for a cloning construct. Rules for subclasses and inheritance would be important for treating standard class-based languages like Modula-3 and Java; perhaps one could develop a formal version of Stata's and Guttag's approach. The addition of concurrency primitives would be more difficult; it would call for a change of formalism, similar to the move from Hoare logic to Owicki-Gries logic [27].

7 Conclusions

In summary, the main outcome of our work is a logic that enables us (at least in principle) to specify and to verify object-oriented programs. To our knowledge, our notations and rules permit proofs that, despite their simplicity, are outside the scope of previous methods. However, our work is only a first step. It has already stimulated some further research, and we hope that it will continue to do so.

Secondarily, we hope that our logic will serve as another datapoint on the relations between types and specifications. In the realm of functional programming, specifications can be seen as a neat generalization of ordinary types (through notions such as dependent types, or in the context of abstract interpretations). In our experience with imperative object-oriented languages, the step from types to specifications is not straightforward; still, type theory is sometimes helpful, for example in suggesting techniques for soundness proofs.

Acknowledgments

Luca Cardelli and Greg Nelson helped in the initial stages of this work. Gary Leavens and Raymie Stata told us about related research. Krzysztof Apt, Frank de Boer, Luca Cardelli, Rowan Davies, and David Naumann made useful comments on drafts of this paper. This work was started at Digital Equipment Corporation's Systems Research Center. Martín Abadi's work is partly supported by the National Science Foundation under Grant CCR-0204162. Last but not least, Martín Abadi's work owes much to Zohar Manna, who taught him verification and many other things.

References

1. M. Abadi and L. Cardelli. *A Theory of Objects*. Springer-Verlag, 1996.
2. M. Abadi and K. R. M. Leino. A logic of object-oriented programs. In *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *Lecture Notes in Computer Science*, pages 682–696. Springer-Verlag, April 1997.
3. K. R. Apt. Ten years of Hoare's logic: A survey—Part I. *ACM Transactions on Programming Languages and Systems*, 3(4):431–483, October 1981.
4. J. Barwise. An introduction to first-order logic. In J. Barwise, editor, *The Handbook of Mathematical Logic*, Studies in Logic and Foundations of Mathematics, pages 5–46. North Holland, 1977.
5. E. M. Clarke. Programming language constructs for which it is impossible to obtain good Hoare axiom systems. *Journal of the ACM*, 26(1):129–147, January 1979.
6. F. S. de Boer. A proof system for the parallel object-oriented language POOL. In M. S. Paterson, editor, *Proceedings of the Seventeenth International Colloquium on Automata, Languages and Programming*, volume 443 of *Lecture Notes in Computer Science*, pages 572–585. Springer-Verlag, 1990.
7. Special issue: formal techniques for Java programs. *Concurrency and Computation: Practice and Experience*, 13(13), November 2001. Edited by S. Eisenbach and G. T. Leavens.

8. C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. *ACM SIGPLAN Notices*, 37(5):234–245, June 2002. Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation.
9. R. W. Floyd. Assigning meanings to programs. In *Proceedings of the Symposium on Applied Math., Vol. 19*, pages 19–32. American Mathematical Society, 1967.
10. R. Harper. A simplified account of polymorphic references. *Information Processing Letters*, 51:201–206, 1994.
11. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–583, October 1969.
12. M. Hofmann and F. Tang. Implementing a program logic of objects in a higher-order logic theorem prover. In *Theorem Proving in Higher Order Logics, 13th International Conference, TPHOLs 2000*, volume 1869 of *Lecture Notes in Computer Science*, pages 268–282. Springer-Verlag, 2000.
13. B. Jacobs and E. Poll. Coalgebras and monads in the semantics of Java. *Theoretical Computer Science*, 291(3):329–349, 2003.
14. C. B. Jones. An object-based design method for concurrent programs. Technical Report UMCS-92-12-1, University of Manchester, 1992.
15. K. Lano and H. Haughton. Reasoning and refinement in object-oriented specification languages. In O. L. Madsen, editor, *Proceedings of the 6th European Conference on Object-Oriented Programming (ECOOP)*, volume 615 of *Lecture Notes in Computer Science*, pages 78–97. Springer-Verlag, June 1992.
16. K. Lano and H. Haughton. *Object-Oriented Specification Case Studies*. Prentice Hall, New York, 1994.
17. G. T. 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.
18. G. T. Leavens. Modular specification and verification of object-oriented programs. *IEEE Software*, pages 72–80, July 1991.
19. K. R. M. Leino. *Toward Reliable Modular Programs*. PhD thesis, California Institute of Technology, 1995. Available as Technical Report Caltech-CS-TR-95-03.
20. K. R. M. Leino. Recursive object types in a logic of object-oriented programs. *Nordic Journal of Computing*, 5(4):330–360, 1998.
21. K. R. M. Leino and G. Nelson. Data abstraction and information hiding. *ACM Transactions on Programming Languages and Systems*, 24(5):491–553, September 2002.
22. X. Leroy. Polymorphic typing of an algorithmic language. Technical report, Institut National de Recherche en Informatique et en Automatique, October 1992. English version of the author’s PhD thesis.
23. B. H. Liskov and J. M. Wing. A behavioral notion of subtyping. *ACM Transactions on Programming Languages and Systems*, 16(6):1811–1841, November 1994.
24. A. R. Meyer and K. Sieber. Towards fully abstract semantics for local variables: Preliminary report. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 191–203, January 1988.
25. P. Müller. *Modular Specification and Verification of Object-Oriented Programs*, volume 2262 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002. PhD thesis, FernUniversität Hagen.
26. P. W. O’Hearn and R. D. Tennent. Parametricity and local variables. *Journal of the ACM*, 42(3):658–709, May 1995.
27. S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 6(4):319–340, 1976.

28. S. Peyton Jones, W. Partain, and A. Santos. Let-floating: moving bindings to give faster programs. In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming (ICFP '96)*, pages 1–12, May 1996.
29. A. M. Pitts and I. D. B. Stark. Observable properties of higher order functions that dynamically create local names, or: What's new? In *Mathematical Foundations of Computer Science, Proc. 18th Int. Symp., Gdańsk, 1993*, volume 711 of *Lecture Notes in Computer Science*, pages 122–141. Springer-Verlag, 1993.
30. A. Poetzsch-Heffter. Specification and verification of object-oriented programs. Habilitationsschrift, Technische Universität München, 1997. Available at www.technik.tu-muenchen.de/persons/poetzsch/habil.ps.gz.
31. A. Poetzsch-Heffter and P. Müller. Logical foundations for typed object-oriented languages. In D. Gries and W.-P. de Roeper, editors, *Programming Concepts and Methods (PROCOMET'98)*, IFIP, pages 404–423. Chapman & Hall, June 1998.
32. B. Reus and T. Streicher. Semantics and logic of object calculi. In *Proceedings of the Seventeenth Annual IEEE Symposium on Logic in Computer Science*, pages 113–122, July 2002.
33. A. Sernadas, C. Sernadas, and J. F. Costa. Object specification logic. *Journal of Logic and Computation*, 5(5):603–630, 1995.
34. R. Stata and J. V. Guttag. Modular reasoning in the presence of subclassing. *ACM SIGPLAN Notices*, 30(10):200–214, October 1995. OOPSLA '95 conference proceedings.
35. F. H.-L. Tang. *Towards feasible, machine-assisted verification of object-oriented programs*. PhD thesis, University of Edinburgh, 2002. Available at www.dcs.ed.ac.uk/home/fhlt/docs/fhlt-thesis.pdf.
36. M. Tofte. Type inference for polymorphic references. *Information and Computation*, 89(1):1–34, November 1990.
37. M. Utting. *An Object-Oriented Refinement Calculus with Modular Reasoning*. PhD thesis, University of New South Wales, 1992. Available at www.cs.waikato.ac.nz/~marku/phd.html.
38. J. van den Berg and B. Jacobs. The LOOP compiler for Java and JML. In *Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001, Proceedings*, volume 2031, pages 299–313, 2001.
39. D. von Oheimb. Hoare logic for Java in Isabelle/HOL. *Concurrency and Computation: Practice and Experience*, 13(13):1173–1214, November 2001.
40. D. von Oheimb and T. Nipkow. Hoare logic for NanoJava: Auxiliary variables, side effects and virtual methods revisited. In *Formal Methods – Getting IT Right*, volume 2391 of *Lecture Notes in Computer Science*, pages 89–105. Springer-Verlag, July 2002.
41. A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, November 1994.
42. A. Yonezawa and M. Tokoro, editors. *Object-oriented Concurrent Programming*. MIT Press, 1987.

Appendix: Soundness (Definitions and Theorems)

The soundness theorem requires several auxiliary notions. Before giving their definitions, we motivate each of them informally.

As part of the soundness theorem, we show that if the verification rules say that a satisfies the specification A and if, operationally, a yields the result v , then v satisfies A . Making precise the assertion that v satisfies A is delicate for several reasons.

One problem is that program variables may occur free in A . We surmount this difficulty by considering specification closures. A specification closure is a specification paired with a stack that gives values for the free program variables of the specification. (See Definition 1.) Instead of saying that v satisfies A , we can say that v satisfies the specification closure (A, S) , where S is the stack used for the execution of a .

Definition 1 (Specification closures).

- *A specification closure is a pair (A, S) where A is a specification and S is a stack such that the free program variables of A are all in the domain of S .*
- *Given a specification A and a stack S , we write AS for the result of replacing the free program variables of A by the corresponding results from S . Similarly, we write TS when T is a formula.*
- *We handle AS and TS as formal expressions by treating all object names as free variables. In particular, we may write the subtyping assertion $A'S' <: AS$; this assertion is defined by the standard rules for subtyping. (We omit $a \Vdash$ in $A'S' <: AS$ in order to stress that this is not a judgment of our verification system.)*

A second problem is that v may be an address in the store (that is, an object name) and the store may contain cycles; cycles impede inductive definitions. We surmount this difficulty by introducing store specifications, which associate specification closures with object names.

Definition 2 (Store specifications).

- *A store specification is a partial function that maps object names (from the set \mathcal{H}) to specification closures.*
- *Given store specifications Σ and Σ' , we write $\Sigma' \succeq \Sigma$ if Σ' extends Σ .*
- *Given a store specification Σ , a specification closure (A, S) , and a result v , $\Sigma \models_0 v : (A, S)$ holds if either A is *Bool* and v is one of *false* and *true*, or if (A, S) is $\Sigma(v)$ and $v \in \mathcal{H}$.*
- *$\Sigma \models v : (B, S)$ holds if there exist B' and S' such that $B'S' <: BS$ and $\Sigma \models_0 v : (B', S')$.*

The verification of a takes place with a particular specification environment, and the execution of a takes place with a particular stack. The soundness theorem assumes that the stack matches the environment, in the sense that if v_i is the value for x_i in the stack, and A_i is its specification in the environment, then v_i satisfies a suitable specification closure (A_i, S_i) .

Definition 3 (Stacks vs. environments). *The relation $\Sigma \models S : E$ is defined inductively by:*

- If Σ is a store specification, then $\Sigma \models \emptyset : \emptyset$.
- If $\Sigma \models S : E$, $\Sigma \models v : (A, S)$, and x is not in E , then $\Sigma \models S.(x \mapsto v) : (E, x : A)$.

For a store specification to be useful, it needs to be consistent with the particular store under consideration. We define this consistency relation without breaking cycles in the store, as follows.

Definition 4 (Stores vs. store specifications). *Given a store σ and a store specification Σ , $\Sigma \models \sigma$ holds if Σ and σ have the same domain, and for every v in their domain, $\Sigma(v)$ has the form (A, S) where A is*

$$[f_i : A_i \quad i \in 1..n, \quad m_j : \zeta(y_j)B_j :: T_j \quad j \in 1..m]$$

and

- for $i \in 1..n$, $\sigma(v)(f_i)$ is defined and $\Sigma \models \sigma(v)(f_i) : (A_i, S)$;
- for $j \in 1..m$, $\sigma(v)(m_j)$ is of the form $\langle \zeta(y_j)b_j, S \rangle$, and $E, y_j : A \Vdash b_j : B_j :: T_j$ for some E such that $\Sigma \models S : E$.

As part of the soundness theorem, we show also that if the verification rules say that a satisfies the transition relation T and if the execution of a with the initial store σ yields the store σ' , then T is true when interpreted as a predicate on the stores σ and σ' , with v as the value of r . Making precise the assertion that T is true is not too difficult, since T is simply a formula in first-order logic.

In order to interpret a first-order-logic formula, all one does is give a domain (a nonempty set), associate relations on this domain with predicate symbols, associate operations on this domain with function symbols, and map variables to elements of the domain [4]. Collectively, the domain, the relations, and the operations are called a structure; the mapping of variables to elements is called an assignment. In the case of T , we define the structure from the stores σ and σ' ; the assignment maps r to v , and maps any other free variables of T to their values in the stack.

Definition 5 (Satisfaction for formulas). *Given two stores σ and σ' , we define a structure, as follows:*

- The domain of the structure is $\{\text{false}, \text{true}\} \cup \mathcal{H} \cup \mathcal{F} \cup \{\perp\}$ (so it includes booleans, object names, field names, and a special, distinct undefined value).
- false , true , and all field names are interpreted as themselves.
- δ is interpreted as the binary function that maps any d_1 and d_2 to $\sigma(d_1)(d_2)$ if this is defined and to \perp otherwise.
- δ' is interpreted as the binary function that maps any d_1 and d_2 to $\sigma'(d_1)(d_2)$ if this is defined and to \perp otherwise.
- alloc is interpreted as the unary predicate that maps any d to true if and only if d is in the domain of σ .
- alloc' is interpreted as the unary predicate that maps any d to true if and only if d is in the domain of σ' .

Given a stack S and a result v , we define an assignment of elements of the domain to variables, as follows:

- r is interpreted as v .
- Any other variable x is interpreted as $S(x)$ if this is defined and as \perp otherwise.

Given a transition relation T , we write $(S, \sigma, \sigma', v) \models T$ if the assignment associated with S and v satisfies T in the structure associated with σ and σ' .

Two simple lemmas, given next without proof, state some of the properties of the notions defined above.

Lemma 1 (Substitution). *If A and B are specifications such that $\Vdash A <: B$, and S is a stack, then $AS <: BS$.*

Lemma 2 (Extension). *If $\Sigma \models S : E$ and $\Sigma' \succeq \Sigma$, then $\Sigma' \models S : E$.*

Theorem 1, given in the main body of this paper, is a special case of Theorem 2, which we restate and prove here.

Theorem 2. *Assume that $\sigma, S \vdash a \rightsquigarrow v, \sigma'$ is provable and that $\Sigma \models \sigma$. If $E \Vdash a : A :: T$ is provable and $\Sigma \models S : E$, then $(S, \sigma, \sigma', v) \models T$ and there exists Σ' such that $\Sigma' \succeq \Sigma$, $\Sigma' \models \sigma'$, and $\Sigma' \models v : (A, S)$.*

Proof. The proof is a direct induction on the derivation of $\sigma, S \vdash a \rightsquigarrow v, \sigma'$. There is a case for each of the rules of the operational semantics.

The consideration of the subsumption rule basically complicates the arguments for all the cases. A different, more complex induction can avoid this complication; see [20]. In any case, unlike in some proofs for Hoare logic [3], the possibility of recursion does not lead to the consideration of approximations to recursive procedures.

Variables

$$\frac{S(x) = v}{\sigma, S \vdash x \rightsquigarrow v, \sigma}$$

Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash x : A :: T$. Since $E \Vdash x : A :: T$, the environment E must contain $x:A'$ for some A' such that $\Vdash A' <: A$, and $\Vdash_{fol} Res(x) \Rightarrow T$. We obtain:

- $(S, \sigma, \sigma, v) \models T$, since $\Vdash_{fol} Res(x) \Rightarrow T$ and $(S, \sigma, \sigma, v) \models Res(x)$ (because $S(x) = v$).
- $\Sigma \succeq \Sigma$, trivially.
- $\Sigma \models \sigma$, by hypothesis.
- $\Sigma \models v : (A, S)$: Since $\Sigma \models S : E$, we must have $\Sigma \models v : (A', S')$ for some prefix S' of S , and hence $\Sigma \models_0 v : (A'', S'')$ for some (A'', S'') such that $A''S'' <: A'S'$. Since $\Vdash A' <: A$, the substitution lemma yields $A'S <: AS$. By transitivity, we obtain $A''S'' <: AS$, since $A'S' = A'S$. Therefore, $\Sigma \models v : (A, S)$.

Constants

$$\frac{}{\sigma, S \vdash \text{false} \rightsquigarrow \text{false}, \sigma} \quad \frac{}{\sigma, S \vdash \text{true} \rightsquigarrow \text{true}, \sigma}$$

We argue the case for *false*. Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash \text{false} : A :: T$. Since $E \Vdash \text{false} : A :: T$, we must have $A = \text{Bool}$ and $\Vdash_{\text{fol}} \text{Res}(\text{false}) \Rightarrow T$. We obtain:

- $(S, \sigma, \sigma, \text{false}) \models T$ since $\Vdash_{\text{fol}} \text{Res}(\text{false}) \Rightarrow T$ and $(S, \sigma, \sigma, \text{false}) \models \text{Res}(\text{false})$.
- $\Sigma \succeq \Sigma$, trivially.
- $\Sigma \models \sigma$, by hypothesis.
- $\Sigma \models \text{false} : (\text{Bool}, S)$, by the definitions.

Conditional

$$\frac{S(x) = \text{false} \quad \sigma, S \vdash a_1 \rightsquigarrow v, \sigma'}{\sigma, S \vdash \text{if } x \text{ then } a_0 \text{ else } a_1 \rightsquigarrow v, \sigma'}$$

$$\frac{S(x) = \text{true} \quad \sigma, S \vdash a_0 \rightsquigarrow v, \sigma'}{\sigma, S \vdash \text{if } x \text{ then } a_0 \text{ else } a_1 \rightsquigarrow v, \sigma'}$$

We argue the case for *false*. Let a be *if* x *then* a_0 *else* a_1 . Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash a : A :: T$. Since $E \Vdash a : A :: T$, we must have $E \Vdash x : \text{Bool} :: \text{Res}(x)$ and there must exist A_1, A'_1, T_1 , and T'_1 such that $E \Vdash a_1 : A_1 :: T_1$, $A_1[\text{false}/x] = A'_1[\text{false}/x]$, $T_1[\text{false}/x] = T'_1[\text{false}/x]$, $\Vdash A'_1 <: A$, and $\Vdash_{\text{fol}} T'_1 \Rightarrow T$. By induction hypothesis, $(S, \sigma, \sigma', v) \models T_1$ and there exists Σ' such that $\Sigma' \succeq \Sigma$, $\Sigma' \models \sigma'$, and $\Sigma' \models v : (A_1, S)$. We obtain:

- $(S, \sigma, \sigma', v) \models T$: Since $(S, \sigma, \sigma', v) \models T_1$ and $S(x) = \text{false}$, we have $(S, \sigma, \sigma', v) \models T_1[\text{false}/x]$. But $T_1[\text{false}/x] = T'_1[\text{false}/x]$, so we have $(S, \sigma, \sigma', v) \models T'_1[\text{false}/x]$. Since $S(x) = \text{false}$, we have $(S, \sigma, \sigma', v) \models T'_1$. Finally, $\Vdash_{\text{fol}} T'_1 \Rightarrow T$ yields $(S, \sigma, \sigma', v) \models T$.
- $\Sigma' \succeq \Sigma$.
- $\Sigma' \models \sigma'$.
- $\Sigma' \models v : (A, S)$: Since $\Sigma' \models v : (A_1, S)$, there exist A' and S' such that $\Sigma' \models_0 v : (A', S')$ and $A'S' <: A_1S$. Since $S(x) = \text{false}$ and $A_1[\text{false}/x] = A'_1[\text{false}/x]$, we obtain $A_1S = A'_1S$. Since $\Vdash A'_1 <: A$, the substitution lemma yields $A'_1S <: AS$. We obtain $A'S' <: A_1S = A'_1S <: AS$, and hence $\Sigma' \models v : (A, S)$.

Let

$$\frac{\sigma, S \vdash a \rightsquigarrow v, \sigma' \quad \sigma', S.(x \mapsto v) \vdash b \rightsquigarrow v', \sigma''}{\sigma, S \vdash \text{let } x = a \text{ in } b \rightsquigarrow v', \sigma''}$$

Let c be *let* $x = a$ *in* b . Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash c : B :: T$. Since $E \Vdash c : B :: T$, we must have $E \Vdash a : A :: R$ and $E, x:A \Vdash b : B' :: T'$ for some R, B' , and T' such that $\Vdash B' <: B$ and $\Vdash_{\text{fol}} R[\check{\sigma}/\check{\sigma}, \text{alloc}/\text{alloc}, x/r] \wedge T'[\check{\sigma}/\check{\sigma}, \text{alloc}/\text{alloc}] \Rightarrow T$; in addition, $E \Vdash B'$ and $E \Vdash T$ is a transition relation, so x does not occur free in either B' or T , and $\check{\sigma}$ and alloc do not occur in T .

By induction hypothesis, $(S, \sigma, \sigma', v) \models R$ and there exists Σ' such that $\Sigma' \succeq \Sigma$, $\Sigma' \models \sigma'$, and $\Sigma' \models v : (A, S)$. Since $\Sigma \models S : E$, the extension lemma yields $\Sigma' \models S : E$. Therefore, $\Sigma' \models S.(x \mapsto v) : (E, x : A)$. (Note that x cannot appear in E because $E, x : A \Vdash b : B' :: T$; hence x is not in the domain of S either.) By induction hypothesis, $(S.(x \mapsto v), \sigma', \sigma'', v') \models T'$ and there exists Σ'' such that $\Sigma'' \succeq \Sigma'$, $\Sigma'' \models \sigma''$, and $\Sigma'' \models v' : (B', S.(x \mapsto v))$. We obtain:

- $(S, \sigma, \sigma'', v') \models T$: Since $(S, \sigma, \sigma', v) \models R$ and $(S.(x \mapsto v), \sigma', \sigma'', v') \models T'$, the two models associated with $(S.(x \mapsto v), \sigma, \sigma', v')$ and $(S.(x \mapsto v), \sigma', \sigma'', v')$ can be extended to a model for $R[\check{\sigma}/\acute{\sigma}, \check{alloc}/\acute{alloc}, x/r] \wedge T'[\check{\sigma}/\acute{\sigma}, \check{alloc}/\acute{alloc}]$. (The interpretation of $\check{\sigma}$ and \check{alloc} in this model is easily determined from σ' .) Since $\Vdash_{fol} R[\check{\sigma}/\acute{\sigma}, \check{alloc}/\acute{alloc}, x/r] \wedge T'[\check{\sigma}/\acute{\sigma}, \check{alloc}/\acute{alloc}] \Rightarrow T$, this is also a model for T . Since $\check{\sigma}$ and \check{alloc} do not occur in T and x does not occur free in T , we conclude that $(S, \sigma, \sigma'', v') \models T$.
- $\Sigma'' \succeq \Sigma$, by transitivity.
- $\Sigma'' \models \sigma''$.
- $\Sigma'' \models v' : (B, S)$: Since $\Sigma'' \models v' : (B', S.(x \mapsto v))$, there exist B'' and S'' such that $\Sigma'' \models_0 v' : (B'', S'')$ and $B''S'' < B'(S.(x \mapsto v))$. Since x does not occur free in B' , we have $B'(S.(x \mapsto v)) = B'S$. Since $\Vdash B' < B$, the substitution lemma yields $B'S < BS$. Therefore, $B''S'' < BS$, so $\Sigma'' \models v' : (B, S)$.

Object construction

$$\frac{S(x_i) = v_i \quad i \in 1..n \quad h \notin \text{dom}(\sigma) \quad h \in \mathcal{H} \quad \sigma' = \sigma.(h \mapsto (f_i \mapsto v_i \quad i \in 1..n, m_j \mapsto \langle \varsigma(y_j)b_j, S \rangle \quad j \in 1..m))}{\sigma, S \vdash [f_i = x_i \quad i \in 1..n, m_j = \varsigma(y_j)b_j \quad j \in 1..m] \rightsquigarrow h, \sigma'}$$

Let c be $[f_i = x_i \quad i \in 1..n, m_j = \varsigma(y_j)b_j \quad j \in 1..m]$. Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash c : A :: T$. Since $E \Vdash c : A :: T$, there exists A' of the form $[f_i : A_i \quad i \in 1..n, m_j : \varsigma(y_j)B_j :: T_j \quad j \in 1..m]$ such that $\Vdash A' < A$, $E \Vdash x_i : A_i :: \text{Res}(x_i)$ for $i \in 1..n$, and $E, y_j : A' \Vdash b_j : B_j :: T_j$ for $j \in 1..m$. In addition, let T' be:

$$\begin{aligned} & \neg \check{alloc}(r) \wedge \acute{alloc}(r) \wedge \\ & (\forall z. z \neq r \Rightarrow (\check{alloc}(z) \equiv \acute{alloc}(z))) \wedge \\ & \acute{\sigma}(r, f_1) = x_1 \wedge \dots \wedge \acute{\sigma}(r, f_n) = x_n \wedge \\ & (\forall z, w. z \neq r \Rightarrow \check{\sigma}(z, w) = \acute{\sigma}(z, w)) \end{aligned}$$

It must be that $\Vdash_{fol} T' \Rightarrow T$. Let Σ' be $\Sigma.(h \mapsto (A', S))$. We obtain:

- $(S, \sigma, \sigma', h) \models T$, since $(S, \sigma, \sigma', h) \models T'$.
- $\Sigma' \succeq \Sigma$, since Σ and σ have the same domain and this domain does not include h .
- $\Sigma' \models \sigma'$: First, Σ' and σ' have the same domain, namely the domain of Σ and σ extended with h . Since $\Sigma \models \sigma$, we need to check conditions only for h ; these conditions are determined by $\Sigma'(h) = (A', S)$.

- For $i \in 1..n$, $\sigma'(h)(f_i)$ is defined, and equals v_i . Since $E \Vdash x : A_i :: \text{Res}(x_i)$, it must be that E contains $x_i : A'_i$ for some A'_i such that $\Vdash A'_i < A_i$. Since $\Sigma \models S : E$ and $S(x_i) = v_i$, we must have $\Sigma \models v_i : (A'_i, S'_i)$ for some prefix S'_i of S , and hence $\Sigma \models_0 v_i : (A''_i, S''_i)$ for some (A''_i, S''_i) such that $A''_i S''_i < A'_i S'_i$. Therefore, $\Sigma' \models_0 v_i : (A''_i, S''_i)$ (because $\Sigma' \succeq \Sigma$). Since $\Vdash A'_i < A_i$, the substitution lemma yields $A''_i S''_i < A_i S$. In addition $A''_i S''_i = A'_i S$. By transitivity, we obtain $A''_i S''_i < A_i S$. It follows that $\Sigma' \models v_i : (A_i, S)$.
 - For $j \in 1..m$, $\sigma'(v)(m_j)$ is of the form $\langle \varsigma(y_j) b_j, S \rangle$, and $E, y_j : A' \Vdash b_j : B_j :: T_j$. In addition, E is such that $\Sigma \models S : E$, and hence $\Sigma' \models S : E$ by the extension lemma.
- $\Sigma' \models h : (A, S)$, since $\Sigma'(h) = (A', S)$ and $A' S < A S$ (because $\Vdash A' < A$ and by the substitution lemma).

Field selection

$$\frac{S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(f) = v}{\sigma, S \vdash x.f \rightsquigarrow v, \sigma}$$

Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash x.f : A :: T$. Since $E \Vdash x.f : A :: T$, there exist B and A' such that B has the form $[\dots f : A' \dots]$, $\Vdash A' < A$, and the environment E contains $x : B$. In addition, $\Vdash_{\text{fol}} \text{Res}(\delta(x, f)) \Rightarrow T$. We obtain:

- $(S, \sigma, \sigma, v) \models T$, since $v = \sigma(S(x))(f)$ so $(S, \sigma, \sigma, v) \models \text{Res}(\delta(x, f))$.
- $\Sigma \succeq \Sigma$, trivially.
- $\Sigma \models \sigma$, trivially.
- $\Sigma \models v : (A, S)$: Since $\Sigma \models S : E$ and $S(x) = h$, we must have $\Sigma \models h : (B, S')$ for some prefix S' of S , and hence $\Sigma(h) = (B'', S'')$ for some (B'', S'') such that $B'' S'' < B S'$. Therefore, B'' has the form $[\dots f : A'' \dots]$ with $A'' S'' = A' S'$. Moreover, $A' S' = A' S$ since A' is a subexpression of B , (B, S') is a specification closure, and S' is a prefix of S . Since $\Sigma \models \sigma$, $\sigma(h)(f) = v$, and $\Sigma(h) = ([\dots f : A'' \dots], S'')$, we have $\Sigma \models v : (A'', S'')$, and hence $\Sigma \models_0 v : (A''', S''')$ for some (A''', S''') such that $A''' S''' < A'' S''$. Since $\Vdash A' < A$, the substitution lemma yields $A' S < A S$. By transitivity, we derive $A''' S''' < A S$. Therefore, $\Sigma \models v : (A, S)$.

Method invocation

$$\frac{S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(m) = \langle \varsigma(y) b, S' \rangle \quad \sigma, S'.(y \mapsto h) \vdash b \rightsquigarrow v, \sigma'}{\sigma, S \vdash x.m \rightsquigarrow v, \sigma'}$$

Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash x.m : A :: T$. Since $E \Vdash x.m : A :: T$, there exist B , A' , T' , and T^+ such that B has the form $[\dots m : \varsigma(y) A' :: T' \dots]$, $\Vdash A'[x/y] < A$, $\Vdash_{\text{fol}} T' \Rightarrow T^+$, and $\Vdash_{\text{fol}} T^+[x/y] \Rightarrow T$, and the environment E contains $x : B$. Since $\Sigma \models S : E$ and $S(x) = h$, we must have $\Sigma \models h : (B, S^+)$ for some prefix S^+ of S , and hence $\Sigma(h) = (B'', S'')$ for some (B'', S'') such that $B'' S'' < B S^+$. We assume (without loss of generality) that y is not in

the domain of S'' or S . Therefore, B'' has the form $[\dots m: \zeta(y)A'' :: T'' \dots]$ with $A''S'' <: A'S^+$ and $T''S'' \Rightarrow T'S^+$. Since $\Sigma \models \sigma$, $\sigma(h)(m)$ contains the stack S' , and $\Sigma(h) = ([\dots m: \zeta(y)A'' :: T'' \dots], S'')$, we obtain that $S'' = S'$. In addition, there exists E' such that $\Sigma \models S' : E'$ and $E', y: B'' \Vdash b : A'' :: T''$. Since $\Sigma(h) = (B'', S'')$ and $S'' = S'$, we obtain that $\Sigma \models S'.(y \mapsto h) : (E', y: B'')$. (Note that y cannot appear in E' because $E', y: B'' \Vdash b : A'' :: T''$.) By induction hypothesis, $(S'.(y \mapsto h), \sigma, \sigma', v) \models T''$ and there exists Σ' such that $\Sigma' \succeq \Sigma$, $\Sigma' \models \sigma'$, and $\Sigma' \models v : (A'', S'.(y \mapsto h))$, and hence there exists (C, U) such that $\Sigma' \models_0 v : (C, U)$ and $CU <: A''(S'.(y \mapsto h))$. We obtain:

- $(S, \sigma, \sigma', v) \models T$: We have $(S'.(y \mapsto h), \sigma, \sigma', v) \models T''$, so $(S''.(y \mapsto h), \sigma, \sigma', v) \models T''$ since $S'' = S'$. From $T''S'' \Rightarrow T'S^+$ we deduce $(S^+. (y \mapsto h), \sigma, \sigma', v) \models T'$. The free variables of T' other than y are included in the domain of S^+ , since (B, S^+) is a specification closure and $\zeta(y)A' :: T'$ is a subexpression of B . In addition, S^+ is a prefix of S . Therefore, $(S^+. (y \mapsto h), \sigma, \sigma', v) \models T'$ implies $(S.(y \mapsto h), \sigma, \sigma', v) \models T'$. We obtain $(S.(y \mapsto h), \sigma, \sigma', v) \models T^+$ since $\Vdash_{fol} T' \Rightarrow T^+$, so $(S, \sigma, \sigma', v) \models T^+[x/y]$ by substitution and since $S(x) = h$, so $(S, \sigma, \sigma', v) \models T$ since $\Vdash_{fol} T^+[x/y] \Rightarrow T$.
- $\Sigma' \succeq \Sigma$.
- $\Sigma' \models \sigma'$.
- $\Sigma' \models v : (A, S)$: We already have $\Sigma' \models_0 v : (C, U)$. In addition, $CU <: A''(S'.(y \mapsto h)) = A''(S''.(y \mapsto h)) <: A'(S^+. (y \mapsto h)) = A'(S.(y \mapsto h)) = A'[x/y]S <: AS$. (The equality $A''(S'.(y \mapsto h)) = A''(S''.(y \mapsto h))$ follows from $S'' = S'$. The relation $A''(S''.(y \mapsto h)) <: A'(S^+. (y \mapsto h))$ is obtained from $A''S'' <: A'S^+$ by substitution. The equality $A'(S^+. (y \mapsto h)) = A'(S.(y \mapsto h))$ holds because $\zeta(y)A' :: T'$ is a subexpression of B , (B, S^+) is a specification closure, and S^+ is a prefix of S . The equality $A'(S.(y \mapsto h)) = A'[x/y]S$ holds because $S(x) = h$. The relation $A'[x/y]S <: AS$ follows from $\Vdash A'[x/y] <: A$ by the substitution lemma.)

Field update

$$\frac{\begin{array}{l} S(x) = h \quad h \in \mathcal{H} \quad \sigma(h)(f) \text{ is defined} \\ S(y) = v \quad \sigma' = \sigma.(h \mapsto \sigma(h).(f \mapsto v)) \end{array}}{\sigma, S \vdash x.f := y \rightsquigarrow h, \sigma'}$$

Suppose $\Sigma \models \sigma$, $\Sigma \models S : E$, and $E \Vdash x.f := y : A :: T$. Since $E \Vdash x.f := y : A :: T$, there exist A' and C such that A' has the form $[\dots f: C \dots]$, $\Vdash A' <: A$, and the environment E contains $x: A'$; and there exists C' such that $\Vdash C' <: C$ and the environment E contains $y: C'$. In addition, let T' be:

$$\begin{array}{l} r = x \wedge \acute{\sigma}(x, f_k) = y \wedge \\ (\forall z, w. \neg(z = x \wedge w = f_k) \Rightarrow \grave{\sigma}(z, w) = \acute{\sigma}(z, w)) \wedge \\ (\forall z. \mathit{alloc}(z) \equiv \mathit{alloc}(z)) \end{array}$$

It must be that $\Vdash_{fol} T' \Rightarrow T$. We obtain:

- $(S, \sigma, \sigma', h) \models T$ since $(S, \sigma, \sigma', h) \models T'$.

- $\Sigma \succeq \Sigma$, trivially.
- $\Sigma \models \sigma'$: Since $\Sigma \models \sigma$, we need to check conditions only for $\sigma'(h)(f)$. Since $\Sigma \models S : E$ and $S(x) = h$, we must have $\Sigma \models h : (A', S')$ for some prefix S' of S , and hence $\Sigma(h) = (A'', S'')$ for some (A'', S'') such that $A''S'' <: A'S'$. Therefore, A'' has the form $[\dots f: C'' \dots]$ with $C''S'' = CS'$. Since $\Sigma \models S : E$ and $S(y) = v$, we must have $\Sigma \models v : (C', U)$ for some prefix U of S , and hence $\Sigma \models_0 v : (B', U')$ for some (B', U') such that $B'U' <: C'U$. Because $\Sigma(h)$ has the form $([\dots f: C'' \dots], S'')$, we need that $\Sigma \models v : (C'', S'')$. We obtain this from $\Sigma \models_0 v : (B', U')$ and $B'U' <: C'U = C'S <: CS = CS' = C''S''$. (The relation $C'S <: CS$ follows from $\Vdash C' <: C$ by the substitution lemma. The equality $CS = CS'$ holds because C is a subexpression of A' , (A', S') is a specification closure, and S' is a prefix of S .)
- $\Sigma \models h : (A, S)$: We have $A''S'' <: A'S'$; moreover, $A'S' = A'S$ since (A', S') is a specification closure and S' is a prefix of S . Since $\Vdash A' <: A$, the substitution lemma yields $A'S <: AS$. By transitivity, we derive $A''S'' <: AS$. Therefore, $\Sigma(h) = (A'', S'')$ yields $\Sigma \models h : (A, S)$.