

Environmental Acquisition Revisited

Richard Cobbe

Matthias Felleisen

PLT at the College of Computer and Information Science
Northeastern University
Boston, MA 02115
cobbe@ccs.neu.edu

September 5, 2006

Northeastern University Technical Report NU-CCIS-06-10

Abstract

Object-oriented programs frequently deal with hierarchical forms of data, often called whole-part arrangements. When in such arrangements the methods in the parts must compute information about their container, the programmer is currently forced to manually maintain a two-way relationship, also known as friend pointers. The maintenance problem becomes particularly complex in the case of deeply nested hierarchies. To eliminate this drudgery of pointer maintenance, Gil and Lorenz proposed a language construct, dubbed environmental acquisition, in 1996. The construct has been adapted and studied in several different contexts over the past decade.

In this paper, we analyze Gil and Lorenz’s language design with a formal model. Specifically, we extend CLASSICJAVA with environmental acquisition and determine what it takes to prove the revised type system sound. Then we study a number of design decisions in this framework and illustrate when a formal model can help (or not) make such decisions. We also discuss how this model helps us understand what a full-scale implementation of type-sound environmental acquisition would require.

1 From Patterns to Language Constructs

Object-oriented programs often manipulate data in the form of nested containers, which of course are equivalent to trees. For example, a program with a graphical user interface (GUI) tends to nest panels within panels inside a window to arrange the atomic widgets. Similarly, typical XML libraries create objects in the shape of trees to represent XML forms. In these contexts, it is common that methods compute information about the nesting of containers, and programming such a method is straightforward. If, however, the method in a containee of such a container hierarchy must compute data about its containers, programs become complicated. Specifically, the programs must establish and maintain knowledge about the container hierarchies at all times.

The problem of managing such hierarchies is the subject of several object-oriented design patterns [8]. In the particular case of nested hierarchies, a programmer should choose a composite pattern and maintain so-called “friend pointers.” Those connect an item to its container and are therefore known as “two-way links” in the UML-oriented pattern community [6, p. 155]. Even though patterns help solve this problem, programmers must still maintain complex relationships among objects manually. For this reason, several people [10, 14] have proposed the idea that patterns suggest new constructs for programming languages. Then the compiler and the run-time system can enforce the pattern’s programming invariants or, alternatively, can signal violations as soon as they recognize them.

Gil and Lorenz [9] recognized the problem of *environmental acquisition*, i.e., the need for methods in containees to compute information about their context, and proposed constructs for specifying whole-part

```

class JComponent extends Container /* ... */ {
    // ...
    public JRootPane getRootPane() {
        return SwingUtilities.getRootPane(this);
    }
    // ...
}

class SwingUtilities /* ... */ {
    // ...
    public static JRootPane getRootPane(Component c) {
        if (c instanceof RootPaneContainer) {
            return ((RootPaneContainer)c).getRootPane();
        }
        for( ; c != null; c = c.getParent()) {
            if (c instanceof JRootPane) {
                return (JRootPane)c;
            }
        }
        return null;
    }
    // ...
}

```

Figure 1: Root panes in Swing

relationships within programs.¹ Their proposal was informal; it neither specified a formal type system for environmental acquisition nor stated a soundness theorem for a specific model. Since then, environmental acquisition has found its way into Python [19, 7] as a part of the Zope project [15]. Also, the second author has used environmental acquisition for a systems administration project where computers, switches, and ethernet cards acquire policies and other network information from their containers [16].

In this paper—a revised and extended version of our POPL paper [2]—we resume and continue Gil and Lorenz’s language design experiment. Specifically, we present a statically-typed version of environmental acquisition for a class-based object-oriented programming language. Our goal is to show language designers how environmental acquisition helps programmers; what it takes to type constructs for environmental acquisition; and when the run-time system needs to check the invariants for acquisition relationships. A formal model, based on CLASSICJAVA [5], provides the framework in which we discuss the design decisions and in which we can prove a type soundness theorem. Finally, we use the model to formulate a challenge for the designers of type systems concerning the elimination of run-time checks.

The paper consists of seven sections. The next section introduces environmental acquisition with concrete examples. Then the third section briefly recalls those elements of CLASSICJAVA that matter for this paper. The fourth section is dedicated to JACQUES, an extension of CLASSICJAVA with constructs for specifying the environmental acquisition of fields and methods; we describe its type system and its formal semantics. In the fifth section we discuss several decisions concerning the design of JACQUES. The sixth section sketches how to scale our model to a full object-oriented language. The seventh section sketches related efforts, their limitations, and how we might incorporate their methods into future versions of JACQUES. The final section suggests some future work, especially on an improved type system for keeping track of acquisitions.

2 Motivating Examples

Environmental acquisition is inseparably tied to the whole-part relationships of an object containment hierarchy. In many cases, the purpose of methods is to propagate information toward the root of such an object tree. Acquisition becomes relevant when an application must propagate information *away* from the root.

¹As Gil and Lorenz point out, an object’s *environment* may incorporate many different kinds of related objects. We restrict ourselves, as they do, to the environment provided by an object’s containers.

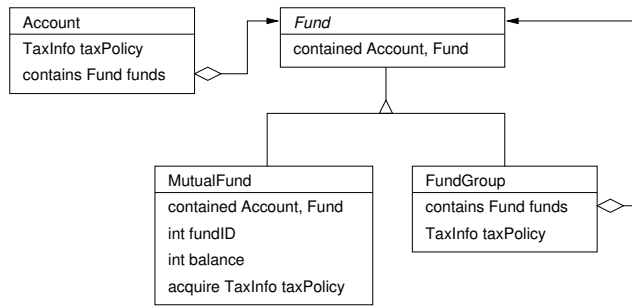


Figure 2: Acquisition example: investment accounts

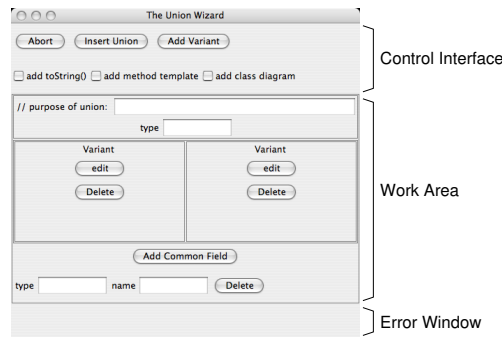


Figure 3: UnionInfo wizard dialog

Consider GUI frameworks, which are typically implemented with the composite pattern. In Java’s Swing toolkit [18], for instance, all graphical components are subclasses of `JComponent`; certain components, like `JPanel`, may contain other components. Every operating-system-level window in a Swing program has a single root pane; this pane contains the window’s menu bar and other capabilities. Therefore, components that need to access the menu bar must have access to their containing root pane. Currently, the `JComponent.getRootPane` method retrieves the root pane by chasing explicit pointers up the panel containment hierarchy, as shown in figure 1. As GUI components are added, destroyed, and moved, the library maintains these pointers. If Java supported environmental acquisition, each individual component could simply acquire a reference to the root pane from its container hierarchy, avoiding the need to write error-prone pointer-management code directly and instead leaving that responsibility to the language.

Following Gamma et al, we take our next example of acquisition from the financial sector. See figure 2 for the UML diagram. The `Account` class represents an individual’s account with a mutual fund firm; each account contains several `Fund` instances. It is frequently the case that an investor may have several types of funds within the same account; some funds may be standard investments, while others may be part of an IRA. To keep these funds separate for tax purposes, it is convenient to group multiple funds into a `FundGroup` instance. Now, transactions on a particular fund must take the fund’s tax policy into account. With environmental acquisition, we simply declare the policy for the account as a whole, allow fund groups to override that policy for their funds, and have each fund acquire the policy from its nearest container.

For a final example, extracted from DrScheme’s ProfessorJ component [12], consider a class-writing wizard for use in a program development environment. This wizard allows a student to define a new class or set of classes by specifying their properties graphically; the `produce` method in the `ClassUnionWizard` class generates the corresponding code. ProfessorJ allows the student to define classes in several different manners. For example, the `ClassInfo` wizard constructs a single class, and the `UnionInfo` wizard constructs a sum-of-products type.

As figure 3 shows, each dialog contains three main visual components: a control interface at the top, an

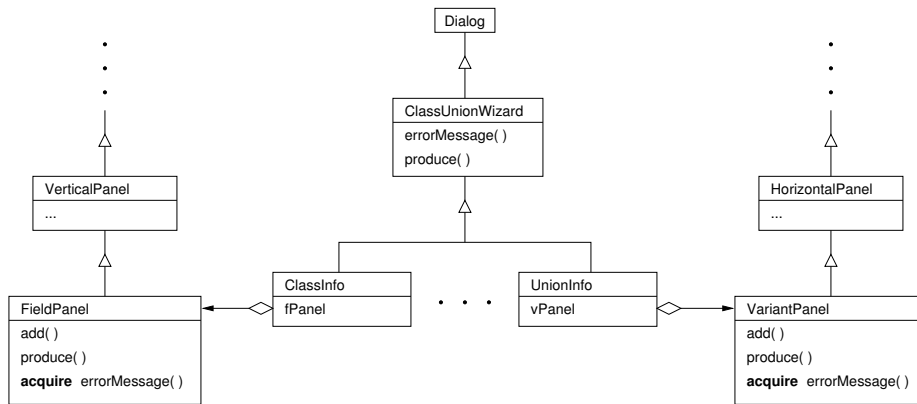


Figure 4: Acquisition example: class wizard

error window at the bottom, and a large work area in the center that is specialized to the specific wizard. For example, the `UnionInfo` wizard, contains a work area with one panel for each variant of the data type being developed. Since these variant panels are complicated in their own right, we have not included this code in `UnionInfo` but removed it to the `VariantPanel` class, instances of which are contained in the main dialog box. Figure 4 displays a UML diagram that summarizes the design.

A wizard typically cannot detect all of the errors in a design until it attempts to emit code via the `produce` method. In the `UnionInfo` class, this method simply invokes the same method on each of its `VariantPanel`s. If one of those detects an error, it must report it by communicating the error message up to the main dialog box for display in its error panel. Since, however, `VariantPanel` inherits from `HorizontalPanel` and not from `ClassUnionWizard`, it cannot use inheritance to access the `errorMessage` method. Instead, it must *acquire* this method from its containers.

This acquisition link also makes it easier to extend `UnionInfo`'s behavior later. In particular, if we wanted to subclass this class to direct all error reports to a log file on disk, in addition to sending them to the window, we could simply place the `VariantPanel` objects into the extended `UnionInfo`; the panels would acquire the extended `errorMessage` method, and we would achieve the desired behavior with no changes to `VariantPanel`.

Without environmental acquisition, we could achieve the same behavior through either of two strategies. First (and currently), we could wrap the `errorMessage` method in a Scheme closure and pass this as an argument to the `VariantPanel` constructor; this panel object could then simply invoke the closure to report an error. Second, we could manage the container pointer manually, and the `VariantPanel` object could chase pointers to find the `errorMessage` method. Unfortunately, both strategies involve significant bookkeeping. It is the purpose of environmental acquisition to eliminate such mundane tasks.

3 Classic Java

CLASSICJAVA [5] is a formal model of the sequential core of Java version 1.4 [11] that includes inheritance, method overriding, and field mutation; the model could easily apply to similar OO languages, like C# [3] or Eiffel [17]. We summarize CLASSICJAVA in this section and refer the reader to the original paper for the full details.

Figure 5 specifies the syntax for CLASSICJAVA programs. The underlined portions of the syntax are added by the type checking and elaboration phase in order to make certain context-sensitive details available to the operational semantics. In particular, evaluating a `super` expression requires knowledge of the location of the expression in the original program in order to perform method lookup correctly. Similarly, field accesses and assignments must know the static type of the object whose field is being accessed in order to resolve shadowed fields correctly.

Figure 6 describes the six type judgments necessary to define CLASSICJAVA's static semantics; we refer

P	=	$defn^* e$
$defn$	=	$class\ c\ extends\ c\ \{field^*\ meth^*\}$
$field$	=	$t\ fd$
$meth$	=	$t\ md(arg^*)\ \{e\}$
arg	=	$t\ var$
e	=	$new\ c\ var\ null\ e : c.fld\ e : c.fld = e$ $ e.md(e^*)\ super \equiv this : c.md(e^*)$ $ cast\ t\ e\ let\ var = e\ in\ e$
var	=	a variable name or <code>this</code>
c	=	a class name or <code>Object</code>
fd	=	a field name
md	=	a method name
t	=	c (a type name)

Figure 5: CLASSICJAVA abstract syntax

$\vdash_p P \Rightarrow P' : t$	program P elaborates to P' with type t
$P \vdash_d defn \Rightarrow defn'$	class definition $defn$ elaborates to $defn'$
$P, t \vdash_m meth \Rightarrow meth'$	method $meth$ in class t elaborates to $meth'$
$P, \Gamma \vdash_e e \Rightarrow e' : t$	e elaborates to e' with exact type t in Γ
$P, \Gamma \vdash_s e \Rightarrow e' : t$	e elaborates to e' and has type t using subsumption in Γ
$P \vdash_t t$	type t exists

Figure 6: CLASSICJAVA type judgments

the reader to the original paper for the formal definition of these judgments. The CLASSICJAVA semantics requires several additional relations; those that are relevant to our work are defined in figure 7. The dynamic semantics, a small-step operational semantics, defines reductions of the form $P \vdash \langle e, \mathcal{S} \rangle \hookrightarrow \langle e', \mathcal{S}' \rangle$, where P denotes the program's class definitions and \mathcal{S} and \mathcal{S}' are stores. Finally, we represent an instance as an ordered pair $\langle c, \mathcal{F} \rangle$, where c is the object's class tag and \mathcal{F} is a finite map from tagged field names (i.e., $c.fld$) to values.

CLASSICJAVA satisfies a type soundness theorem:

Theorem 1 (1: Type Soundness for CLASSICJAVA) *If $\vdash_p P \Rightarrow P' : t$ and $P' = defn_1 \dots defn_n e$, then either*

- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle obj, \mathcal{S} \rangle$ and $\mathcal{S}(obj) = \langle t', \mathcal{F} \rangle$ and $t' \leq_P t$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle null, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle error: bad\ cast, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle error: dereferenced\ null, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \uparrow$.

4 Jacques

JACQUES is an extension of CLASSICJAVA that supports environmental acquisition of both fields and methods. In this section, we first describe JACQUES's syntax, type system, and operational semantics; then we state

\prec_P	immediate subclass $c \prec_P c' \iff \text{class } c \text{ extends } c' \{ \dots \} \text{ is in } P$
\leq_P	subclass \leq_P is defined to be the reflexive-transitive closure of \prec_P
\in_P	field is declared in a class $\langle c.f.d, t \rangle \in_P c \iff \text{class } c \dots \{ \dots t \text{ fd } \dots \} \text{ is in } P$
\in_P	method is declared in a class $\langle md, (t_1 \dots t_n \rightarrow t), (var_1 \dots var_n), e \rangle \in_P c \iff$ $\text{class } c \dots \{ \dots t \text{ md}(t_1 \text{ var}_1 \dots t_n \text{ var}_n) \{ e \} \dots \} \text{ is in } P$
\in_P	field is contained in a class $\langle c'.fd, t \rangle \in_P c \iff \langle c'.fd, t \rangle \in_P c' \text{ and } c' = \min\{c'' \mid c \leq_P c'' \text{ and } \exists t'. \langle c''.fd, t' \rangle \in_P c''\}$
\in_P	method is contained in a class $\langle md, T, V, e \rangle \in_P c \iff \langle md, T, V, e \rangle \in_P c' \text{ and}$ $c' = \min\{c'' \mid c \leq_P c'' \text{ and } \exists e', V'. \langle md, T, V', e' \rangle \in_P c''\}$

Figure 7: CLASSICJAVA relations

and sketch a proof of type soundness for JACQUES. In the following section, we discuss our design decisions and how the formal model helped making them.

JACQUES subsumes almost all of CLASSICJAVA, except for field shadowing. This simplifies the description of field acquisition. In section 6, we revisit the issue.

4.1 Syntax

To support acquisition along containment links between objects, we must have some way to distinguish between those fields that represent containment in the sense of a containment hierarchy and those that do not. Therefore, in JACQUES, the programmer must explicitly tag those fields that represent a containment relationship. JACQUES also requires explicit acquisition, i.e., classes must explicitly declare the fields they acquire from their environment. Figure 8 extends CLASSICJAVA's class definition syntax with these annotations.

The `contains` and `acquires` keywords divide the field definitions into three groups. Consider the class definition

```
class A ... {
    X x
    contains Y y
    acquires Z z ...
}
```

and let `a` be an instance of `A`. Fields in the first two groups (here, `x` and `y`) are parts of `a`, and fields in the last group (`z`) are not directly in `a` but are rather acquired from `a`'s context. Further, only objects in fields in the second group (here, `y`) are considered to be *contained* in `a`, and therefore only they may acquire fields from `a`.

Similarly, the `acquires` keyword splits the method definitions into two groups. Methods in the first group behave as in CLASSICJAVA. Methods in the second group, on the other hand, are acquired from the object's context; therefore, their definitions need not include anything beyond their signatures.

Finally, the (possibly empty) sequence of classes specified after the `contained` keyword places restrictions on the types of objects that may contain objects of the class being defined. In particular, if an instance `a` of class `c` is contained in anything, the container's type must be a subtype of one of the classes specified in `c`'s container list. As a consequence, if this list is empty, then `a` cannot be contained (in the spirit of environmental acquisition) in any object.

$$\begin{aligned}
defn &= \text{class } c \text{ extends } c^* \text{ contained } c^* \{ \\
&\quad \text{field}^* \text{ contains } \text{field}^* \text{ acquires } \text{field}^* \\
&\quad \text{meth}^* \text{ acquires } \text{acqmeth}^* \} \\
acqmeth &= t \text{ md}(t^*)
\end{aligned}$$

Figure 8: JACQUES syntax extensions

$$\begin{aligned}
ClassFieldsOK(P) &\text{ no field definition shadows a field defined in a superclass} \\
&\langle t_1, fd, s_1 \rangle \in_P c_1 \text{ and } \langle t_2, fd, s_2 \rangle \in_P c_2 \text{ and } c_1 \neq c_2 \implies c_1 \not\leq_P c_2 \\
containers_P(c) &\text{ set of all possible containers of instances of class } c \\
containers_P(\text{Object}) &= \emptyset \\
containers_P(c) &= \{c_1, \dots, c_n\} \text{ where class } c \dots \text{ contained } c_1 \dots c_n \{ \dots \} \text{ is in } P \\
acqFields_P(c) &\text{ set of all acquired fields in class } c \\
acqFields_P(\text{Object}) &= \emptyset \\
acqFields_P(c) &= \{ \langle t, fd, \text{acquired} \rangle \mid \langle t, fd, \text{acquired} \rangle \in_P c \} \\
&\text{ where class } c \dots \{ \dots \} \text{ is in } P \\
acqMethods_P(c) &\text{ set of all acquired methods in class } c \\
acqMethods_P(\text{Object}) &= \emptyset \\
acqMethods_P(c) &= \{ \langle md, T, V, e, \text{acquired} \rangle \mid \langle md, T, V, e, \text{acquired} \rangle \in_P c \} \\
&\text{ where class } c \dots \{ \dots \} \text{ is in } P \\
\sqsubset_P &\text{ class may be contained} \\
c_1 \sqsubset_P c_2 &\iff \text{there exists } c' \in containers_P(c_1) \text{ such that } c_2 \leq_P c'
\end{aligned}$$

Figure 9: JACQUES relations

4.2 Types

Both the static and dynamic semantics must be able to distinguish between the three kinds of fields and the two kinds of methods. Therefore, we adjust CLASSICJAVA's abstract representation of fields and methods. A field definition is now a triple $\langle t, fd, s \rangle$, where t is the field's type, fd is its name, and s is a tag indicating the kind of field; s is one of **normal**, **contained**, or **acquired**.

Similarly, a method definition is a 5-tuple $\langle md, T, V, e, s \rangle$, where md , T , V , and e are the method's name, type signature, argument names, and body expression, as in CLASSICJAVA, and s is a tag indicating where the method originated; s may be either **normal** or **acquired**. (The third and fourth elements of the tuple are meaningless for acquired methods, but we leave them in to simplify the definition of \in_P and other relations.)

Finally, we also use an adjusted instance representation: $\langle c, ctnr, \mathcal{F} \rangle$ where c is the class tag, $ctnr$ is a (possibly null) reference to the instance's container, and \mathcal{F} is a finite map from field names to values. We also define three selector functions: $classTag$, $ctnr$, and $fields$ map an instance triple to its class tag, container reference, and field function, respectively.

JACQUES's semantics requires many of the relations and predicates defined in figure 7, with appropriate modifications to several, including \in_P and \leq_P . The additional necessary predicates are straightforward; their definitions are provided in figures 9 and 10.

We use the idea of the class containment graph G_P to describe the possible object containment relationships in a JACQUES program. Figure 10 provides the formal definition of G_P . Intuitively, there is a path from class c_1 to class c_2 in G_P if an instance of c_1 can be contained in an instance of c_2 according to the static annotations in the program source. (Not all paths in G_P correspond to potential containment relationships, but these extra paths do not present any problems.)

The predicates $canAcq_P$ and $canAcqMeth_P$, also defined in figure 10, specify whether a class c may safely acquire a specific field or method. The class c may safely acquire the field fd with type t if one of the

$\prec_{:P}$ subtyping
 $t \prec_{:P} t' \iff t \leq_P t'$
 $(t_1 \dots t_m \rightarrow t) \prec_{:P} (t'_1 \dots t'_n \rightarrow t') \iff (m = n \text{ and } t \prec_{:P} t' \text{ and } t'_i \prec_{:P} t_i \text{ for all } i \in [1, n])$

G_P Class containment graph
 $G_P = (V_P, E_P)$
 where $V_P = \text{dom}(\prec_P) \cup \{\text{Object}\}$
 and $E_P = \{(c_1, c_2) \in (V_P)^2 \mid c_2 \in \text{containers}_P(c_1) \text{ or } c_2 \prec_P c_1\}$

$\text{canAcq}_P(c, \langle t, fd, s \rangle)$ class c may safely acquire field fd with type t from its containers
 for all paths c_0, c_1, \dots in G_P , where $c_0 = c$
 if there exists an $i \geq 1$ such that $\langle t', fd, _ \rangle \in_P c_i$ and $t' \prec_{:P} t$
 then there must exist a $j \in [1, i)$ such that
 $\langle t'', fd, s \rangle \in_P c_j$ and $t'' \prec_{:P} t$ and $s \neq \text{acquired}$.

$\text{canAcqMeth}_P(c, \langle md, (t_1 \dots t_n \rightarrow t), V, e, s \rangle)$
 (class c can safely acquire method md with the given type)
 for all paths c_0, c_1, \dots in G_P , where $c_0 = c$
 if there exists an $i \geq 1$ such that $\langle md, (t'_1 \dots t'_m \rightarrow t'), _, _, _ \rangle \in_P c_i$ and
 $t'_1 \dots t'_m \rightarrow t' \prec_{:P} t_1 \dots t_n \rightarrow t$
 then there must exist a $j \in [1, i)$ such that
 $\langle md, (t''_1 \dots t''_n \rightarrow t''), _, _, s \rangle \in_P c_j$ and
 $(t''_1 \dots t''_n \rightarrow t'') \prec_{:P} (t_1 \dots t_n \rightarrow t)$ and
 $s \neq \text{acquired}$.

Figure 10: Acquisition safety conditions

PROG

$$\frac{\begin{array}{c} \text{ClassesOnce}(P) \quad \text{MethodOncePerClass}(P) \quad \text{FieldOncePerClass}(P) \\ \text{CompleteClasses}(P) \quad \text{WellFoundedClasses}(P) \quad \text{ClassMethodsOK}(P) \quad \text{ClassFieldsOK}(P) \\ P \vdash_{\text{d}} \text{defn}_i \Rightarrow \text{defn}'_i \text{ for all } i \in [1, n] \quad P, \emptyset \vdash_{\text{e}} e \Rightarrow e' : t \quad \text{where } P = \text{defn}_1 \dots \text{defn}_n \text{ e} \end{array}}{\vdash_{\text{p}} \text{defn}_1 \dots \text{defn}_n \text{ e} \Rightarrow \text{defn}'_1 \dots \text{defn}'_n \text{ e}' : t}$$

DEFN

$$\frac{\begin{array}{c} P \vdash_{\text{t}} t_j \text{ for all } j \in [1, m] \quad P \vdash_{\text{t}} c_j \text{ for all } j \in [1, n] \quad P, c \vdash_{\text{m}} \text{meth}_i \Rightarrow \text{meth}'_i \text{ for all } i \in [1, p] \\ \forall f \in \text{acqFields}_P(c) . \text{canAcq}_P(c, f) \quad \forall m \in \text{acqMethods}_P(c) . \text{canAcqMeth}_P(c, m) \\ \text{containers}_P(c') \subseteq \text{containers}_P(c) \quad t_i \sqsubseteq_P c \text{ for all } i \in [k+1, \ell] \end{array}}{P \vdash_{\text{d}} \text{class } c \text{ extends } c' \text{ contained } c_1, \dots, c_n \Rightarrow \text{class } c \text{ extends } c' \text{ contained } c_1, \dots, c_n \\ \left\{ \begin{array}{l} t_1 \text{ fd}_1 \dots t_k \text{ fd}_k \\ \text{contains } t_{k+1} \text{ fd}_{k+1} \dots t_\ell \text{ fd}_\ell \\ \text{acquires } t_{\ell+1} \text{ fd}_{\ell+1} \dots t_m \text{ fd}_m \\ \text{meth}_1 \dots \text{meth}_p \\ \text{acquires } \text{acqmeth}_1 \dots \text{acqmeth}_q \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} t_1 \text{ fd}_1 \dots t_k \text{ fd}_k \\ \text{contains } t_{k+1} \text{ fd}_{k+1} \dots t_\ell \text{ fd}_\ell \\ \text{acquires } t_{\ell+1} \text{ fd}_{\ell+1} \dots t_m \text{ fd}_m \\ \text{meth}'_1 \dots \text{meth}'_p \\ \text{acquires } \text{acqmeth}_1 \dots \text{acqmeth}_q \end{array} \right\}}$$

METHOD

$$\frac{P \vdash_{\text{t}} t \quad P \vdash_{\text{t}} t_i \text{ for all } i \in [1, n] \quad P, \{\text{this} \mapsto t_0, \text{var}_1 \mapsto t_1, \dots, \text{var}_n \mapsto t_n\} \vdash_{\text{s}} e \Rightarrow e' : t}{P, t_0 \vdash_{\text{m}} t \text{ md}(t_1 \text{ var}_1 \dots t_n \text{ var}_n) \{e\} \Rightarrow t \text{ md}(t_1 \text{ var}_1 \dots t_n \text{ var}_n) \{e'\}}$$

$$\frac{\text{SUBTYPE} \quad P, \Gamma \vdash_{\text{e}} e \Rightarrow e' : t \quad t' <:_P t}{P, \Gamma \vdash_{\text{s}} e \Rightarrow e' : t'} \quad \frac{\text{TYPE EXISTENCE} \quad t \in \text{dom}(\prec_P) \cup \{\text{Object}\}}{P \vdash_{\text{t}} t}$$

Figure 11: JACQUES type elaboration rules

following conditions holds for all paths in the class containment graph that begin with c :

1. The field fd is not defined in any class on the path. If any of their subclasses define fd , they define it with a compatible type.
2. Alternatively, the first class on the path that defines the field fd defines it with a compatible type. Further, if fd is defined in any subclass of a class that appears earlier on the path, it must be defined with a compatible type.

The class c may acquire the method md under similar circumstances; the only significant difference is in the notion of a “compatible type.”

Surprisingly, these relations allow the programmer to define a class that acquires a field that is not provided by any class in the program. While this may seem overly permissive, it is the most reasonable solution. On the one hand, it has proven challenging to formulate alternative canAcq_P and canAcqMeth_P predicates that accomodate cycles in the class containment graph but forbid paths without matches. On the other hand, even if JACQUES performed these stronger static checks, objects could still have null parent pointers, and the run-time system would still have to perform appropriate checks.

Figures 11 and 12 provide JACQUES’s type rules. These are based on CLASSICJAVA’s type rules, with modifications to allow acquisition. Because JACQUES does not support field shadowing, neither the [GET] nor the [SET] rules annotate field reference or assignment expressions, and we add $\text{ClassFieldsOK}(P)$ to [PROG]’s antecedents.² Further, [SET] prohibits assignments to acquired fields; see section 5.5 for the rationale for this restriction.

The [DEFN] rule contains most of the changes that support acquisition. First, it verifies that class c ’s containers provide the necessary fields and methods with compatible types. Second, c ’s containers must be

²The other predicates in the [PROG] rule are taken *mutatis mutandis* from CLASSICJAVA.

$$\begin{array}{c}
\text{CTOR} \quad \frac{P \vdash_{\mathbf{t}} c}{P, \Gamma \vdash_{\mathbf{e}} \text{new } c \Rightarrow \text{new } c : c} \qquad \text{VARREF} \quad \frac{\text{var} \in \text{dom}(\Gamma) \quad \Gamma(\text{var}) = t}{P, \Gamma \vdash_{\mathbf{e}} \text{var} \Rightarrow \text{var} : t} \qquad \text{NULL} \quad \frac{P \vdash_{\mathbf{t}} t}{P, \Gamma \vdash_{\mathbf{e}} \text{null} \Rightarrow \text{null} : t} \\
\\
\text{FIELDREF} \quad \frac{P, \Gamma \vdash_{\mathbf{e}} e \Rightarrow e' : t' \quad \langle t, \text{fd}, _ \rangle \in_P t'}{P, \Gamma \vdash_{\mathbf{e}} e.\text{fd} \Rightarrow e'.\text{fd} : t} \qquad \text{FIELDSET} \quad \frac{P, \Gamma \vdash_{\mathbf{s}} e_1 \Rightarrow e'_1 : t' \quad \langle t, \text{fd}, s \rangle \in_P t' \quad s \neq \text{acquired} \quad P, \Gamma \vdash_{\mathbf{s}} e_2 \Rightarrow e'_2 : t}{P, \Gamma \vdash_{\mathbf{e}} e_1.\text{fd} = e_2 \Rightarrow e'_1.\text{fd} = e'_2 : t} \\
\\
\text{CALL} \quad \frac{P, \Gamma \vdash_{\mathbf{e}} e \Rightarrow e' : c \quad P, \Gamma \vdash_{\mathbf{s}} e_i \Rightarrow e'_i : t_i \text{ for all } i \in [1, n] \quad \langle \text{md}, (t_1 \dots t_n \rightarrow t), _ , _ , _ \rangle \in_P c}{P, \Gamma \vdash_{\mathbf{e}} e.\text{md}(e_1 \dots e_n) \Rightarrow e'.\text{md}(e'_1 \dots e'_n) : t} \\
\\
\text{SUPER} \quad \frac{P, \Gamma \vdash_{\mathbf{e}} \text{this} \Rightarrow \text{this} : c' \quad c' \prec_P c \quad \langle \text{md}, (t_1 \dots t_n \rightarrow t), _ , _ , _ \rangle \in_P c \quad P, \Gamma \vdash_{\mathbf{s}} e_i \Rightarrow e'_i : t_i \text{ for all } i \in [1, n]}{P, \Gamma \vdash_{\mathbf{e}} \text{super}.\text{md}(e_1 \dots e_n) \Rightarrow \text{super} \equiv \text{this} : c.\text{md}(e'_1 \dots e'_n) : t} \\
\\
\text{WIDECAST} \quad \frac{P, \Gamma \vdash_{\mathbf{s}} e \Rightarrow e' : t}{P, \Gamma \vdash_{\mathbf{e}} \text{cast } t \text{ } e \Rightarrow \text{cast } t \text{ } e' : t} \qquad \text{NARROWCAST} \quad \frac{P, \Gamma \vdash_{\mathbf{e}} e \Rightarrow e' : t' \quad t' \leq_P t}{P, \Gamma \vdash_{\mathbf{e}} \text{cast } t \text{ } e \Rightarrow \text{cast } t \text{ } e' : t} \\
\\
\text{LET} \quad \frac{P, \Gamma \vdash_{\mathbf{e}} e_1 \Rightarrow e'_1 : t_1 \quad P, \Gamma[\text{var} \mapsto t_1] \vdash_{\mathbf{e}} e_2 \Rightarrow e'_2 : t}{P, \Gamma \vdash_{\mathbf{e}} \text{let } \text{var} = e_1 \text{ in } e_2 \Rightarrow \text{let } \text{var} = e'_1 \text{ in } e'_2 : t}
\end{array}$$

Figure 12: JACQUES type elaboration rules, continued

a superset of c 's superclass's containers, to preserve behavioral subtyping.³ Third, this rule ensures that all of the classes that can be in c 's `contained` fields allow themselves to be contained in c , as required by the soundness proof. The [DEFN] rule could also verify that class A can contain instances of class B if and only if B can be contained in A , and that class A may be contained in at least one other class if it acquires any fields or methods. While warnings about such violations are likely to be useful to the programmer, they are not required for type soundness.

4.3 Semantics

JACQUES's semantics is based on that of CLASSICJAVA. Figure 13 defines the JACQUES reduction relation; figure 14 provides additional supplementary functions. These reductions also perform run-time checks to ensure that acquisition is well-defined. First, we require that there be no containment cycles in the program; the [ct-set] and [cycle-set] rules enforce this invariant and signal a “container cycle” if it is violated. Second, the [ct-set] and [x-ct-set] rules ensure that, on assignment to a `contained` field, the field's new contents are not already contained elsewhere; they signal an “already contained” error otherwise. Finally, because any object's container reference can be null, the [aget], [acall], and [asuper] rules must ensure the presence of a container that provides the acquired field or method; the [xaget], [xacall], and [xasuper] rules signal an “incomplete context” error otherwise.

4.4 Jacques Soundness

JACQUES also satisfies a type soundness theorem:

Theorem 2 (2: Type Soundness for JACQUES) *If*

$$\vdash_P P \Rightarrow P' : t \text{ and } P' = \text{defn}_1 \dots \text{defn}_n e,$$

then either

- $P' \vdash \langle e, \emptyset \rangle \uparrow$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{obj}, \mathcal{S} \rangle$ and $\text{classTag}(\mathcal{S}(\text{obj})) \leq_P t$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{null}, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{error: bad cast}, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{error: dereferenced null}, \mathcal{S} \rangle$;

or

- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{error: incomplete context}, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{error: already contained}, \mathcal{S} \rangle$; or
- $P' \vdash \langle e, \emptyset \rangle \hookrightarrow^* \langle \text{error: container cycle}, \mathcal{S} \rangle$.

A comparison of theorem 1 and theorem 2 indicates that JACQUES can throw additional exceptions that do not occur in CLASSICJAVA programs. The last three exceptions arise from acquisition, and we cannot avoid them without a complex analysis, a more sophisticated type system, or different linguistic mechanisms. Put differently, we have succeeded in making acquisition just as safe as ML arrays; many errors can be caught statically, but some run-time checks remain necessary.

The soundness proof follows the standard method [20] as applied to CLASSICJAVA [5]. In addition to the subject reduction and progress lemmas, we also require some definitions and supporting lemmas beyond those used in the CLASSICJAVA soundness proof; they are stated below.

³As we move down the inheritance hierarchy, the set of possible containers must increase monotonically. However, a subclass may also add new acquired fields or methods; this further constrains the set of possible containers. There is no inconsistency hidden here; it is simply a trade-off that the programmer must consider during the design of class hierarchies.

$e = \dots \mid obj$	
$v = obj \mid null$	
$E = [] \mid E.fd \mid E.fd = e \mid v.fd = E \mid E.md(e\dots) \mid v.md(v\dots E e\dots)$	
$\mid \text{super} \equiv v : c.md(v\dots E e\dots) \mid \text{cast } t E \mid \text{let } var = E \text{ in } e$	
$P \vdash \langle E[\text{new } c], \mathcal{S} \rangle \hookrightarrow \langle E[obj], \mathcal{S}[obj \mapsto \langle c, null, \mathcal{F} \rangle] \rangle$	[new]
$\text{where } obj \notin \text{dom}(\mathcal{S}) \text{ and } \mathcal{F} = \{fd \mapsto null \mid \langle t, fd, s \rangle \in_P c \text{ and } s \neq \text{acquired}\}$	
$P \vdash \langle E[obj.fd], \mathcal{S} \rangle \hookrightarrow \langle E[\mathcal{F}(fd)], \mathcal{S} \rangle$	[get]
$\text{where } \mathcal{F} = \text{fields}(\mathcal{S}(obj)) \text{ and } fd \in \text{dom}(\mathcal{F})$	
$P \vdash \langle E[obj.fd], \mathcal{S} \rangle \hookrightarrow \langle E[\text{getAcqField}_P(\mathcal{S}, obj, fd)], \mathcal{S} \rangle$	[aget]
$\text{where } \langle t, fd, \text{acquired} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \text{getAcqField}_P(\mathcal{S}, obj, fd) \neq \text{notFound}$	
$P \vdash \langle E[obj.fd], \mathcal{S} \rangle \hookrightarrow \langle \text{error: incomplete context}, \mathcal{S} \rangle$	[xaget]
$\text{where } \langle t, fd, \text{acquired} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \text{getAcqField}_P(\mathcal{S}, obj, fd) = \text{notFound}$	
$P \vdash \langle E[obj.fd = v], \mathcal{S} \rangle \hookrightarrow \langle E[v], \text{updateField}(\mathcal{S}, obj, fd, v) \rangle$	[set]
$\text{where } \langle t, fd, \text{normal} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$P \vdash \langle E[obj.fd = null], \mathcal{S} \rangle \hookrightarrow \langle E[null], \text{updateCtdField}(\mathcal{S}, obj, fd, null) \rangle$	[n-ct-set]
$\text{where } \langle t, fd, \text{contained} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$P \vdash \langle E[obj.fd = v], \mathcal{S} \rangle \hookrightarrow \langle \text{error: already contained}, \mathcal{S} \rangle$	[x-ct-set]
$\text{where } \langle t, fd, \text{contained} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } v \in \text{dom}(\mathcal{S}) \text{ and } \text{ctnr}(\mathcal{S}(v)) \neq null$	
$P \vdash \langle E[obj.fd = v], \mathcal{S} \rangle \hookrightarrow \langle \text{error: container cycle}, \mathcal{S} \rangle$	[cycle-set]
$\text{where } v \in \text{dom}(\mathcal{S}) \text{ and } \langle t, fd, \text{contained} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \text{ctnr}(\mathcal{S}(v)) = null \text{ and } \text{canReach}_P(\mathcal{S}, obj, v)$	
$P \vdash \langle E[obj.fd = v], \mathcal{S} \rangle \hookrightarrow \langle E[v], \text{updateCtdField}(\mathcal{S}, obj, fd, v) \rangle$	[ct-set]
$\text{where } v \in \text{dom}(\mathcal{S}) \text{ and } \langle t, fd, \text{contained} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \text{ctnr}(\mathcal{S}(v)) = null \text{ and } \neg \text{canReach}_P(\mathcal{S}, obj, v)$	
$P \vdash \langle E[obj.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle E[e_b[obj/\text{this}, v_1/var_1, \dots, v_n/var_n]], \mathcal{S} \rangle$	[call]
$\text{where } \langle md, _, (var_1, \dots, var_n), e_b, \text{normal} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$P \vdash \langle E[obj.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle E[e_b[obj'/\text{this}, v_1/var_1, \dots, v_n/var_n]], \mathcal{S} \rangle$	[acall]
$\text{where } \langle md, _, _, e, \text{acquired} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \langle obj', (var_1, \dots, var_n), e_b \rangle = \text{getAcqMethod}_P(\mathcal{S}, obj, md)$	
$P \vdash \langle E[obj.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle \text{error: incomplete context}, \mathcal{S} \rangle$	[xacall]
$\text{where } \langle md, _, _, e, \text{acquired} \rangle \in_P \text{classTag}(\mathcal{S}(obj))$	
$\text{and } \text{getAcqMethod}_P(\mathcal{S}, obj, md) = \text{notFound}$	
$P \vdash \langle E[\text{super} \equiv obj : c'.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle E[e[obj/\text{this}, v_1/var_1, \dots, v_n/var_n]], \mathcal{S} \rangle$	[super]
$\text{where } \langle md, _, (var_1, \dots, var_n), e, \text{normal} \rangle \in_P c'$	
$P \vdash \langle E[\text{super} \equiv obj : c'.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle E[e[obj'/\text{this}, v_1/var_1, \dots, v_n/var_n]], \mathcal{S} \rangle$	[asuper]
$\text{where } \langle md, _, _, _, \text{acquired} \rangle \in_P c'$	
$\text{and } \langle obj', (var_1, \dots, var_n), e \rangle = \text{getAcqMethod}_P(\mathcal{S}, obj, md)$	
$P \vdash \langle E[\text{super} \equiv obj : c'.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle \text{error: incomplete context}, \mathcal{S} \rangle$	[xasuper]
$\text{where } \langle md, _, _, _, \text{acquired} \rangle \in_P c' \text{ and } \text{getAcqMethod}_P(\mathcal{S}, obj, md) = \text{notFound}$	
$P \vdash \langle E[\text{cast } t' obj], \mathcal{S} \rangle \hookrightarrow \langle E[obj], \mathcal{S} \rangle$	[cast]
$\text{where } \mathcal{S}(obj) = \langle c, \text{ctnr}, \mathcal{F} \rangle \text{ and } c \leq_P t'$	
$P \vdash \langle E[\text{let } var = v \text{ in } e], \mathcal{S} \rangle \hookrightarrow \langle E[e[v/var]], \mathcal{S} \rangle$	[let]
$P \vdash \langle E[\text{cast } t' obj], \mathcal{S} \rangle \hookrightarrow \langle \text{error: bad cast}, \mathcal{S} \rangle$	[xcast]
$\text{where } \mathcal{S}(obj) = \langle c, \text{ctnr}, \mathcal{F} \rangle \text{ and } c \not\leq_P t'$	
$P \vdash \langle E[\text{cast } t' null], \mathcal{S} \rangle \hookrightarrow \langle \text{error: dereferenced null}, \mathcal{S} \rangle$	[ncast]
$P \vdash \langle E[null.fd], \mathcal{S} \rangle \hookrightarrow \langle \text{error: dereferenced null}, \mathcal{S} \rangle$	[nget]
$P \vdash \langle E[null.fd = v], \mathcal{S} \rangle \hookrightarrow \langle \text{error: dereferenced null}, \mathcal{S} \rangle$	[nset]
$P \vdash \langle E[null.md(v_1 \dots v_n)], \mathcal{S} \rangle \hookrightarrow \langle \text{error: dereferenced null}, \mathcal{S} \rangle$	[ncall]

Figure 13: JACQUES operational semantics

$$\begin{aligned}
getAcqField_P(\mathcal{S}, \text{null}, fd) &= \text{notFound} \\
getAcqField_P(\mathcal{S}, obj, fd) &= \begin{cases} \mathcal{F}(fd) & \text{if } fd \in \text{dom}(\mathcal{F}) \\ getAcqField_P(\mathcal{S}, ctnr, fd) & \text{otherwise} \\ \text{where } \langle _, ctnr, \mathcal{F} \rangle = \mathcal{S}(obj) \end{cases} \\
getAcqMethod_P(\mathcal{S}, \text{null}, md) &= \text{notFound} \\
getAcqMethod_P(\mathcal{S}, obj, md) &= \begin{cases} \langle obj, (var_1, \dots, var_n), e_b \rangle & \text{if } \langle md, _, (var_1, \dots, var_n), e_b, \text{normal} \rangle \in_P c \\ getAcqMethod_P(\mathcal{S}, ctnr, md) & \text{otherwise} \\ \text{where } \langle c, ctnr, _ \rangle = \mathcal{S}(obj) \end{cases} \\
updateField(\mathcal{S}, obj, fd, v) &= \mathcal{S}[obj \mapsto \langle c, ctnr, \mathcal{F}[fd \mapsto v] \rangle] \\
&\quad \text{where } \langle c, ctnr, \mathcal{F} \rangle = \mathcal{S}(obj) \\
updateCtdField(\mathcal{S}, obj, fd, v) &= updateField(\mathcal{S}_2, obj, fd, v) \\
&\quad \text{where } \langle c, ctnr, \mathcal{F} \rangle = \mathcal{S}(obj) \text{ and} \\
&\quad \mathcal{S}_1 = updateContainer(\mathcal{S}, \mathcal{F}(fd), \text{null}) \text{ and} \\
&\quad \mathcal{S}_2 = updateContainer(\mathcal{S}_1, v, obj) \\
updateContainer(\mathcal{S}, v, ctnr) &= \begin{cases} \mathcal{S} & \text{if } v = \text{null} \\ \mathcal{S}[v \mapsto \langle c, ctnr, \mathcal{F} \rangle] & \text{otherwise} \\ \text{where } v \in \text{dom}(\mathcal{S}) \text{ and } \langle c, _, \mathcal{F} \rangle = \mathcal{S}(v) \end{cases} \\
canReach_P(\mathcal{S}, obj, v) &\iff obj = v \text{ or } (ctnr \neq \text{null} \text{ and } canReach_P(\mathcal{S}, ctnr, v)) \\
&\quad \text{where } ctnr = ctnr(\mathcal{S}(obj))
\end{aligned}$$

Figure 14: JACQUES supplementary functions

Definition 1 (Object Containment) We write $P, \mathcal{S} \vdash obj_1 \sqsupseteq obj_2$ to indicate that obj_1 contains obj_2 in store \mathcal{S} and program P . Formally:

$$\begin{aligned} P, \mathcal{S} \vdash obj_1 \sqsupseteq obj_2 &\iff \\ \mathcal{S}(obj_1) &= \langle c, ctnr, \mathcal{F} \rangle \text{ and } \mathcal{F}(fd) = obj_2 \\ &\text{for some } fd \text{ such that } \langle _, fd, \text{contained} \rangle \in_P c. \end{aligned}$$

Definition 2 (Object Containment Graph) For a given store \mathcal{S} , we construct a directed graph $G_{\mathcal{S}}$, called the object containment graph, as follows:

$$G_{\mathcal{S}} = (V, E)$$

where

$$V = \text{rng}(\mathcal{S}) \text{ and } E = \{(x, y) \in V^2 \mid ctnr(x) = y\}.$$

In other words, there is a node in $G_{\mathcal{S}}$ for each object in \mathcal{S} , and an edge from each node to its container.

Next we need to adapt the environment-store consistency relation from the proof of theorem 1 to the new context.

Definition 3 (Environment-Store Consistency) We write $P, \Gamma \vdash_{\sigma} \mathcal{S}$ if the type environment Γ and the store \mathcal{S} are consistent with one another, given program P . Formally:

$$\begin{aligned} P, \Gamma \vdash_{\sigma} \mathcal{S} &\iff \\ &obj \in \text{dom}(\Gamma) \implies obj \in \text{dom}(\mathcal{S}) && (esc-1) \\ \wedge \text{dom}(\mathcal{S}) &\subseteq \text{dom}(\Gamma) && (esc-2) \\ \wedge G_{\mathcal{S}} &\text{ has no cycles and is finite} && (esc-3) \\ \wedge [\mathcal{S}(obj) = \langle c, ctnr, \mathcal{F} \rangle \implies && \\ &\Gamma(obj) = c && (esc-4) \\ \wedge \text{dom}(\mathcal{F}) = \{fd \mid \langle _, fd, s \rangle \in_P c, s \in \{\text{normal, contained}\}\} && (esc-5) \\ \wedge \text{rng}(\mathcal{F}) &\subseteq \text{dom}(\mathcal{S}) \cup \{\text{null}\} && (esc-6) \\ \wedge (\mathcal{F}(fd) \in \text{dom}(\mathcal{S}) \text{ and } \langle c', fd, _ \rangle \in_P c) \implies && (esc-7) \\ &\text{classTag}(\mathcal{S}(\mathcal{F}(fd))) \leq_P c' \\ \wedge (\mathcal{F}(fd) = \mathcal{F}(fd') \neq \text{null} \text{ and } \langle t, fd, \text{contained} \rangle \in_P c \text{ and} && (esc-8) \\ &\langle t', fd', \text{contained} \rangle \in_P c) \implies fd = fd' \\ \wedge ctnr \in \text{dom}(\mathcal{S}) \iff P, \mathcal{S} \vdash ctnr \sqsupseteq obj && (esc-9) \\ \wedge ctnr \in \text{dom}(\mathcal{S}) \implies && (esc-10) \\ &\exists c' \in \text{containers}_P(c). \text{classTag}(\mathcal{S}(ctnr)) \leq_P c' \\ \wedge ctnr \in \text{dom}(\mathcal{S}) \cup \{\text{null}\} && (esc-11) \end{aligned}$$

Most of the clauses of this definition are present in CLASSICJAVA; clauses 3 and 8–11 are new for JACQUES. In order, they place the following requirements on the store:

- (esc-3): The object containment graph must be a finite set of finite trees. (The definition of $G_{\mathcal{S}}$ also implies that these trees are disjoint.)
- (esc-8): No object may have two contained fields that refer to the same object. This allows us to set an object's container reference to null when we remove it from its container without breaking any other invariants.
- (esc-9): The object obj 's container reference is $ctnr$ if and only if $ctnr$ contains obj . (As a consequence, if $ctnr$ is null, then no object may contain obj .)
- (esc-10): If obj has class tag c and container obj' , then obj' must have a class tag that is a subtype of one of class c 's containers.
- (esc-11): An object's container reference is either a valid object reference or null.

Lemma 1 Let P be a well-typed program, and let Γ and \mathcal{S} be a type environment and store, respectively, such that $P, \Gamma \vdash_{\sigma} \mathcal{S}$. For every path from obj_1 to obj_n in $G_{\mathcal{S}}$, where $n > 1$, there exists a path $c_1, \dots, c_2, \dots, c_n$ in G_P , such that $\mathcal{S}(obj_i) = \langle c_i, _, _ \rangle$ for all $i \in [1, n]$, and such that for all $i \in [1, n]$, we can decompose the sub-path c_i, \dots, c_{i+1} into $c_{i0}, c_{i1}, \dots, c_{im}$, where $c_{i0} = c_i$ and $c_{im} = c_{i+1}$ and $c_{i1} \in \text{containers}_P(c_{i0})$ and $c_{i(j+1)} \prec_P c_{ij}$ for all $j \in [2, m]$. We say that any path in G_P that corresponds to a path in $G_{\mathcal{S}}$ and that decomposes in this manner is well-structured.

Proof 1 Strong induction on n , the length of the path.

Base case: $n = 2$.

Let $\langle c_1, _, _ \rangle = \mathcal{S}(obj_1)$ and $\langle c_2, _, _ \rangle = \mathcal{S}(obj_2)$. By (esc-10), there exists a class $c' \in \text{containers}_P(c_1)$ such that $c_2 \leq_P c'$. Since \leq_P is the reflexive-transitive closure of \prec_P , there exists a sequence of classes c'_1, \dots, c'_m such that $c_2 \prec_P c'_m \prec_P \dots \prec_P c'_1$ and $c'_1 = c'$. By construction of G_P , then, $c_1, c'_1, \dots, c'_m, c_2$ is a path in G_P , and by the definition above, it is well-structured.

Inductive step: Decompose the path obj_1, \dots, obj_n into two paths obj_1, obj_2 and obj_2, \dots, obj_n . By induction, there exist well-structured paths c_1, \dots, c_2 and c_2, \dots, c_n in G_P , such that $\mathcal{S}(obj_1) = \langle c_1, _, _ \rangle$, $\mathcal{S}(obj_2) = \langle c_2, _, _ \rangle$, and $\mathcal{S}(obj_n) = \langle c_n, _, _ \rangle$. The concatenation of these paths $c_1, \dots, c_2, \dots, c_n$ is a path in G_P , and it is also well-structured. ■

Lemma 2 Let P be a well-typed program, and let Γ and \mathcal{S} be a type environment and a store, respectively, such that $P, \Gamma \vdash_{\sigma} \mathcal{S}$. If

- $\mathcal{S}(obj) = \langle c, _, _ \rangle$, and
- $\langle t, fd, \text{acquired} \rangle \in_P c$, and
- $\text{canAcq}_P(c, \langle t, fd, \text{acquired} \rangle)$, and
- $\text{getAcqField}_P(\mathcal{S}, obj, fd) \Downarrow v$ where $v \neq \text{notFound}$,

then either

- $v = \text{null}$, or
- $v \in \text{dom}(\mathcal{S})$ and $\mathcal{S}(v) = \langle c_v, _, _ \rangle$ and $c_v <:_P t$.

Similarly, if

- $\mathcal{S}(obj) = \langle c, _, _ \rangle$, and
- $\langle md, (t_1 \dots t_n \rightarrow t), V, e, \text{acquired} \rangle \in_P c$, and
- $\text{canAcqMeth}_P(c, \langle md, (t_1 \dots t_n \rightarrow t), V, e, \text{acquired} \rangle)$, and
- $\text{getAcqMethod}_P(\mathcal{S}, obj, md) \Downarrow \langle v, (var_1, \dots, var_n), e_b \rangle$,

then

- $\mathcal{S}(v) = \langle c_v, _, _ \rangle$, and
- $\langle md, (t'_1 \dots t'_n \rightarrow t'), (var_1, \dots, var_n), e_b, _ \rangle \in_P c_v$, and
- $(t'_1 \dots t'_n \rightarrow t') <:_P (t_1 \dots t_n \rightarrow t)$.

Proof 2 We provide the proof only for fields; that for methods is similar.

Since $\text{getAcqField}_P(\mathcal{S}, obj, fd) \Downarrow v$ and $v \neq \text{notFound}$, there must exist a path $obj_0, obj_1, \dots, obj_n$ in $G_{\mathcal{S}}$, where $obj_0 = obj$ and $\mathcal{S}(obj_i) = \langle c_i, _, _ \rangle$ for all $i \in [0, n]$ and $\langle _, fd, s_i \rangle \in_P c_i \implies s_i = \text{acquired}$ for all $i \in [0, n]$ and $\langle t', fd, s' \rangle \in_P c_n$ and $s' \neq \text{acquired}$.

By lemma 1, there exists a well-structured path $c_0, \dots, c_1, \dots, c_n$ in G_P , where $c_0 = c$. Assume for the sake of a later contradiction that $t' \not<:_P t$. Then $\text{canAcq}_P(c, \langle t, fd, \text{acquired} \rangle)$ requires that there be some class $c' \neq c_n$ on this well-structured path such that $\langle t'', fd, s \rangle \in_P c'$, $t'' <:_P t$, and $s \neq \text{acquired}$. Because $\langle t, fd, \text{acquired} \rangle \in_P c_0$, we know that $c' \neq c_0$. Therefore, three cases remain:

1. $c' = c_i$ for some $i \in [1, n]$. This contradicts our earlier conclusion that none of the intervening classes c_1, \dots, c_{n-1} define the field fd .

$$\begin{aligned} \text{redex} &= \text{new } c \mid v.\text{fd} \mid v.\text{fd} = v \mid v.\text{md}(v^*) \\ &\mid \text{super} \equiv v : \underline{c}.\text{md}(v^*) \mid \text{cast } t \ v \mid \text{let } \text{var} = v \text{ in } e \end{aligned}$$

Figure 15: JACQUES redexes

2. $c' = c_n$. This contradicts our assumption that c_n defines the field with an incompatible type.
3. c' is on a path c_i, \dots, c_{i+1} for some $i \in [0, n)$, but $c' \neq c_i$ and $c' \neq c_{i+1}$. Because this path is well-structured, $c_{i+1} \leq_P c'$. Because P is well-typed, then $\text{ClassFieldsOK}(P)$ must hold, which requires that $\langle t'', \text{fd}, s \rangle \in_P c_{i+1}$. This violates our earlier assumptions. If $i + 1 < n$, then an intermediate class provides the field. Otherwise, $i + 1 = n$, which violates our assumption that c_n provides the field with an incompatible type.

Therefore, our assumption $t' \prec_P t$ must be false, and $t' <_P t$. Then $v = \mathcal{F}_n(\text{fd})$, where $\mathcal{S}(\text{obj}_n) = \langle c_n, _, \mathcal{F}_n \rangle$. If $\mathcal{F}_n(\text{fd}) = \text{null}$, the lemma is satisfied. Otherwise, $\mathcal{F}_n(\text{fd}) \in \text{dom}(\mathcal{S})$ by (esc-6), and (esc-7) then requires that $\mathcal{S}(\mathcal{F}_n(\text{fd})) = \langle c_v, _, _ \rangle$, where $c_v <_P t'$. This implies $c_v <_P t$, as required. ■

The next lemma uses the predicate superOk from CLASSICJAVA.

Definition 4 (Well-Formed Super Calls)

$$\text{superOk}(e) \iff \text{For all subexpressions } \text{super} \equiv e_0 : \underline{c}.\text{md}(e_1 \dots e_n) \text{ of } e, \text{ either } e_0 = \text{this} \text{ or there exists an object reference } \text{obj} \text{ such that } e_0 = \text{obj}.$$

As with CLASSICJAVA, we must also extend our typing relation to elaborated expressions and situations that may only arise at run-time. The extension mimics CLASSICJAVA; the extended judgment is written $P, \Gamma \vdash_{\underline{e}} e : t$, and we elide the precise definition for space reasons. With this last definition, we can state a subject reduction lemma for JACQUES; the proof is provided as an appendix.

Lemma 3 (3: Subject Reduction) *If $P, \Gamma \vdash_{\underline{e}} e : t$ and $P, \Gamma \vdash_{\sigma} \mathcal{S}$ and $\text{superOk}(e)$ and $P \vdash \langle e, \mathcal{S} \rangle \leftrightarrow \langle e', \mathcal{S}' \rangle$, then either e' is an error configuration or there exists a type environment Γ' such that*

1. $P, \Gamma' \vdash_{\underline{e}} e' : t$, and
2. $P, \Gamma' \vdash_{\sigma} \mathcal{S}'$, and
3. $\text{superOk}(e')$.

To finish the proof of the soundness theorem, we must state and prove a Progress Lemma. For this proof, we need the definition of a redex, given in figure 15, and a unique decomposition lemma.

Lemma 4 (4: Unique Decomposition) *For any closed expression e that is not a value, there must exist a context E and a redex e_r such that $E[e_r] = e$. Further, if $E[e_r] = e$ and $E'[e'_r] = e$, then $E = E'$ and $e_r = e'_r$.*

The proof is a straightforward structural induction on the term e .

Lemma 5 (5: Progress) *If e is closed, $P, \Gamma \vdash_{\underline{e}} e : t$, $P, \Gamma \vdash_{\sigma} \mathcal{S}$, and $\text{superOk}(e)$, then either e is a value or there exist e', \mathcal{S}' such that $P \vdash \langle e, \mathcal{S} \rangle \leftrightarrow \langle e', \mathcal{S}' \rangle$, and e' is closed.*

Proof 3 (of lemma 5) The case where e is a value follows directly, so assume that e is not a value. By lemma 4, we know that there exist E and e_r such that $e = E[e_r]$. The proof proceeds by case analysis on e_r . Most of the cases are as in CLASSICJAVA; we show only those cases that differ here.

- $e_r = v.fd$. If $v = \text{null}$, then $[nget]$ applies. Otherwise, we know from [FIELDREF] that $P, \Gamma \vdash_{\mathbf{e}} v : t'$ and $\langle t, fd, s \rangle \in_P t'$. Let $\langle c, ctnr, \mathcal{F} \rangle = \mathcal{S}(v)$. If $s = \text{normal}$ or $s = \text{contained}$, then by (esc-5), $fd \in \text{dom}(\mathcal{F})$, and $[get]$ applies. Otherwise, $s = \text{acquired}$, and we must evaluate $getAcqField_P(\mathcal{S}, v, fd)$. There are three possibilities:

1. $getAcqField_P(\mathcal{S}, v, fd) = v'$, so $[aget]$ applies, and $P \vdash \langle E[v.fd], \mathcal{S} \rangle \leftrightarrow \langle E[v'], \mathcal{S} \rangle$. We know that $E[v.fd]$ is closed, and substituting v' for $v.fd$ neither removes binding forms nor adds variable references, so $E[v']$ is closed as required.
2. $getAcqField_P(\mathcal{S}, v, fd) = \text{notFound}$, and $[xaget]$ applies.
3. $getAcqField_P(\mathcal{S}, v, fd) \uparrow$.

However, this last case cannot happen. The function $getAcqField_P$ traverses $G_{\mathcal{S}}$, starting from $\mathcal{S}(v)$, and stops either if there are no more outgoing edges or if it finds the desired field. The definition of $G_{\mathcal{S}}$ implies that each node has at most one outgoing edge, so there is exactly one path in the graph that starts with v . And from (esc-3), we know that this path is of finite length, so $getAcqField_P$ always terminates, and case 3 is impossible.

- $e_r = v_1.fd = v_2$. This case proceeds much as the previous case; we highlight the differences. First, by [FIELDSET], we know that fd is in the appropriate class and is not an acquired field. The only interesting case occurs when fd is a contained field. If $v_2 = \text{null}$, then $[n-ct-set]$ applies. Otherwise, there are three interesting cases:
 1. $ctnr(\mathcal{S}(v_2)) \neq \text{null}$; $[x-ct-set]$ applies, and the result of the reduction is closed by the same logic as in the previous case.
 2. $ctnr(\mathcal{S}(v_2)) = \text{null}$ and $canReach_P(\mathcal{S}, v_1, v_2)$; $[cycle-set]$ applies.
 3. $ctnr(\mathcal{S}(v_2)) = \text{null}$ and $\neg canReach_P(\mathcal{S}, v_1, v_2)$; $[ct-set]$ applies.

The $canReach_P$ predicate also traverses a path in $G_{\mathcal{S}}$, so it must converge to a value by the same logic as used above with $getAcqField_P$.

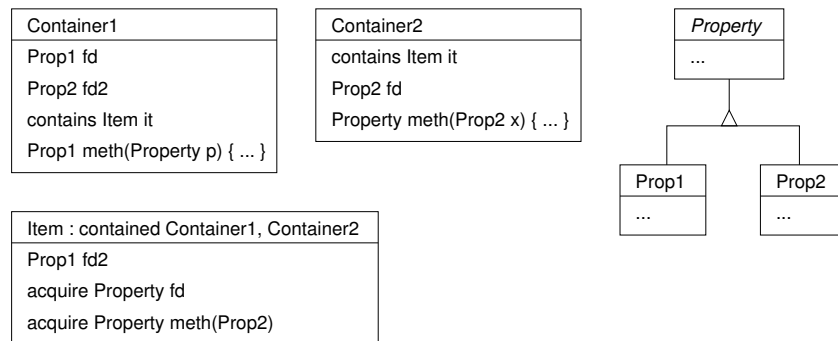
- $e_r = v.md(v_1 \dots v_n)$. Parallel to the field reference case above; one of $[call]$, $[acall]$, $[xacall]$, or $[ncall]$ applies. If $[call]$ or $[acall]$ applies, then this reduction replaces the original method call expression with $e_b[v/\text{this}, v_1/var_1, \dots, v_n/var_n]$, where e_b is the body of the method in question and var_1, \dots, var_n are the method's formal parameters. The expression e_b was originally typechecked, by the [METHOD] rule, in a type environment Γ , where $\text{dom}(\Gamma) = \{\text{this}, var_1, \dots, var_n\}$. Therefore the free variables of e_b must be a subset of $\text{dom}(\Gamma)$, and so the result of the substitution must be closed, and therefore e' is closed as well.
- $e_r = \text{super} \equiv v : c.md(v_1 \dots v_n)$. Similar to the previous case, except that $superOk(e)$ implies that $v \neq \text{null}$, so one of $[super]$, $[asuper]$, or $[xasuper]$ applies. ■

5 Jacques Design Issues

In developing JACQUES, we encountered several interesting design problems. We rejected some design alternatives because they break type safety; for others, we based our decisions on pragmatics. In this section, we discuss these design choices and provide rationales. We use the program fragment in figure 16 as a running example.

5.1 Containment Cycles

The object containment graph cannot have cycles, as Gil and Lorenz explain. This does not mean, however, that the *class* containment graph should be acyclic as well. As stated in section 2, acquisition is often useful in conjunction with the composite pattern. It is frequently essential in this pattern that objects be allowed to contain other instances of their own class (or a superclass). For a concrete example, consider a simple recursive list implementation, as in figure 17. In this case, a `Node` instance must be allowed to contain



```

class Main {
    Container1 aCtnr;
    Container2 aCtnr2;
    Item anItem;
    ...
    Object run () {
        ...
        aCtnr.it = anItem;    // anItem contained in aCtnr
        ...
    }
    ...
}
  
```

Figure 16: JACQUES running example

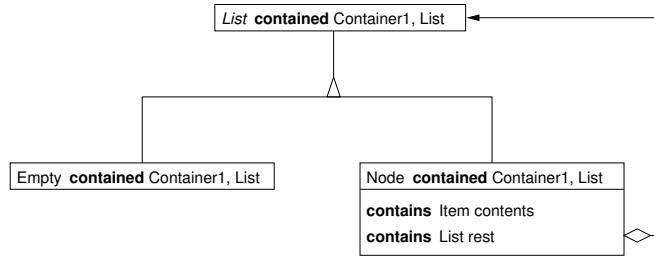


Figure 17: Recursive lists and acquisition

instances of *List* and particularly other *Nodes*. Otherwise, acquisition is stopped at the first *Node* instance, and elements of the list cannot acquire anything from *Container1*.

In general, forbidding this kind of cycle precludes the use of acquisition in *any* inductively-defined data structure or instance of the composite pattern, severely restricting acquisition’s usefulness. Therefore, we must allow cycles within the class containment graph. Acquisition remains well-defined as long as the *object* containment graph has no cycles, and the *[ct-set]* reduction enforces this restriction at runtime.

This design choice is orthogonal to type safety. Our soundness proof in section 4.4 makes no assumptions about whether G_P contains cycles.

5.2 Acquisition by Value and by Name

In their original description of acquisition, Gil and Lorenz do not address field assignment and its interaction with acquisition. With assignment, it becomes important to determine when acquisition takes place. We considered two options, dubbed “acquisition-by-value” and “acquisition-by-name,” by analogy with call-by-value and call-by-name. The primary difference between the two is the point at which the acquiring object receives the value of the container’s field.

In our running example (figure 16) with acquisition-by-value, the acquiring object *anItem* receives the value of its acquired fields when it is placed in its container *aCtnr*. Subsequent modifications to *aCtnr*’s fields are not automatically visible to *anItem*. With acquisition-by-name, on the other hand, *anItem* receives the value of its acquired fields anew each time the fields are referenced. Therefore, modifications to *aCtnr*’s fields are automatically visible to *anItem*. In both cases, type soundness is preserved as long as the value in *aCtnr.f*d has a type compatible with that declared for *anItem.f*d.

The choice of acquisition-by-value raises two questions:

1. If we remove *anItem* from its container *aCtnr*, does it preserve the field values it acquired from *aCtnr*, or do those fields now become undefined?
2. If we then add *anItem* to another container *aCtnr'*, does it preserve its existing field values, or does it replace them with the values then current in its new container?

With acquisition-by-name, in contrast, the answers to these questions follow directly from the definition above. This ambiguity with acquisition-by-value leads us to believe that acquisition-by-name is the better mechanism, from both a semantic as well as a pragmatic perspective.

5.3 Types of Acquired Fields

In the example of figure 16, *anItem* acquires the field *fd* from its environment, specifically the container *aCtnr*. The class *Container1* defines *fd* to have type *Prop1*, but *Item* acquires it with *Property*’s interface. Since *Prop1* is a subtype of *Property*, this acquisition is safe. In general, acquiring classes may expect a more general type

for their acquired fields than their environments actually provide. For increased flexibility, the definition of JACQUES allows just this sort of type variance through the $canAcq_P$ predicate.⁴

5.4 Types of Acquired Methods

In our example, `Item` acquires the method `meth` with argument type `Prop2` and result type `Property`. This is safe so long as the method that is actually executed as a result of calls to `anItem.meth` (that is, the method provided by the context) allows a more general type for its arguments and returns a value of a more specific type. So, since $Prop2 \leq_P Property$ and $Prop1 \leq_P Property$, `Item` can safely acquire `meth` from `Container1`. The definition of $canAcqMeth_P$ allows this flexibility. This is only safe, however, because JACQUES does not support method overloading. As we scale acquisition to a language with method overloading, we may have to restrict method acquisition to use invariant types.

5.5 Assignment to Acquired Fields

Allowing assignment to acquired fields with “contravariant” typing as described above is not type-safe. In our running example, an assignment such as

```
anItem.fd = new Prop2()
```

would type-check because `anItem.fd` is declared to have type `Property`. However, since `anItem` acquires the field `fd` from `aCtnr`, and `aCtnr.fd` has type `Prop1`, allowing assignments such as the above would break type safety, as references to `aCtnr.fd` would reduce to a value of type `Prop2`. There are three possible strategies that would allow assignment to `anItem.fd` while preserving type safety:

1. The type of an acquired field must match exactly in the acquiring class and the container. In our example, this would correspond to requiring `Item` to acquire the field `fd` with type `Prop1` (which would in turn prevent `Container2` from containing `Item`).
2. We could restrict the assignment statement while preserving variance in acquired field types. In this case, the value on the right-hand side of an assignment must have exactly the same type as the left-hand side; subtyping would no longer be allowed here.
3. We could preserve the variance in acquired field types and simply disallow assignments to acquired fields. In our running example, an assignment to `anItem.fd` would be disallowed because this field is acquired from another object, but assignments to `aCtnr.fd` would be legal. That `aCtnr.fd` is acquired by another object is not relevant.

The second option has several drawbacks. First, as long as we preserve subsumption throughout the rest of the language, this restriction could only be checked dynamically, which would complicate the semantics yet further. Second, it would restrict the types allowed on the left-hand side of an assignment statement. In Java, for instance, the left-hand side of such an assignment statement could not have an interface type. Because interfaces can never be instantiated directly, it would be impossible for the value of the right-hand side to have exactly the same type as the left. This would severely constrain assignment statements. Finally, it would forbid subsumption in a particular context. Not only would this create asymmetry in the language design, it would do so in a particularly restrictive fashion, due to the frequent occurrence of assignment statements in object-oriented programs.

Therefore, we must choose between the first and third options. JACQUES implements the third option, as we conjecture that it is more pragmatically useful than the first. Without the benefit of experience with acquisition in software projects, though, we cannot answer this question definitively, and we need to revisit this issue once we have gained more experience.

⁴At first glance, this situation appears similar to, though not exactly the same as, the traditional notion of contravariance: we have $Prop1 \leq_P Property$ but $Item \not\leq_P Container1$, so the idea of one type getting larger while the other gets smaller is not applicable.

5.6 Changing Containers

With the ability to mutate fields in existing objects comes the ability to mutate the object containment graph after it has been established. While this ability may be necessary in some applications, we want to preserve the invariant that `aCtnr` contains `anItem` if and only if `anItem`'s container is `aCtnr`. This allows both objects, but particularly `aCtnr`, to make useful assumptions about the object containment graph at runtime.

Allowing programs, however, to move an object from one container to another in a single assignment statement, can violate the invariant above. If, in our running example, we execute the assignment `aCtnr2.it = anItem`, then `anItem`'s container reference will point to `aCtnr2`, even though `aCtnr` still contains `anItem`.

For this reason, we prohibit programs from moving an object from one container to another in a single assignment statement. Instead, we require the programmer to remove the object from its container before placing it in the new one. In our example, the programmer must first execute `aCtnr.it = null`, thus setting `anItem`'s container reference to `null`, before adding `anItem` to `aCtnr2`. This restriction is enforced by the $[x\text{-ct}\text{-set}_j]$ reduction, which requires an object's container pointer to be `null` before it can be added to a container. An alternative would be to follow Alan Kay's maxim of replacing assignment with higher-level forms of mutation operations [13], specifically a `switch-container` operation.

5.7 Handling incomplete environments

As described above in the discussion of theorem 2, JACQUES cannot statically ensure that references to `anItem.fid` are well-defined, because `anItem` may not be in a container that provides the field `fid`. This arises because `anItem` is placed into its container by assigning it to the field `aCtnr.it`, which is only possible after `anItem` is instantiated. The JACQUES type system does not ensure that code executed between the instantiation and the assignment does not refer to any acquired fields.

An alternative strategy for constructing the object containment graph that would avoid this problem is to require the programmer to supply `aCtnr` as an argument to `anItem`'s constructor.⁵ To prevent the programmer from removing `anItem` from `aCtnr` later, we would also disallow assignments to `contains` fields and instead require a `switch-container` primitive operation that moves an object from one container to another. Of course, to ensure statically that `anItem` is always in a container, we would also have to ensure, statically, that the constructor argument and the operands to `switch-container` are not `null`, by using a type system similar to that described by Fähndrich [4].

Unfortunately, using such a type system has a number of drawbacks. First, object construction leads to significant complications, as Fähndrich and Leino discuss, because all object fields are initialized to `null`, even those that are declared to have non-`null` types. Second, the authors do not discuss how the non-`null` type annotations “poison” the rest of the program, much as `const` declarations in C++ can require changing large numbers of other, only tangentially related, classes. Third, requiring a non-`null` reference as a constructor argument to `anItem` requires programmers to construct all containment hierarchies by beginning with the root and moving towards the leaves. Experience with an early implementation of JACQUES suggests that this requirement can be extremely awkward in cases where another construction mechanism might be more appropriate. In particular, it disallows the construction pattern commonly used in functional languages, which builds the leaves first and constructs the hierarchy from there.

Because of these drawbacks, we believe that supplying the container as a constructor argument and requiring it to be non-`null` is a bigger problem, practically speaking, than the possibility of “incomplete context” exceptions. Clearly, we require additional practical experience with acquisition to decide this issue.

5.8 Forwarding and Delegation Semantics

Gil and Lorenz discuss the choice between forwarding and delegation semantics for acquired methods. To define these terms as applied to environmental acquisition, consider our running example, in which `anItem` acquires the method `meth` from `aCtnr`; recall that `Item` $\not\leq_P$ `Container1`. Under forwarding semantics, when `Main` invokes `anItem.meth`, this is bound to `aCtnr` during the execution of the method's body. With delegation semantics, this is instead bound to `anItem`.

⁵David Lorenz proposed this alternative in a private conversation.

By type rule [METH], the body of `Container1.meth` is elaborated under the assumption that `this` has type `Container1` (or any subtype). Hence, the body can refer to arbitrary features of `Container1`, such as `fd2`. If, instead, `this` were suddenly bound to `anItem` at runtime, references to properties of `Container1` may become undefined or ill-typed. Specifically, a reference to `this.fd2` would reduce to a value of type `Prop1` rather than `Prop2` as expected. We thus conclude that forwarding semantics for environmental acquisition is the only safe possibility in a statically typed language.

5.9 Guarded Field References

As defined, JACQUES throws an exception upon a reference to an acquired field if the object in question is not in an environment that provides the expected field. In their original formulation, Gil and Lorenz propose guarded field references as a protection against this exception. A guarded field reference is simply a field reference expression that includes a default value to be used if the object is not in a suitable environment. If the compiler can prove that, for a particular field reference, the object is always in a suitable environment, the programmer may omit the default.

Since JACQUES's type system cannot prove that objects are always in a suitable environment for *any* field reference, *all* references would have to have guards in order to preserve safety. As a result, proving type soundness with guards is not markedly different from proving it with the incomplete context exception. Further, we conjecture that in practice, most programmers would specify `null` as their guard expression, thus hiding the incomplete context exception within a null pointer exception and thus severely hindering their own debugging efforts. Therefore, we have chosen to omit guards and to raise an exception in cases where an object's environment does not suffice. If practical experience suggests that such guards are useful, we may add them to the language later.

6 Scaling to Full Java

JACQUES is a model of the core of a sequential OO language with acquisition; we have deliberately omitted certain features of full Java [11] and C# [3] to keep the presentation and the proofs simple. In order to extend acquisition to one of these languages, we have to consider how it interacts with other language features not in our model, and we have to explore how it interacts with practical implementation issues.


6.1 Acquisition and Java Features

To scale JACQUES's design to full Java or C#, we must address two features of the larger language that our model does not include: concurrency and field shadowing. It seems that acquisition can co-exist with concurrency without any more than the normal difficulties associated with concurrent programs. In particular, if one thread changes the object containment graph while another is traversing that graph to look up an acquired field value, the results are undefined. It might be possible to define an acquired field reference to be an atomic operation via a Java `synchronized` block, but this seems unnecessary; the difficulties that can arise in situations like this are no worse than those which may arise in concurrent Java programs that involve heavy pointer manipulation.

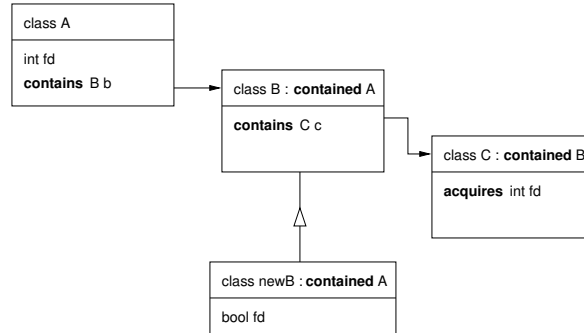
Java-style field shadowing presents a more complicated challenge. In the case where an object acquires a shadowed or shadowing field from its container, it is no longer clear which field should be visible to the acquiring object. Consider the following class definitions:

```
class A extends Object ... {X x ... }
class B extends A ... {X x ... }
class C extends Object contained A { ... acquires X x ... }
```

If `c`, an instance of `C`, is contained within `b`, an instance of `B`, it is most consistent with Java's current behavior if `c` acquires `A.x` rather than `B.x`, as `c` acquires properties from its container as viewed through `A`'s interface. The language should, however, provide the programmer with a way to override this default when necessary, much as Java allows the programmer to cast the object in a field reference expression to select which overloaded field is used.



The original program



The augmented program

Figure 18: Subclassing, acquisition, and incremental compilation

6.2 Incremental Compilation

By “incremental compilation,” we refer to the ability to add a class to a pre-existing program that has already been compiled (and thus type-checked) without needing to recompile the classes from the original program. For environmental acquisition to support this, we must ensure that adding a new class cannot violate invariants that have already been verified on the existing classes. Acquisition makes this challenging, particularly it requires the type checker to verify certain properties of all classes that can contain other objects, plus all of their subclasses. The following example demonstrates the problem.

In the original program (figure 18, top), the [DEFN] rule uses $canAcq_P$ to ensure that C’s acquisition of the field `fd` is safe; in part, this requires that B has no subclasses that define a field `fd` with an incompatible type. In the augmented program (figure 18, bottom), the new class `newB` violates this invariant. This violation, however, is not detected when applying the [DEFN] rule to the `newB` class. Therefore, a naïve solution would have to repeat the type-checking on all classes that could be contained in `newB` instances, directly or indirectly. In the worst case, this would require type-checking the entire program again.

We can solve this problem by modifying the type checker to annotate each class on a containment path with information regarding the fields and methods whose definitions are constrained. In the example above, the type checker would annotate the class B with the field constraint `int fd`. Then, when it checks the new class `newB`, the type checker would proceed up the inheritance hierarchy, verifying that `newB` does not include any conflicting field and method definitions.

Finally, adding a new class as part of a containment path cannot violate any acquisition invariants in the original program—for the simple reason that a programmer cannot add a single class that is directly part of a containment path. If there is a containment relationship between two classes, the containing class must name the contained class and vice versa. Therefore, adding either class alone causes type-checking to fail with an undefined-class error. As a result, the programmer must add all of the classes in a single containment path in a single increment, at which point the type checker can ensure the necessary invariants for all classes on the path.

7 Related Work

The recent work in ownership types [1] introduces constraints on the object composition graph in order to preserve certain desirable invariants, much as we do. In such a type system, the programmer annotates all field declarations, local variable declarations, method arguments, and method return types with ownership information. The type system then ensures that all paths from the root of the object containment graph to an object must pass through the object’s owner. Consequently, if an object is owned by its container, only its container may refer to it. This restricts object aliasing and makes it easier to reason about object-oriented programs in the face of field mutation.

While ownership types and types for environmental acquisition both restrict the shape of the object containment graph, the similarities end there. Even though ownership constraints could help ensure that an object has at most one container, we believe that preventing aliasing to the degree achieved by their system is overly restrictive. In JACQUES, we explicitly allow an object to possess a reference to the contents of another object’s contained field, so long as this additional reference does not imply a containment relationship. Further, we could remove the need for the “already contained” exception without complicating the type system further by disallowing direct assignment to contained fields and providing a `switch-container` operation that manipulates the object containment graph while preserving all desired invariants. Finally, ownership types do not appear to help our existing system’s most severe drawback, namely the inability to detect at compile time whether an object will be in an environment that provides all of the necessary acquired fields and methods.

8 Conclusions and Future Work

In this paper, we have resumed Gil and Lorenz’s language design experiment in environmental acquisition. We have extended CLASSICJAVA to provide a formal operational semantics and a formal type system for a language that includes acquisition, thus placing the mechanism on a firm theoretical grounding.

We have also explored several of the design alternatives in this formal context, and we have found that type soundness means that many of these alternatives can only be resolved in certain ways. Soundness alone is not sufficient to resolve all of the design choices, however, so we clearly need more experience with acquisition to make the appropriate choices.

This research suggests three additional projects. First, we need to implement acquisition by modifying an existing class system, so that we can conduct program design experiments in the context of a full programming language. Second, we must search existing class libraries and frameworks for instances of environmental acquisition via pointer chasing or callbacks to collect a wide range of examples for this feature. Third, we must investigate the use of more advanced type systems for acquisition. These include the ability to infer certain types, such as the list of possible containers for a class, and the use of resource-aware type systems to ensure that certain exceptions (most notably the “incomplete context” exception) cannot be generated.

APPENDIX

Lemma 6 (3: Subject Reduction) *If $P, \Gamma \vdash_{\underline{e}} e : t$ and $P, \Gamma \vdash_{\sigma} \mathcal{S}$ and $\text{superOk}(e)$ and $P \vdash \langle e, \mathcal{S} \rangle \leftrightarrow \langle e', \mathcal{S}' \rangle$, then either e' is an error configuration or there exists a type environment Γ' such that*

1. $P, \Gamma' \vdash_{\underline{e}} e' : t$, and
2. $P, \Gamma' \vdash_{\sigma} \mathcal{S}'$, and
3. $\text{superOk}(e')$.

Proof 4 (of lemma 3) Since e reduces, there must exist a context E and a redex e_r such that $e = E[e_r]$. The proof proceeds by structural induction on E ; the base case requires a case analysis on the reduction $P \vdash \langle e_r, \mathcal{S} \rangle \leftrightarrow \langle e'_r, \mathcal{S}' \rangle$. Most of the details are as in the CLASSICJAVA soundness proof; we present only the interesting cases here.

Reduction [new]: $P \vdash \langle \text{new } c, \mathcal{S} \rangle \hookrightarrow \langle \text{obj}, \mathcal{S}' \rangle$, where $\text{obj} \notin \text{dom}(\mathcal{S})$, $\mathcal{S}' = \mathcal{S}[\text{obj} \mapsto \langle c, \text{null}, \mathcal{F} \rangle]$, and $\mathcal{F} = \{fd \mapsto \text{null} \mid \langle _, fd, s \rangle \in_P c \wedge s \neq \text{acquired}\}$. This proof case is important because it establishes most of the store invariants.

Let $\Gamma' = \Gamma[\text{obj} \mapsto c]$. We have $P, \Gamma \vdash_{\mathbf{e}} \text{new } c : c$, and, by construction, $P, \Gamma' \vdash_{\mathbf{e}} \text{obj} : c$.

We must also show that $P, \Gamma' \vdash_{\sigma} \mathcal{S}'$, one clause at a time. We show only the interesting cases.

- (*esc-1*) For any object $\text{obj}' \in \text{dom}(\Gamma')$, either $\text{obj}' = \text{obj}$ or $\text{obj}' \in \text{dom}(\Gamma)$. In either case, $\text{obj}' \in \text{dom}(\mathcal{S}')$.
- (*esc-2*) By construction, $\text{dom}(\mathcal{S}') = \text{dom}(\mathcal{S}) \cup \{\text{obj}\}$, and $\text{dom}(\Gamma') = \text{dom}(\Gamma) \cup \{\text{obj}\}$. Since $\text{dom}(\mathcal{S}) \subseteq \text{dom}(\Gamma)$, $\text{dom}(\mathcal{S}') \subseteq \text{dom}(\Gamma')$ as required.
- (*esc-3*) $G_{\mathcal{S}'}$ adds one node and no edges to $G_{\mathcal{S}}$, which is itself finite and acyclic. Therefore $G_{\mathcal{S}'}$ must also be finite and acyclic.
- (*esc-4*) Consider an object $\text{obj}' \in \text{dom}(\mathcal{S}')$ such that $\text{obj}' \neq \text{obj}$ and $\mathcal{S}'(\text{obj}') = \langle c', p', \mathcal{F}' \rangle$. By [new], $\mathcal{S}(\text{obj}') = \langle c', p', \mathcal{F}' \rangle$, which implies that $\Gamma(\text{obj}') = c'$. Then $\Gamma'(\text{obj}') = c'$, as required. For obj , we have $\mathcal{S}'(\text{obj}) = \langle c, \text{null}, \mathcal{F} \rangle$ and $\Gamma'(\text{obj}) = c$ by construction, so the necessary relationship holds.
- (*esc-5*) For obj , \mathcal{F} has the correct domain by construction. For all other $\text{obj}' \in \text{dom}(\mathcal{S}')$, $\text{obj}' \in \text{dom}(\mathcal{S})$ so each such object's field function has the correct domain in the old store. Since this reduction does not update any object except obj , the field functions retain the correct domains.
- (*esc-7*) This clause remains true for all objects in $\text{dom}(\mathcal{S})$, as this rule does not modify any of their fields. This condition is trivially true for obj , as all of its fields are null.
- (*esc-9*) \implies : let obj' be an object in $\text{dom}(\mathcal{S}')$ such that $\mathcal{S}'(\text{obj}') = \langle _, \text{ctnr}, _ \rangle$ and $\text{ctnr}' \neq \text{null}$; as a result of this condition, $\text{obj}' \neq \text{obj}$. Since this reduction does not change any container references, $\mathcal{S}(\text{obj}') = \mathcal{S}'(\text{obj}')$. By (*esc-13*), $\text{ctnr} \in \text{dom}(\mathcal{S})$, and by (*esc-9*), $P, \mathcal{S} \vdash \text{ctnr} \sqsupseteq \text{obj}'$. Since this reduction does not modify any fields in existing objects, $P, \mathcal{S}' \vdash \text{ctnr} \sqsupseteq \text{obj}'$, as required.
 \Leftarrow : Let obj' be an object in $\text{dom}(\mathcal{S}')$ such that there exists a $\text{ctnr}' \in \text{dom}(\mathcal{S}')$ such that $P, \mathcal{S}' \vdash \text{ctnr}' \sqsupseteq \text{obj}'$. Since all of obj' 's fields are null, $\text{ctnr}' \neq \text{obj}$. Because this reduction does not modify any fields in existing objects, we must have $P, \mathcal{S} \vdash \text{ctnr}' \sqsupseteq \text{obj}'$. Therefore, by (*esc-9*), $\mathcal{S}(\text{obj}') = \langle _, \text{ctnr}', _ \rangle$. Since this reduction does not modify any container references, $\mathcal{S}'(\text{obj}') = \langle _, \text{ctnr}', _ \rangle$, as required.
- (*esc-10*) As before, this clause remains true for all objects in $\text{dom}(\mathcal{S})$ because this reduction rule does not modify any container references. Since obj does not have a container, this condition is trivially satisfied.

Finally, obj has no **super** subexpressions, so *superOk*(obj) holds trivially.

Reduction [aget]: $P \vdash \langle \text{obj}.fd, \mathcal{S} \rangle \hookrightarrow \langle \text{getAcqField}_P(\mathcal{S}, \text{obj}, fd), \mathcal{S} \rangle$, where $\mathcal{S}(\text{obj}) = \langle c, p, \mathcal{F} \rangle$ and $\langle t_{fd}, fd, \text{acquired} \rangle \in_P c$ and $\text{getAcqField}_P(\mathcal{S}, \text{obj}, fd) \neq \text{notFound}$.

Let $v = \text{getAcqField}_P(\mathcal{S}, \text{obj}, fd)$, and let $\Gamma' = \Gamma$.

Since P is well-typed, [DEFN] requires that $\text{canAcq}_P(c, \langle t_{fd}, fd, \text{acquired} \rangle)$ be true. Since

$$\text{getAcqField}_P(\mathcal{S}, \text{obj}, fd) \neq \text{notFound}$$

by assumption, lemma 2 holds, and either $v = \text{null}$ or $\mathcal{S}(v) = \langle c_v, _, _ \rangle$, where $c_v <:_P t$. If $v = \text{null}$, then $P, \Gamma' \vdash_{\mathbf{e}} v : t$ trivially. Otherwise, we have $P, \Gamma' \vdash_{\mathbf{e}} v : t$ by (*esc-4*).

The reduction does not change either Γ or \mathcal{S} , so environment-store consistency is trivially preserved. In addition, the reduction does not perform any substitutions, so *superOk*(e') trivially follows.

Reduction [ct-set]: $P \vdash \langle \text{obj}.fd = v, \mathcal{S} \rangle \hookrightarrow \langle v, \text{updateCtdField}(\mathcal{S}, \text{obj}, fd, v) \rangle$, where $v \in \text{dom}(\mathcal{S})$ and $\langle _, fd, \text{contained} \rangle \in_P \text{classTag}(\mathcal{S}(\text{obj}))$ and $\text{ctnr}(\mathcal{S}(v)) = \text{null}$ and $\neg \text{canReach}_P(\mathcal{S}, \text{obj}, v)$.

Let

$$\begin{aligned} \Gamma' &= \Gamma \\ \mathcal{S}' &= \text{updateCtdField}(\mathcal{S}, \text{obj}, fd, v) \\ \mathcal{F} &= \text{fields}(\mathcal{S}(\text{obj})) \\ \mathcal{F}' &= \text{fields}(\mathcal{S}'(\text{obj})). \end{aligned}$$

In particular, \mathcal{S} and \mathcal{F} are the store and obj 's fields before the reduction, and \mathcal{S}' and \mathcal{F}' are the store and the fields after the reduction. Therefore, $\mathcal{F}(fd)$ is the value of the field before the assignment, and v is its value afterwards. Note that $\mathcal{F}(fd) \neq obj$ by $(esc-3)$, and $v \neq obj$, or $canReach_P(\mathcal{S}, obj, v)$ would be true.

We have $P, \Gamma \vdash_{\underline{e}} obj.f d = v : t$, so $P, \Gamma \vdash_{\underline{e}} v : t$ by [FIELDSET]. By $(esc-4)$, then, $classTag(\mathcal{S}(v)) = c$ where $c <_P t$. Since $updateCtdField$ does not modify any object's class tag, $classTag(\mathcal{S}'(v)) = c$, so we have $P, \Gamma' \vdash_{\underline{e}} v : t$ as required.

Now ensure environment-store consistency. By the definition of $updateCtdField$, $\text{dom}(\mathcal{S}') = \text{dom}(\mathcal{S})$ and $\text{dom}(\mathcal{F}) = \text{dom}(\mathcal{F}')$. (Again, we omit the straightforward cases for space reasons.)

- $(esc-3)$ This reduction creates $G_{\mathcal{S}'}$ from $G_{\mathcal{S}}$ through two operations. First, if $\mathcal{F}(fd) \neq \text{null}$, then it removes the edge from $\mathcal{F}(fd)$ to obj , which cannot create any cycles. Then, it adds an edge from v to obj . Since $G_{\mathcal{S}}$ has no cycles, this new edge can create a cycle only when there is a path from obj to v in $G_{\mathcal{S}}$. But since $canReach_P(\mathcal{S}, obj, v)$ is false, this is impossible, and therefore $G_{\mathcal{S}'}$ has no cycles. Finally, $G_{\mathcal{S}'}$ and $G_{\mathcal{S}}$ have the same number of nodes, so $G_{\mathcal{S}'}$ is finite.
- $(esc-6)$ This reduction only modifies the fields of obj ; this property is preserved for all other objects. By construction, $\text{rng}(\mathcal{F}') \subseteq \text{rng}(\mathcal{F}) \cup \{v\}$. Since $v \in \text{dom}(\mathcal{S})$, $\text{rng}(\mathcal{F}) \cup \{v\} \subseteq \text{dom}(\mathcal{S}) \cup \{\text{null}\}$. Since $\text{dom}(\mathcal{S}') = \text{dom}(\mathcal{S})$, we have $\text{rng}(\mathcal{F}') \subseteq \text{dom}(\mathcal{S}') \cup \{\text{null}\}$ as required.
- $(esc-7)$ The only interesting case is the field fd within obj ; no other fields are modified. From $P, \Gamma \vdash_{\underline{e}} obj.f d = v : t$ and [FIELDSET], we have $\langle t, fd, _ \rangle \in_P c$, where $\Gamma(obj) = c$ and $\mathcal{S}(obj) = \langle c, p, \mathcal{F} \rangle$. Further, $P, \Gamma \vdash_{\underline{e}} v : t$. This satisfies the clause.
- $(esc-8)$ This clause holds true for all objects other than obj and for all fields of obj other than fd . For this case, assume that there exists a field $fd' \neq fd$ such that $\mathcal{F}(fd') = v$, which implies that $\mathcal{F}(fd') = v$ as well. Then $P, \mathcal{S} \vdash obj \supseteq v$ and by $(esc-12)$, $ctnr(\mathcal{S}(v)) = obj$. But this violates one of the conditions for this reduction, so this clause's antecedents cannot be true and the clause is thus trivially satisfied.
- $(esc-9) \implies$: Let obj' be an object in $\text{dom}(\mathcal{S}')$ such that $\mathcal{S}'(obj') = \langle _, ctnr, _ \rangle$ and $ctnr \neq \text{null}$. Case analysis on obj' :

- $obj' = \mathcal{F}(fd)$. Cannot happen because $\mathcal{S}'(\mathcal{F}(fd)) = \langle _, \text{null}, _ \rangle$.
- $obj' = obj$, the object whose field is updated. Let $\langle _, p, _ \rangle = \mathcal{S}'(obj')$. Because $G_{\mathcal{S}}$ is acyclic, $obj' \neq \mathcal{F}(fd)$, and because $\neg canReach_P(\mathcal{S}, obj', v)$, $obj' \neq v$. Therefore obj' is not one of the two objects whose containers are modified by this reduction, and $ctnr(\mathcal{S}'(obj')) = ctnr(\mathcal{S}(obj')) = p$. Therefore, we have $P, \mathcal{S} \vdash p \supseteq obj'$. Now, since $p \neq obj'$ (again because $G_{\mathcal{S}}$ is acyclic), this reduction does not modify any of p 's fields, so we must have $P, \mathcal{S}' \vdash p \supseteq obj'$ as required.
- $obj' = v$, the new value of the field. By [ct-set], $ctnr = obj$ and $\mathcal{F}'(fd) = v$, so $P, \mathcal{S}' \vdash obj \supseteq v$ and thus $P, \mathcal{S}' \vdash obj \supseteq obj'$ as required.
- obj' is none of $\mathcal{F}(fd)$, obj , or v . Since this reduction only changes the container references of $\mathcal{F}(fd)$ and v , we must have $ctnr(\mathcal{S}(obj')) = ctnr(\mathcal{S}'(obj'))$. Therefore, $P, \mathcal{S} \vdash ctnr \supseteq obj'$. If $ctnr \neq obj$, then $ctnr$'s fields are not modified, so $P, \mathcal{S}' \vdash ctnr \supseteq obj'$. Otherwise, if $ctnr = obj$, then there must exist some field $fd' \neq fd$ such that $\mathcal{F}(fd') = obj'$, where fd' is a contained field. Since this reduction only changes the field fd , $\mathcal{F}'(fd') = obj'$, and thus $P, \mathcal{S}' \vdash ctnr \supseteq obj'$ as required.

\Leftarrow : Let obj_1 and obj_2 be arbitrary objects in $\text{dom}(\mathcal{S}')$ such that $P, \mathcal{S}' \vdash obj_1 \supseteq obj_2$. Case analysis on obj_1 :

- $obj_1 = \mathcal{F}(fd)$. Then $obj_1 \neq obj$, or $G_{\mathcal{S}}$ would have a cycle. Since this reduction does not modify obj_1 's fields, $P, \mathcal{S} \vdash obj_1 \supseteq obj_2$ and therefore $ctnr(\mathcal{S}(obj_2)) = obj_1$. Further, $obj_2 \neq \mathcal{F}(fd)$; otherwise (because this reduction does not modify $\mathcal{F}(fd)$'s fields) $G_{\mathcal{S}}$ would have had a cycle. Similarly, $obj_2 \neq v$; otherwise we'd have $ctnr(\mathcal{S}(v)) = obj_1$, but we know that $ctnr(\mathcal{S}(v)) = \text{null}$. Therefore, this reduction does not modify obj_2 's container pointer, and $ctnr(\mathcal{S}'(obj_2)) = obj_1$, as required.

- $obj_1 = obj$. By this reduction and (esc-8), we know that $\neg(P, \mathcal{S}' \vdash obj \sqsupseteq \mathcal{F}(fd))$, so $obj_2 \neq \mathcal{F}(fd)$. Further, $obj_2 \neq obj$; otherwise $G_{\mathcal{S}'}$ would have a cycle, and we previously proved that impossible. If $obj_2 = v$, then this reduction ensures that $ctnr(\mathcal{S}'(v)) = obj$, as required. Otherwise, since we do not modify any of obj 's other fields, we must have $P, \mathcal{S} \vdash obj_1 \sqsupseteq obj_2$, which implies that $ctnr(\mathcal{S}(obj_2)) = obj_1$. Since this reduction does not modify obj_2 's parent pointer, $ctnr(\mathcal{S}'(obj_2)) = obj_1$, as required.
- $obj_1 = v$. This reduction does not modify obj_1 's fields, so we must have $P, \mathcal{S} \vdash obj_1 \sqsupseteq obj_2$, and therefore that $ctnr(\mathcal{S}(obj_2)) = obj_1$. Now, because $P, \mathcal{S} \vdash obj \sqsupseteq \mathcal{F}(fd)$, we know that $ctnr(\mathcal{S}(\mathcal{F}(fd))) = obj$. Since $obj_1 \neq obj$, $obj_2 \neq \mathcal{F}(fd)$. And $obj_2 \neq v$, or $G_{\mathcal{S}}$ would have a cycle. Therefore, this reduction does not change obj_2 's parent pointer, so $ctnr(\mathcal{S}'(obj_2)) = obj_1$ as required.
- Finally, $obj_1 \notin \{\mathcal{F}(fd), obj, v\}$. Since this reduction does not modify obj_1 's fields, we have $P, \mathcal{S} \vdash obj_1 \sqsupseteq obj_2$ and therefore $ctnr(\mathcal{S}(obj_2)) = obj_1$. Because $obj_1 \neq obj$, $obj_2 \neq \mathcal{F}(fd)$. And because $ctnr(\mathcal{S}(v)) = \text{null}$, $obj_2 \neq v$. Therefore, this reduction does not modify obj_2 's parent pointer, and $ctnr(\mathcal{S}'(obj_2)) = obj_1$ as required.

(esc-10) This reduction changes the parent pointers of exactly two objects, $\mathcal{F}(fd)$ and v ; the property is trivially preserved for all others. By the definition of the reduction, $ctnr(\mathcal{S}'(\mathcal{F}(fd))) = \text{null}$, so the requirement for the first object is trivially satisfied. For the second object, $ctnr(\mathcal{S}'(v)) = obj$. Let $c_1 = \text{classTag}(\mathcal{S}'(obj))$ and $c_2 = \text{classTag}(\mathcal{S}'(v))$. By [ct-set], we know that $\langle t, fd, \text{contained} \rangle \in_P c_1$; by [FIELDSET], $c_2 \leq_P t$. By [DEFN], $t \sqsupseteq_P c_1$, and also $\text{containers}_P(c_2) \supseteq \text{containers}_P(t)$. Therefore, $c_2 \sqsupseteq_P c_1$ as required.

(esc-11) The only objects whose parent pointers are modified by this reduction are v and $\mathcal{F}(fd)$. By construction, we have $ctnr(\mathcal{S}'(v)) = obj \in \text{dom}(\mathcal{S}')$, and $ctnr(\mathcal{S}'(\mathcal{F}(fd))) = \text{null}$.

Therefore, $P, \Gamma' \vdash_{\sigma} \mathcal{S}'$.

Finally, since the result of this reduction is the value v , it cannot contain any super expressions, and $\text{superOk}(v)$ is true.

Reduction [acall]: We have

$$\begin{aligned}
P \vdash \langle obj.md(v_1 \dots v_n), \mathcal{S} \rangle &\hookrightarrow \langle e_b[obj'/\text{this}, v_1/var_1, \dots, v_n/var_n], \mathcal{S} \rangle, \\
&\text{where } \langle md, (t_1 \dots t_n \rightarrow t_{md}), _, e, \text{acquired} \rangle \in_P \text{classTag}(\mathcal{S}(obj)) \\
&\text{and } \langle obj', (var_1, \dots, var_n), e_b \rangle = \text{getAcqMethod}_P(\mathcal{S}, obj, md).
\end{aligned}$$

Take $\Gamma' = \Gamma$. Since P is well-typed, [DEFN] requires that

$$\text{canAcqMeth}_P(c, \langle md, (t_1 \dots t_n \rightarrow t_{md}), V, e, \text{acquired} \rangle).$$

By [acall], $\text{getAcqMethod}_P(\mathcal{S}, obj, md) = \langle obj', (var_1, \dots, var_n), e_b \rangle$. Therefore lemma 2 applies, so

- $\mathcal{S}(obj') = \langle c_v, _, _ \rangle$, and
- $\langle md, (t'_1 \dots t'_n \rightarrow t'), (var_1, \dots, var_n), e_b, _ \rangle \in_P c_v$, and
- $(t'_1 \dots t'_n \rightarrow t') <:_P (t_1 \dots t_n \rightarrow t_{md})$.

By [CALL], $P, \Gamma \vdash_{\underline{e}} obj.md(v_1 \dots v_n) : t$ implies that $t = t_{md}$ and that

$$P, \Gamma \vdash_{\underline{e}} v_i : t_i \text{ for } i \in [1, n].$$

By the definition of subtyping, this implies that $P, \Gamma \vdash_{\underline{e}} v_i : t'_i$ for all $i \in [1, n]$. Now, by [DEFN] and [METHOD], $\langle md, (t'_1 \dots t'_n \rightarrow t), (var_1, \dots, var_n), e_b, _ \rangle \in_P c_v$ requires that

$$P, \Gamma[\text{this} \mapsto c_v, var_1 \mapsto t'_1, \dots, var_n \mapsto t'_n] \vdash_{\underline{e}} e_b : t'.$$

Therefore, by a substitution lemma,⁶

$$P, \Gamma \vdash_{\underline{e}} e_b[obj'/this, v_1/var_1, \dots, v_n/var_n] : t',$$

and by subtyping, $P, \Gamma \vdash_{\underline{e}} e_b[obj'/this, v_1/var_1, \dots, v_n/var_n] : t_{md}$ as required.

This reduction does not change either Γ or \mathcal{S} , so environment-store consistency is trivially preserved. It may, however, introduce new **super** expressions, though all of these introduced expressions come from e_b , and all of these expressions were annotated by the elaboration phase. Therefore, all of the **super** expressions in e_b are annotated with **this**. The only part of the substitution that can have any effect on these expressions is $e_b[obj'/this]$, which replaces the annotation with obj' . Therefore *superOk* holds. ■

2 Acknowledgments

We thank David Lorenz for taking the time to discuss his prior work with us. Mitch Wand provided invaluable assistance with JACQUES's soundness proof. Erik Ernst pointed out an important omission in the earlier, informal versions of *canAcq_P* and *canAcqMeth_P* [2]. David Richter suggested the possibility of using type systems for non-nullable pointers, as described in section 5.7. Finally, we thank the POPL 2005 anonymous reviewers for their comments, insights, and suggestions.

References

- [1] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *Proceedings of the 13th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 48–64. ACM Press, 1998.
- [2] Richard Cobbe and Matthias Felleisen. Environmental acquisition revisited. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 14–25, New York, NY, USA, 2005. ACM Press.
- [3] ECMA. Standard ECMA-334: C# language specification. Available on the web as <http://www.ecma-international.org/publications/files/ecma-st/ECMA-334.pdf>, December 2002.
- [4] Manuel Fähndrich and K. Rustan M. Leino. Declaring and checking non-null types in an object-oriented language. In *OOPSLA '03: Proceedings of the 18th annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 302–312. ACM Press, 2003.
- [5] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. A programmer's reduction semantics for classes and mixins. In *Formal Syntax and Semantics of Java*, volume 1523 of *Springer Lecture Notes in Computer Science*, pages 241–269. Springer-Verlag, 1999. Preliminary version appeared in *Proceedings of Principles of Programming Languages*, 1998.
- [6] Martin Fowler and Kendall Scott. *UML Distilled: Applying the Standard Object Modeling Language*. Addison-Wesley, Reading, MA, 1997.
- [7] Jim Fulton. Acquisition. <http://olympus.het.brown.edu/doc/python-extclass/Acquisition.html>, 1998.
- [8] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, MA, 1995.
- [9] Joseph Gil and David H. Lorenz. Environmental acquisition: a new inheritance-like abstraction mechanism. In *Proceedings of the Eleventh Annual Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 214–231. ACM Press, 1996.
- [10] Joseph (Yossi) Gil and David H. Lorenz. Object technology: Design patterns and language design. *IEEE Computer*, 31(3):118–120, March 1998.

⁶The substitution lemma takes the standard form; it is omitted here for space reasons.

- [11] James Gosling, Bill Joy, and Guy Steele, Jr. *The Java(TM) Language Specification*. Addison-Wesley, 1996.
- [12] Kathryn E. Gray and Matthew Flatt. ProfessorJ: A gradual intro to Java through language levels. In *OOPSLA Educators' Symposium*, 2003.
- [13] Alan C. Kay. The early history of Smalltalk. In *The Second ACM SIGPLAN Conference on History of Programming Languages*, pages 69–95. ACM Press, 1993.
- [14] Shriram Krishnamurthi, Yan-David Erlich, and Matthias Felleisen. Expressing structural properties as language constructs. In *European Symposium on Programming*, volume 1576 of *Springer Lecture Notes in Computer Science*, pages 258–272, March 1999.
- [15] Amos Latteier, Michel Pelletier, Chris McDonough, and Peter Sabaini. *The Zope Book*. SAMS, 2001. Also available on-line at <http://zope.org/Documentation/Books/ZopeBook>.
- [16] Mark Logan, Matthias Felleisen, and David Blank-Edelman. Environmental acquisition in network management. In *Proceedings of LISA 2002: Sixteenth Systems Administration Conference*, pages 175–184. USENIX Association, 2002.
- [17] Bertrand Meyer. *Eiffel: The Language*. Prentice Hall, 1992.
- [18] Sun Microsystems. Java foundation classes (JFC/Swing). <http://java.sun.com/products/jfc/index.jsp>, 1998.
- [19] Guido van Rossum. Python. <http://www.python.org/>, 1991.
- [20] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994. First appeared as Technical Report TR160, Rice University, 1991.