Ownership Confinement Ensures Representation Independence for Object-Oriented Programs

June 6, 2005

ANINDYA BANERJEE
Department of Computing and Information Sciences
Kansas State University
Manhattan KS 66506 USA
and
DAVID A. NAUMANN
Department of Computer Science
Stevens Institute of Technology
Hoboken NJ 07030 USA

Dedicated to the memory of Edsger W. Dijkstra.

Representation independence formally characterizes the encapsulation provided by language constructs for data abstraction and justifies reasoning by simulation. Representation independence has been shown for a variety of languages and constructs but not for shared references to mutable state; indeed it fails in general for such languages. This paper formulates representation independence for classes, in an imperative, object-oriented language with pointers, subclassing and dynamic dispatch, class oriented visibility control, recursive types and methods, and a simple form of module. An instance of a class is considered to implement an abstraction using private fields and so-called representation objects. Encapsulation of representation objects is expressed by a restriction, called confinement, on aliasing. Representation independence is proved for programs satisfying the confinement condition. A static analysis is given for confinement that accepts common designs such as the observer and factory patterns. The formalization takes into account not only the usual interface between a client and a class that provides an abstraction but also the interface (often called "protected") between the class and its subclasses.

Categories and Subject Descriptors: D.3.3 [**Software**]: Programming Languages—*Language Constructs and Features*; F.3.1 [**Theory of Computation**]: Logics and Meanings of Programs—*Specifying and Verifying and Reasoning about Programs*

General Terms: Languages, Verification

Additional Key Words and Phrases: Alias control, confinement, relational parametricity, simulation, data refinement

---

This is an expanded and revised version of a paper originally appearing in *ACM Symposium on Principles of Programming Languages*, 2002.
Banerjee was supported by NSF grants EIA-9806835, CCR-0209205, ITR-0326577 and NSF Career award CCR-0093080/CCR-0296182.
Naumann was supported by NSF grants INT-9813854, CCR-0208984, CCF-0429894, and a grant from the New Jersey Commission on Science and Technology.
Banerjee and Naumann were also supported by EPSRC grant GR/S03539, "Abstraction, Confinement and Heap Storage", for which we thank the principal investigator, Peter O'Hearn.
Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.
© 20YY ACM 0004-5411/20YY/0100-0001 $5.00

1. INTRODUCTION

You have implemented a class [Dahl and Nygaard 1966; Arnold and Gosling 1998], FIFO, whose instances are FIFO queues with public methods enqueue and dequeue as well as method size that reports the number of elements in the queue. The class, implemented in some Java-like object-oriented language, is part of a library and is used by many programs, most unknown to you. The queue is *represented* using a singly linked chain of nodes that point to elements of the queue. There is also a sentinel node [Cormen et al. 1990]. Each instance of FIFO has a field num with the number of nodes and a field snt that references the sentinel. You realize that a simpler, more efficient implementation can be provided without the sentinel, using two fields, head and tail, pointing to the end nodes in the chain. You revise method size to return num instead of num-1 and revise the other methods suitably. You are guided to the necessary revisions by thinking about the correspondence, sometimes called a *simulation relation*, between the representations for the two versions.

Can the revisions affect the behavior of clients, that is, programs that use class FIFO in some way or other? The answer would be yes, if some client determined the number of nodes by reading field num directly. A client that refers to field name snt would no longer compile. But you have taken care to *encapsulate* the queue's representation: the fields are declared to be private. By using programming language constructs like private fields you aim to ensure that client programs depend only on the *abstraction* provided by the class, not on its representation. If client behavior is independent from the representation of FIFO, it is enough for you to ensure equivalent visible behavior of the revised methods.

For scalable systems, scalable system-building tools, and scalable development methods, abstraction is essential. For reasoning about a single component, e.g., a class, module, or local block, abstraction makes it possible to consider other components in terms of their behavioral interface rather than their internal representation. Abstraction is needed for the automated reasoning embodied in static analysis tools [Cousot and Cousot 1977] and it is needed for formal and informal reasoning about functional correctness during development and evolution [Milner 1971; Hoare 1972]. Modular reasoning has always been a central issue in software engineering and in static analysis. With the ascendancy of mobile code it has become absolutely essential. For example, it is possible for clients of FIFO to be linked to it only at runtime, so it is impossible to check all uses to determine whether the revisions affect them.

The need for flexible but robust encapsulation mechanisms to support data abstraction has been one of the driving forces in the evolution of programming language design, from type safety and scoped local variables to module and abstract data type constructs [Liskov and Guttag 1986]. There is a rich theoretical literature on the subject (e.g., [Plotkin 1973; Reynolds 1974; Donahue 1979; Haynes 1984; Reynolds 1984; He et al. 1986; Mitchell 1996; Lynch and Vaandrager 1995; de Roever and Engelhardt 1998]). Many different language constructs have been studied. There is considerable variation in the details of these theories, partly because the intended applications vary from justifying general tools for program analysis and transformation to justifying proof rules to be applied to specific programs as in the FIFO example. The common thread is that two implementations of

PSfrag replacements

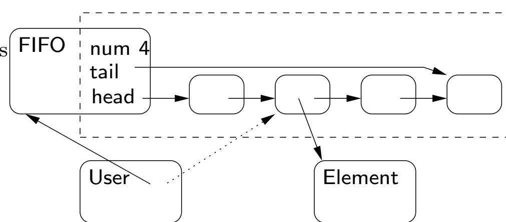

Fig. 1. A FIFO object with its encapsulated representation: private fields and nodes of a list (within the dashed rectangle). One element of the queue is shown as well as a user of the queue, but other objects and references are omitted. The dotted reference is an example of representation exposure.

a component are linked by a simulation relation between the two representations.

Unfortunately, these theories are inadequate for object-oriented programs. They deal well with the encapsulation of data structures that correspond directly to some language construct, such as modules, local variables, or private fields. But the FIFO example also involves encapsulation of a data structure composed of heap cells and pointers, including aliasing with the tail field as depicted in Fig. 1.

The problem is that encapsulation provided by language constructs often runs afoul of aliasing. For variables and parameters, aliasing can be prevented through syntactic restrictions that are tolerable in practice (and often assumed in formal logics and theories). Aliasing via pointers is an unavoidable problem in object oriented programming where shared mutable objects are pervasive. Yet unintended aliasing can be catastrophic. A version of the Java access control system was rendered insecure because a leaked reference to an internal data structure made it possible to forge crytographic authentication [Vitek and Bokowski 2001]. In simply typed languages, types offer limited help: variables x, y are not aliased if they have different types. Even this help is undercut by subclass polymorphism: in Java, a variable x of type **Object** can alias y of any type.

The ubiquity and practical significance of the issue is articulated well in the manifesto of Hogg et al. [1992]. A number of subsequent papers in the object-oriented programming literature propose disciplines to control aliasing. Of particular relevance are disciplines that impose some form of *ownership confinement* that restricts access to designated "representation objects" (*reps* for short) except via their "owners", to prevent *representation exposure* [Leino and Nelson 2002]. A good survey on confinement, especially ownership, can be found in the dissertation of Clarke [2001]; see also Lea [2000],Vitek and Bokowski [2001], Clarke et al. [2001], Müller and Poetzsch-Heffter [2000b], Boyland [2001], Aldrich et al. [2002], and Sect. 12.1.

The contribution of this paper is a theory of representation independence for encapsulation of data in the heap, using ownership confinement. We follow Reynolds [1984] in calling our main result an *abstraction theorem.* Some readers may prefer the term *relational parametricity.*

The literature on confinement is largely concerned with static or dynamic checks to ensure invariance of various confinement properties. One of our contributions is to show how established semantic techniques can be used to evaluate confinement disciplines. To prove our abstraction theorem, we use a semantic formulation of confinement. To show that the notion is not too restrictive or too difficult to enforce, we give a modular, syntax-directed static analysis for confinement. It accepts some example programs that embody important object-oriented design patterns. We do not claim, however, that it is definitive or comprehensive.

There are a number of ways in which abstractions can be expressed using constructs of contemporary object-oriented languages, including modules, classes, local variables, object instances, not to mention heap structures such as object groups. We treat the most common situation: an instance of some class is viewed as representing an abstraction, possibly using some other objects as part of its representation.

We are aware of no previous results on representation independence that address encapsulation of objects in the heap. Thus it is tempting to present the ideas in the setting of a simple idealized language, say a simple imperative language with pointers to mutable heap cells. But this would leave open some challenging issues, such as how class-based scoping rules fit with instance-based abstraction. We have chosen to consider a rich imperative object-oriented language with class-based visibility, inheritance and dynamic binding, type casts and tests, recursive types, and other features sufficient to model many programs that fit common design patterns such as observer and factory [Gamma et al. 1995].

Previous work on representation independence has been concerned with relating two versions of a component with respect to programs that use the component. For a class, the designer of a class needs to consider not only users (the *client interface*) but also subclasses (the *protected interface*). This is a source of complication in our treatment of confinement and, to a lesser extent, in our treatment of representation independence. Our results consider replacement of one version of a class (the "owner") by another with the same public interface, in the context of arbitrary classes that use it and/or are subclasses of it. On one hand, we cannot allow direct access to the owner's fields from subclass code, i.e., "protected" visibility. This would entail, in general, coordinated revision of multiple classes together, whereas for simplicity we consider revision of a single class. On the other hand, we do treat access to reps from subclass code.

Overview and readmap. The organization of the paper is intended to make it possible for readers with different interests to skip some parts and still get the gist of the results.

Sect. 2 introduces the language for which our results are proved and describes a simple example with which we review the formalization of representation independence using simulation relations. The example is extended to one showing how representation independence can be invalidated by leaked references to reps. The section concludes with an informal statement of our abstraction theorem.

Sect. 3 discusses more elaborate examples that typify object-oriented programs. A version of a Meyer-Sieber [1988] example shows how higher order programs can be expressed. Versions of the observer pattern [Gamma et al. 1995] illustrate challenges in formulating robust but practical notions of confinement. The section concludes with an informal description of our notion of ownership confinement.

Sect. 4 formalizes the syntax and typing rules. Sect. 5 gives a surprisingly simple denotational semantics in the manner of Strachey [2000]. The reader is expected to be familiar with elementary domain theory and fixpoints but nothing beyond what is found in introductory textbooks [Davey and Priestley 1990; Winskel 1993].

Confinement, the semantic notion, is defined formally in Sect. 6. Sect. 7 gives the first main result, an abstraction theorem for confined programs. Sect. 8 shows

in detail how the theorem applies to the examples in Sect. 3. Sect. 9 considers examples of the interface between an owner class and its subclasses. We formalize the situation where the owner is declared in a module that contains some other classes, in order to give a notion of confinement that takes into account the subclasses of an owner —while avoiding the complexity of formalizing a module construct or protected scope in the full generality of, say, Java. Sect. 10 proves a second abstraction theorem, for this extended language and for a generalized notion of simulation needed for owner subclasses. Sect. 11 wraps up the technical development by defining a static analysis for confinement that accepts the examples of Sections 2, 3, 8, and 9; soundness with respect to (semantic) confinement is shown. Sect. 12 discusses related work and open challenges.

The companion technical report [Banerjee and Naumann 2004a] gives a number of additional examples, proofs, and proof cases. Several examples use constructors and calls to super class methods; these are explained in the paper, but are formalized only in the technical report.

Differences from the preliminary version. Outgoing references, from representation objects to client objects, were disallowed in the preliminary version of this paper [Banerjee and Naumann 2002]. We conjectured that they could be allowed if restricted to read-only access as in [Müller and Poetzsch-Heffter 2000b; Leino and Nelson 2002]. Here we allow them without restriction, as is needed to handle examples such as the observer pattern where observers may well change state in response to events. The other major additions are as follows: module-scoped methods, the generalized abstraction theorem, extensive worked examples, and the static analysis for confinement.

In [Banerjee and Naumann 2002] we discuss simulation proofs of the equivalence of “security passing style” [Wallach et al. 2000] with the lazy “stack inspection” implementation of Java’s privilege-based access control mechanism [Gong 1999], and then extend our language to include access control. We give an abstraction theorem for this extended language. It was this study that led us to the main results but in retrospect it seems tangential and is omitted.

2. REPRESENTATION INDEPENDENCE

We begin this expository section with a very simple example of representation independence, contrived mainly to introduce the Java-like language that we will use. Building on this example we show how pointer aliasing can invalidate representation independence. We conclude with an informal statement of the main results. Sect. 3 deals with more challenging examples.

2.1 A first example

The concrete syntax for classes is based on that of Java [Arnold and Gosling 1998] but using more conventional notation for simple imperative constructs. Keywords are typeset in bold font and comments are preceded by double slash. A program consists of a collection of class declarations like the following one.

```
class Bool extends Object {
  bool f;                              // private field
  constr { skip }                      // public constructor
  unit set(bool x){ self.f := x }      // public method
  bool get(){ result := self.f }       // public method
}
```

There are two associated methods: set takes a boolean parameter and returns nothing; get takes no parameter and returns a boolean value. Methods are considered to be public, that is, visible to methods in all classes. (Module-scoped methods are added in Sect. 10.) Every method has a return type; the primitive type **unit**, with only a single value (*it*), corresponds to Java's "void" and is used for methods like set that are called only for their effect on state.

Instances of class Bool have a field f of (primitive) type bool. A field f is accessed in an expression of the form e.f, and in particular self.f is used for fields of the current object; a bare identifier like x is either a parameter or a local variable. The distinguished variable result provides the return value; it is initialized with the default for its type (*false* for **bool** and *nil* for class types). Fields are considered to be private, that is, visible only to the methods declared in the class. Visibility is class-based, as in many mainstream object-oriented languages: an object can directly access the private fields of another object of the same class.

When a new object is constructed, each field is initialized with the default value for its type. Then the constructor commands are executed: the constructors declared in superclasses are executed before the declared one which, is designated by keyword **constr**. In subsequent examples we omit the constructor if it is **skip** and we refrain from considering constructors with parameters.

The observable behavior of a Bool object can be achieved using an alternate implementation in which the logical complement is stored in a field:

```
class Bool extends Object {
  bool f;
  constr { self.f := true }
  unit set(bool x){ self.f := ¬x }
  bool get(){ result := ¬(self.f) } }
```

We do not formalize class types ("interfaces" in Java) separately from class declarations. Class names are used as types and we use the term *class* loosely to mean the name of a declared class. But we are concerned with relating comparable versions of a class: as in the example above, a comparable version has the same name and methods with the same names and signatures.

We claim that no client program using Bool can distinguish one implementation from the other; thus we are free to replace one by the other. Of course this is not the case if we consider aspects of client behavior such as real time or the size of object code —but these are not at the level of abstraction of source code. Moreover, input and output for end users is of some limited types like **int** or String. If a Bool could be output directly, say displayed in binary on the screen, then an end user could distinguish between the implementations. So we consider only clients that use Bool objects in temporary data structures and not as input or output data.

An example of such a client is method main in the following class. It declares a local variable b of type Bool, with scope beginning at the keyword **in**. In the absence of explicit braces, the scope of a local variable extends to the end of the method body.

```
class Main extends Object {
  String inout;
  unit main(){ Bool b := new Bool in
                    if ...self.inout... then b.set(true) else b.set(false) fi;
                    self.inout := convertToString(b.get()) } }
```

We may consider method main as a main program for which the observable state consists of field inout. Its final value depends on some condition "...self.inout..." on its initial value. No object of type Bool is reachable in the state of a Main object after invocation of main, so there is no observable difference between its behavior using one implementation of Bool and its behavior using the other.

The claim is that we need not consider specific clients; there is no use of Bool that can distinguish between the two implementations. The standard reasoning goes as follows.

(1) Suppose o is an object of type Bool for the first implementation and o′ an object for the second. The correspondence between their states is described by the following *local coupling relation*: $o.f = \neg(o'.f)$.

(2) This relation has the *simulation property*:
—it holds initially (once the constructor has been executed), and
—if the two versions of set (respectively, get) are executed from related states then the outcomes are related. (As we consider sequential programs, the outcome is the updated heap and the return value if any.)
In short, the relation is *established* by the constructor and *preserved* by the methods of Bool.

(3) To consider client programs we must consider program states consisting of local variables (and parameters) along with the heap, which may contain many instances of Bool as well as other objects. For states, we define the *induced coupling relation* for a given local coupling. Primitive values and locations are related by equality (later we refine this to a bijection, to account for differences in allocation.) A pair of heaps are related if there is a one-to-one correspondence between Bool objects such that they are pairwise related by the local coupling of (1), and everything else is related by equality.
The induced coupling relation is preserved by all commands in methods of all classes. This is the *abstraction theorem*.

(4) For a pair of states related by the (induced) coupling, if no Bool objects are reachable then the states are equal. This is the *identity extension lemma*, which follows from the definition of the induced coupling. Identity extension confirms that the chosen notion of coupling relation is suited to the chosen form of encapsulation.

It is a consequence of (3) and (4) that the two implementations cannot be distinguished by a client that does not input or output Bool objects. Any initial state

for such a client is related to itself, by (4). We can consider an execution of the client using either of the two implementations of Bool; the final states are related, according to (3). And thus they are equal, by (4).

For program refinement, identity can be replaced by inequality in step (4). In this paper we do not emphasize refinement, but the requisite adaptation of our results is straightforward. For applications in program analysis, other relations are used in step (4), e.g., for secure information flow the relation expresses equivalence from the point of low-security observers [Volpano et al. 1996].

The abstraction theorem is a non-trivial property of the language. It would fail, for example, if the language had constructs that allowed client programs to read the private fields of Bool —or to enumerate the names of the private fields, or to query the number of boolean fields that are currently true. In fact, similar facilities can be found in some reflection libraries and in the implementation of Java's inner classes, but are considered to be flaws [Bhowmik and Pugh 1999].

Familiar operations on pointers, however, can also violate abstraction. For example, with pointer arithmetic one can distinguish between two representations that differ only in the size of storage used (e.g., representing a boolean value using one bit of an integer versus one bit of a character). Even in the absence of pointer arithmetic, shared references lead to the following problem.

2.2 Representation exposure

Consider the following class OBool which provides functionality similar to that of Bool, in fact using Bool. For clarity we have chosen different method names, to emphasize that we are not comparing this class with Bool.

```
class OBool extends Object {
  Bool g;
  constr { self.g := new Bool; self.g.set(true) }
  unit setg(bool x){ self.g.set(x) }
  bool getg(){ result := self.g.get() } }
```

Here is an alternate implementation of OBool.

```
class OBool extends Object {
  Bool g;
  constr { self.g := new Bool; self.g.set(false) }
  unit setg(bool x){ self.g.set(¬ x) }
  bool getg(){ result := ¬(self.g.get()) } }
```

To describe the connection between the two implementations a suitable local coupling (recall (1) in Sect. 2.1) is the following relation between an object state o for the first implementation of OBool and o′ for the alternate one:

$$\mathsf{o.g} \neq \mathit{nil} \neq \mathsf{o'.g} \ \wedge\ \mathsf{o.g.f} = \neg(\mathsf{o'.g.f}) \quad . \qquad (*)$$

Invocations of setg and getg preserve this relation. For these implementations, it is not just a private field that is to be encapsulated, but also the object referenced by that field. This is apparent in the coupling (∗) which involves both.

To describe the roles of the objects involved, we call class OBool an *owner* class. Its instances "own" objects of class Bool, their representation objects, which are

called *reps* for short. Together, an owner and its reps constitute what we call an *island*, following Hogg [1991] (cf. Fig. 1).

Here is a suitable client for OBool.

```
class Main extends Object {
  String inout;
  unit main(){ OBool z := new OBool in
               if ...self.inout... then z.setg(true) else z.setg(false) fi;
               self.inout := convertToString(z.getg()) } }
```

This does not distinguish between the two implementations of OBool nor does it violate the intended encapsulation boundary. But suppose we add to both versions of OBool the following method which "leaks" a reference to the rep object.

```
Bool bad(){ result := self.g }
```

The method gives its caller an alias to the object pointed to by the private field g. This makes the location of the encapsulated object visible to clients. In and of itself, access to this location is not harmful.[1] Like the other methods, method bad preserves $(*)$. But a client class C can exploit the leak as in the following command.

```
OBool z := new OBool in
Bool w := z.bad() in if w.get() then skip else abort fi
```

The command aborts if the new OBool is an object o' for the second implementation of OBool, but it does not abort for an object o for the first implementation. An attempt to argue using the steps in Sect. 2.1 breaks down because this difference in behavior violates the abstraction theorem, step (3): the induced coupling is not preserved by the above code.

Identity extension, step (4), also fails. For a related pair o,o' of OBool objects we have $o.g = o'.g$ but $o.g.f \neq o'.g.f$. Thus the relation is not the identity for the client to which, owing to method bad, the reps are visible.

The client in the example above does happen to preserve the relation $(*)$, up to the point where the program does or does not diverge, because it does not alter the state of the objects it accesses. For an example where the abstraction theorem, step (3), fails with terminating computations, consider the following client command.

```
OBool z := new OBool in Bool w := z.bad() in w.set(true)
```

This does not preserve $(*)$. To see why, suppose o, o' are a related pair of OBool objects assigned to z and satisfying $(*)$. After the assignment to w, the effect of w.set(true) is to make $o.g.f = o'.g.f$ (both sides are true), contrary to the relation $(*)$ which requires $o.g.f = \neg(o'.g.f)$. This is very different from the effect of z.setg(true).

The examples show that both ingredients of representation independence — identity extension and preservation— can fail if a rep is leaked. The challenge is to confine pointers in a way that disallows harmful leaks and thus admits a robust

[1] To make this clear, one could assume that, for both versions of OBool, the Bool object is allocated at the same location. The assumption can be formalized by adding $o.g = o'.g$ to the local coupling $(*)$. Another justification is given in Sect. 10 where we show formally how the language is "parametric in locations".

representation independence property —without imposing impractical restrictions. The challenge is made more difficult by various features of Java-like languages, for example, type casts. We consider casts now; other challenges are deferred to Sect. 3.

Suppose we change the return type for method bad, attempting to hide the type of the rep object.

Object bad(){ result := self.g }

Class **Object** is the root of the subclassing hierarchy so by subsumption it allows references to objects of any class. The client can use a cast, "(Bool)", to assert that the result of z.bad() has type Bool. (In a state where the assertion is false, the cast would cause abortion.)

OBool z := **new** OBool **in**
Bool w := (Bool)(z.bad()) **in if** w.get() **then skip else abort fi**

Again, the client is dependent on representation.

Note that the cast could not be used if the scope of class name Bool did not include the client. This suggests a focus on modules ("packages" in Java) for confinement of pointers, as has been studied by Vitek and Bokowski [2001] among others (see Sect. 12). But in our example the field has private scope, each rep is associated with a single owner, and the coupling relation is expressed in terms of a single owner. Our results account for this sort of instance-based encapsulation. Instance-based encapsulation facilitates more local or modular reasoning —in particular, a local coupling pertains to a single instance of the owner class. It is suited to many common design patterns, as we illustrate in the sequel, and it is similar to the value-oriented notions used for representation independence in functional languages [Reynolds 1984; Mitchell 1986; 1991].

2.3 Overview of results

In the examples above, class OBool is viewed as providing an abstraction. It is just as sensible to consider Bool as providing an abstraction for which OBool is a client. We do not annotate programs with a fixed designation of owners and reps. Rather, we study how to reason about a class, say *Own*, that one has chosen to view as an abstraction with encapsulated representation. Instances of any subclass of *Own* are also considered to be owners. A second class, say *Rep*, is designated as the type of reps for *Own* (reps can be any subclass of *Rep*). In practice, *Rep* could be an interface or class type, and there could be multiple *Rep* classes; these generalizations are straightforward but would complicate the formalization.

A complete program is a closed collection of class declarations, called a *class table*.[2] We consider an idealized Java-like language similar to the sequential fragment of C++ (without pointer arithmetic), Modula-3, Oberon, C#, Eiffel, and other class-based languages. It includes subclassing and dynamic dispatch, class oriented visibility control, recursive types and methods, type casts and tests (Java's `instanceof`), and a simple form of module.

[2]Our theory provides for modular reasoning about a single class in the context of an arbitrary class table. But it simplifies the semantic model to assume that the class table is closed.

Roughly speaking, a class table CT is *confined*, for Own and Rep, if all of its methods preserve heap confinement. A *confined heap* is one where the objects can be partitioned into some owner islands (recall Fig. 1) along with a block of client objects as in Fig. 5. Furthermore, there are no references from clients to reps. (We use the term *client* for all objects except owners and reps.)

Sect. 3 discusses confinement in more detail and the formal definitions are the subject of Sect. 6. The full significance of the definitions does not become clear until Sect. 9 where we study subclasses of Own: an object of such a type inherits the methods and private fields of Own, which manipulate reps. To be useful, owner subclasses must have some access to reps. On the other hand, as mentioned in Sect. 1, full access is not granted in our formulation since we treat revision of a single class.

Our objective is to compare versions of Own that may use different reps. We say CT and CT' are *comparable* if they are identical except for having different versions of class Own, and those two versions declare the same public methods. The two versions of Own may well use different rep classes, say Rep and Rep'. Without loss of generality, our formalization has both Rep and Rep' present in CT and in CT'.

A key question is how to formalize local couplings, step (1) of the proof method outlined in Sect. 2.1. To allow useful data structures, we need to allow representations to include pointers to client objects (e.g., elements of the FIFO queue in Fig. 1). But if a coupling depends on the state of some object that can be updated by clients, it is difficult to reason modularly about the two versions of Own. We have chosen to use relations that depend only on encapsulated state. Put differently: those things on which a coupling depends are considered as part of the island. Although other alternatives merit study, this one makes for transparent application of the formal results to interesting examples (this is done in Sects. 8 and 9). Moreover, it is straightforward to define the induced coupling.

A *local coupling* is a relation between a pair of owner islands for comparable CT and CT'. A simple example is given by $(*)$ above in Sect. 2.2. More interesting is the observer example, discussed in Sect. 3, which uses a linked list of client objects (the observers). In Fig. 7 on page 37, a local coupling is depicted in which the observer objects occur as dangling pointers from the corresponding islands. The point is that both versions are manipulating the same observer objects in the same way, including the invocation of methods on those objects. So the state of the observer objects is not relevant in the local coupling —nor could it be, if the reasoning is to be carried out in a modular way independent of the particular clients.

In a related pair of islands, both owners have the same class, which may well be a proper subclass of Own.

The induced *coupling relation* for heaps relates h to h' just if there are confining partitions of h and h' for which corresponding islands are pairwise related by the local coupling. Moreover, there is an exact correspondence between client objects in h and h'. Primitive values are related by equality. Locations are related by an arbitrary bijective renaming, which is needed to account for differences in allocation behavior.

The induced coupling is a *simulation* if it is preserved by the methods of class Own in CT and in CT'. A method declared in one version of Own may be inherited

in the other version; it is the behavior of those methods that matters.

The abstraction theorem says that the induced coupling is preserved by all methods of all classes, provided that it is a simulation and both class tables are confined. The identity extension lemma says that the induced coupling is the identity, after garbage collection, for client states in which no owners are reachable.

Sect. 7 gives the formal definitions for coupling and simulation in the special case where locations of objects other than reps are related by equality. The abstraction and identity extension results are proved there in detail. Sect. 10 generalizes the definitions to allow an arbitrary bijection on locations; abstraction and identity extension are proved for the general case. The special case is of interest because it is simpler and adequate for some non-trivial examples like those of Sect. 3 (as shown in Sect. 8). Examples that require the general case are given in Sect. 9; they are subclasses of *Own* that construct reps and pass them to methods of *Own* as in the factory pattern [Gamma et al. 1995]. Notation is more complicated for the general case but the proofs are not very different from the special case.

These results are proved in terms of a semantic formulation of confinement; indeed, the details of this formulation come directly from what is needed in the proofs. Sect. 11 gives a syntax-directed static analysis: typing rules that characterize *safe* programs and a soundness proof that safety implies confinement. Our objective is to round out the story by showing how confinement can be achieved in practice, not to give a definitive treatment of static analyses. But our analysis accepts many natural examples and the constraints are clearly motivated in the proof of soundness. The analysis is modular: It does not require code annotations and the only constraint it imposes on client programs is that they cannot manufacture reps.

3. OWNERSHIP CONFINEMENT

This section considers two examples of representation independence. The first is an object-oriented version of an example given by Meyer and Sieber [1988] as a challenge for semantics of Algol. It illustrates the expressiveness of object-oriented constructs, specifically the use of *callbacks* which go against the hierarchical calling structure which typifies the simplest forms of procedural and data abstraction.

The second example is an instance of the observer pattern [Gamma et al. 1995] which is widely used in object-oriented programs. In addition to callbacks it involves a non-trivial data structure and outgoing references from representation objects to clients. Note that we use the term *client* not just for objects that use the abstraction (owner class) of interest, by instantiating it or calling its methods, but for all objects except instances of the owner and its reps.

This section concludes with an overview of our semantic notion of confinement, Sect. 3.3, to which some readers may want to jump now.

3.1 Callbacks

Meyer and Sieber [1988] consider the following pair of Algol commands:

$$\mathbf{var}\ \mathsf{n} := 0;\ \mathsf{P}(\mathsf{n} := \mathsf{n}{+}2);\ \mathbf{if}\ \mathsf{n}\ \mathbf{mod}\ 2 = 0\ \mathbf{then\ abort\ else\ skip\ fi} \qquad (*)$$

$$\mathbf{var}\ \mathsf{n} := 0;\ \mathsf{P}(\mathsf{n} := \mathsf{n}{+}2);\ \mathbf{abort} \qquad (\dagger)$$

Both invoke some procedure P, passing to it the command $\mathsf{n} := \mathsf{n}{+}2$ that acts on

local variable n. (That is, P is passed a parameterless procedure whose calls have the effect n := n+2.) For any P, the commands are equivalent: both abort. The reason is that in the first example n is invariably even: P is declared somewhere not in the scope of n so the variable can only be affected by (possibly repeated) executions of n := n+2 and this maintains the invariant.

The difficulty in formalizing this argument is due to the difficulty of capturing the semantics of lexically scoped local variables and procedures in a language where local variables can be free in procedures that can be passed as arguments to other procedures. It appears even more difficult, and remains an open problem, to cope with assignment of such procedures to variables (see Sect. 12.1).

Now we consider a Java-like adaptation of the example, due to Peter O'Hearn. In place of local variable n it uses a private field g in a class A. Instead of passing the command n := n+2 as argument, an A-object passes a reference to itself; this gives access to a public method inc that adds 2 to the field.

```
class A extends Object {
  int g; // (the default integer value is 0)
  unit callP(C y){ y.P(self); if self.g mod 2 = 0 then abort else skip fi }
  unit inc(){ self.g := self.g + 2 } }
```

In the context of this class and some declaration of class C with method P, the Algol command (∗) corresponds to the command

$$\mathsf{C\ y} := \mathbf{new}\ \mathsf{C}\ \mathbf{in}\ \mathsf{A\ x} := \mathbf{new}\ \mathsf{A}\ \mathbf{in}\ \mathsf{x.callP(y)} \tag{‡}$$

This aborts because after calling y.P, method callP aborts. The command (†) also corresponds to (‡) but in the context of an alternative implementation of class A:

```
class A extends Object {
  int g;
  unit callP(C y){ y.P(self); abort }
  unit inc(){ self.g := self.g + 2 } }
```

In Example 8.3, we use the abstraction theorem to prove equivalence of the two versions using local coupling relation $o.g = o'.g \wedge o.g \ \mathbf{mod}\ 2 = 0$. This relation is preserved by arbitrary P because P can affect the private field g only by calls to inc.[3]

The example illustrates what are known as *callbacks* in object-oriented programs. When an A-object invokes y.P(self) it passes a reference to itself, by which y may invoke a method on the A-object which is in the middle of executing method callP —a callback to A. If in (‡) we replace x.callP(y) by x.callP(self), and assume that (‡) is a constituent of a method of class C, then we get a callback to C.

[3]As Reynolds [1978] shows (see also [Reddy 2002]), instance-based object-oriented constructs can be expressed in Algol-like languages, but the latter are in some ways significantly more powerful. The Java version of the example can be seen as giving an explicit closure (method inc) to represent the command n := n+2. The fact that we can prove the result in a simple semantic model can be explained by saying the language is defunctionalized [Reynolds 1972; Banerjee et al. 2001] and lacks true higher order constructs. So P ranges over more limited procedures than in Algol.

The point of the Algol example is modular reasoning about (∗) and (†) independent from the definition of P. For the object-oriented version we can also consider reasoning independent from subclasses of A. If instead of (‡) we consider a method

```
unit m(C y, A x){ x.callP(y) }
```

then there is the possibility that m is passed an argument x of some subtype of A that overrides inc. By dynamic binding, the overriding implementation would be invoked by callP. One might expect that, for modular reasoning, we need to require that inc and callP satisfy the conditions of behavioral subclassing [Liskov and Wing 1994; Dhara and Leavens 1996]. But note that the subclass cannot directly access private field g, so no matter what the overriding methods do they cannot violate the invariant that g is even. Representation independence does not depend on behavioral subclassing. But behavioral subclassing is essential for modular reasoning, in particular to prove the simulation property (step (2) in Sect. 2.1).

3.2 The observer pattern

In this subsection we consider variations on the often-used Observer design pattern [Gamma et al. 1995], which involves a non-trivial recursive data structure using multiple rep objects and outgoing references to client objects. Further variations are given in Sect. 9.

We focus attention on the abstraction provided by an Observable object (sometimes called the “subject” role). It maintains a list of so-called observers (“views”) to be notified when some event occurs. Its public method add allows the addition of an observer object to the list. The public method notifyAll represents the event of interest; for our purposes, its effect is simply to invoke method notify on each observer in the list. What notify does is not relevant, so long as it is confined. This example is similar to “collection classes” which hold a set of references to client objects; a comparison method may be invoked on the clients for sorting etc.

In the first version of the observer example, Fig. 2, most of the work is done by the owner class Observable, which uses rep class Node to store observers in a singly linked list. Versions in more object-oriented style can be found in the companion technical report [Banerjee and Naumann 2004a].

Fig. 3 gives example client classes AnObserver and Main. Class AnObserver records notifications in its state. Method main constructs and initializes an Observable, installs an observer, ob, and invokes notifyAll; upon termination, ob.count = 1 and no Observable is reachable.

Fig. 4 gives another version of Observable, using a sentinel node [Cormen et al. 1990], for the sake of an example. In Sect. 8 we show equivalence of the versions of Figs. 2 and 4 as an application of the abstraction theorem and identity extension. The coupling relation describes the correspondence between a pair of lists, one with and one without a sentinel node (see Fig. 7). It is enough to say that the same Observer locations are stored in the lists, in the same order: The state of the Observer is not relevant —nor could it be in a modular treatment, as class Observer has no fields. To reason about outgoing calls, namely to notify, it is enough to show that the two implementations make the same calls. Those calls may lead to calls back to the Observable, but encapsulation ensures that those calls are the only way the behavior of notify can depend on, or affect, the Observable.

```
class Observer extends Object { // "abstract class" to be overridden in clients
  unit notify(){ abort } }

class Node extends Object { // rep for Observable
  Observer ob;
  Node nxt; // next node in list
  unit setOb(Observer o){ self.ob := o }
  unit setNext(Node n){ self.nxt := n }
  Observer getOb(){ result := self.ob }
  Node getNext(){ result := self.nxt } }

class Observable extends Object { // owner
  Node fst; // first node in list
  unit add(Observer ob){ Node n := new Node; n.setOb(ob); n.setNext(self.fst); self.fst := n }
  unit notifyAll(){ Node n := self.fst; while n ≠ null do n.getOb().notify(); n := n.getNext() od } }
```

Fig. 2. First version of observer pattern, in procedural style.

```
class AnObserver extends Observer {
  int count;
  unit notify(){ self.count := self.count+1 } }

class Main extends Object {
  AnObserver ob;
  unit main(){
    ob := new AnObserver; Observable obl := new Observable; obl.add(ob); obl.notifyAll() } }
```

Fig. 3. Example client for Observable.

```
class Node2 extends Object { // rep for Observable
  Observer ob;
  Node2 nxt;
  unit setOb(Observer o){ self.ob := o }
  unit setNext(Node2 n){ self.nxt := n }
  Observer getOb(){ result := self.ob }
  Node2 getNext(){ result := self.nxt } }
class Observable extends Object // owner {
  Node2 snt; // sentinel node pointing to list
  constr { self.snt := new Node2 }
  unit add(Observer ob){
    Node2 n := new Node2; n.setOb(ob); n.setNext(self.snt.getNext()); self.snt.setNext(n); }
  unit notifyAll(){
    Node2 n := self.snt.getNext(); while n ≠ null do n.getOb().notify(); n := n.getNext() od } }
```

Fig. 4. Version of observable that uses sentinel node, in procedural style

Except for the bad method of Sect. 2.1, all of the examples discussed so far satisfy the confinement conditions discussed next.

3.3 Confinement

Fig. 5 illustrates instance-based owner confinement; in this case Nodes are confined to their owning Observable. Following Hogg [1991], we use the term *island* for the

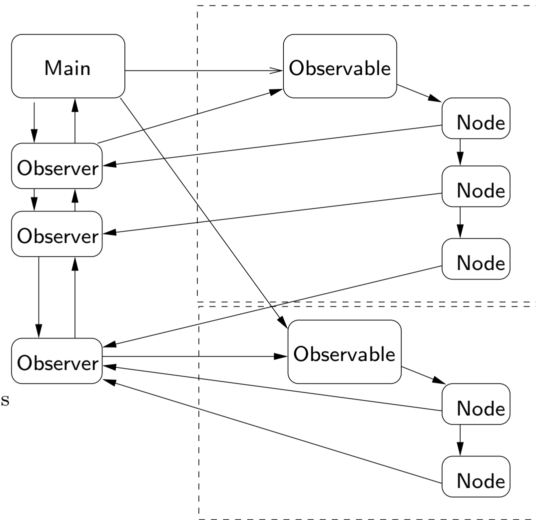

Fig. 5. Confinement example. Rounded boxes are instances of the indicated class. Solid arrows represent allowed pointers. Dashed boxes indicate owner islands, each consisting of one owner and its reps.

PSfrag replacements

sub-heap consisting of an owner and its reps. Dashed lines in the Figure depict two islands. Our notion of owner confinement imposes four conditions on islands; here are the first three:

(1) there are no references from a client object to a rep;
(2) there are no references from an owner to reps in a different island;
(3) there are no references from a rep to reps in a different island.

The Figure exhibits most allowed references, but we also allow an owner and its reps to reference another owner (see Fig. 6 on page 28). An example is given in Sect. 9.1. Note that heap confinement is a state predicate. The full definition, formalized in Sect. 6, deals with preservation of this predicate by commands and with leaks via parameter passing in outgoing method calls from island to client.

In class-based languages with inheritance, there is a subclass (or "protected") interface in addition to the public one. This raises the possibility of expressing encapsulation of reps for not only (instances of) the owner class but also its subclasses. Because we study the replacement of a single class, not a class together with its subclasses, we must treat subclasses like clients in that fields they declare may *not* point to reps. To the list of conditions above we add:

(4) references from an owner's fields to its reps are only in the private fields of the owner class.

In order not to abandon the expressiveness of subclassing, however, we allow subclass methods to manipulate reps: they may be constructed, stored in local variables, and passed to the owner. This fits well with the factory pattern [Gamma et al. 1995] which allows owner behavior to be adapted in owner subclasses without violating encapsulation. To balance the paper, we have deferred the relevant examples to Sect. 8. Sect. 12 has further discussion of the "protected" interface.

Confinement is formulated using class names. Two incomparable class names, *Own* and *Rep*, are designated. An object is considered to be an owner (respectively,

a rep) if its type is *Own* (resp. *Rep*) or a subtype thereof. Incomparability is a mild restriction that enforces a widely-followed discipline of distinguishing between rep objects (e.g., nodes in a linked list) and objects representing abstractions (e.g., a list). The technical benefit of incomparability is that if C and D are incomparable, which we write $C \not\lesseqgtr D$, then an expression of type C never has a value of type D.

We aim for clear separation between the semantic property needed for the abstraction theorem —restrictions on the heap as described above— and the syntactic conditions used by the static analysis to enforce the confinement property. For perspicuity, the separation is not absolute: the "semantic" property includes conditions on method signatures. For example, we impose the restriction that the return type of a public owner method is incomparable to *Rep*.

Our use of types to formulate alias restrictions allows heterogeneous data structures, but is slightly restrictive in that there is a single common superclass for all reps and that class cannot also be used by clients. (Generalizations are mentioned in Sect. 2.3.) A more substantial restriction is due to the fact that class **Object** is comparable to all classes. Because Java lacks parametric polymorphism, **Object** is often used to express generics, e.g., a list containing elements of arbitrary type. A method to enumerate the list would have return type **Object**, which violates our restriction on owner methods. This restriction could be dropped in favor of more sophisticated conditions to ensure that no rep is returned (see Sect. 12). But in practice many generics have some sort of constraint expressed by a class or interface type —like Observer in our examples, or Comparable for data structures that depend on an ordering. These do not run afoul of our restriction. In any case, the use of **Object** for generics is widely deplored because it undercuts the benefits of typing; parametric types are clearly preferable.

Some works on confinement have considered all the confinement properties intended to be satisfied by a program, using hierarchical notions of ownership [Clarke et al. 2001; Müller 2002]. For example, a Set could own the header of a list which in turn owns the nodes of the list. This is not necessary for our purposes (see Sect. 12). To analyse the abstraction provided by the set, we would consider both the header and nodes to be reps. On the other hand, to replace one header implementation by another, Set is irrelevant; we choose *Own* to be the header and *Rep* for the nodes.

What our formalization does not allow is for a single class to be considered as both an abstraction to be revised and as an internal representation thereof. An instance of *Own* cannot be part of the island of another instance of *Own*; in other words, islands are not nested. Reps for one instance of *Own* can point to another instance of *Own*, allowing, for example, a collection to contain collections of the same kind. But a collection cannot use another collection of the same kind as part of its encapsulated representation. This restriction is not needed for soundness of simulation. Nested islands can be handled, at the cost of a healthiness condition on couplings, saying that the coupling holds recursively on nested islands (this is worked out by Cavalcanti and Naumann [2002, Def. 5] for a language without sharing). The restriction simplifies our theory and fits comfortably with the notations we borrow from Separation Logic. It precludes, for example, a class *List* that serves both as the interface used by clients and also as the recursive node type. Such examples are considered poor programming practice.

4. SYNTAX

This section formalizes the language, for which purpose we adapt some notations from Featherweight Java [Igarashi et al. 2001].[4] To avoid burdening the reader with straightforward technicalities we deliberately confuse concrete syntax with abstract syntax. We do not distinguish between classes and class types. We confuse syntactic categories with names of their typical elements. Barred identifiers like $\overline{T}$ indicate finite lists, e.g., $\overline{T}\,\overline{f}$ stands for a list $\overline{f}$ of field names with corresponding types $\overline{T}$. The bar has no semantic import; $\overline{T}$ has nothing to do with T.

The grammar is based on given sets of class names (with typical element C and including at least **Object**), field names (f), method names (m), and names (x) for parameters and local variables. In most respects self and result are like any other variables but self cannot be the target of assignment.

Grammar

T	$::=$	$\mathbf{bool} \mid \mathbf{unit} \mid C$	data type
CL	$::=$	$\mathbf{class}\ C\ \mathbf{extends}\ C\ \{\ \overline{T}\ \overline{f};\ \overline{M}\ \}$	class declaration
M	$::=$	$T\ m(\overline{T}\ \overline{x})\ \{S\}$	method declaration
S	$::=$	$x := e \mid e.f := e$	assign to variable, to field
	$\mid$	$x := \mathbf{new}\ C$	object construction
	$\mid$	$x := e.m(\overline{e})$	method call
	$\mid$	$T\ x := e\ \mathbf{in}\ S$	local variable block
	$\mid$	$\mathbf{if}\ e\ \mathbf{then}\ S\ \mathbf{else}\ S\ \mathbf{fi} \mid S;\ S$	conditional, sequence
e	$::=$	$x \mid \mathbf{null} \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{it}$	variable, constant
	$\mid$	$e.f \mid e = e$	field access, equality test
	$\mid$	$e\ \mathbf{is}\ C \mid (C)\ e$	type test, cast

Class **Object** has no fields, no methods, and no proper superclass. Additional primitive types, such as integers, can be treated in the same way as **bool** and **unit** (integers can also be represented, e.g., in unary using linked lists).

In the formal language, expressions do not have side effects. Object construction, **new**, occurs only as a command $x := \mathbf{new}\ C$ that assigns to a local variable. Method calls are not expressions but rather occur in special assignments $x := e.m(\overline{e})$ to allow both heap effects and a return value.

Remark* 4.1 (*syntactic sugar) In examples we use several abbreviations:

—A method call command $e.m(\overline{e})$, e.g., self.g.set(true), abbreviates a call assigning to an otherwise unused local variable.

—Assignment of a new object to a field abbreviates a local block assigning the new object to a variable that is then assigned to the field.

—Object construction in local variable initialization abbreviates initialization to **null** followed by object construction.

[4] But the languages differ, e.g., ours has imperative features and private fields.

—Methods that return values but do not mutate state are used in expressions, e.g., the argument in $\mathsf{self.inout} := \mathsf{convertToString(z.getg())}$ and the target object in $\mathsf{n.getOb().notify()}$. These are easily desugared using fresh variables and suitable assignments.

—**skip** abbreviates some no-op assignment $x := x$.

Since methods can be defined recursively, we omit loops.[5] □

A program is given as a *class table* CT, a finite partial function sending class name C to its declaration $CT(C)$ which may make mutually recursive references to other classes. Well formed class tables are characterized using typing rules which are expressed using some auxiliary functions that in turn depend on the class table, as is needed to allow mutual recursion. Consider a declaration

$$CT(C) = \textbf{class } C \textbf{ extends } D \; \{ \; \overline{T}_1 \; \overline{f}; \; \overline{M} \; \} \;\; .$$

For the direct superclass of C, we define $\mathit{super}\, C = D$. Let M be in the list $\overline{M}$ of method declarations, with

$$M = T \; m(\overline{T}_2 \; \overline{x}) \; \{S_2\} \;\; .$$

We record the typing information by defining $\mathit{mtype}(m, C) = \overline{T}_2 {\rightarrow} T$. (Note that $\overline{T}_2 {\rightarrow} T$ is not a data type in the language.) For the parameter names we define $\mathit{pars}(m, C) = \overline{x}$. If m has no declaration in $CT(C)$ but $\mathit{mtype}(m, D)$ is defined then m is an inherited method, for which we define $\mathit{mtype}(m, C) = \mathit{mtype}(m, D)$ and $\mathit{pars}(m, C) = \mathit{pars}(m, D)$. For the declared fields, we define $\mathit{type}(\overline{f}, C) = \overline{T}_1$ and $\mathit{dfields}\, C = (\overline{f} : \overline{T}_1)$. Here $\overline{f} : \overline{T}_1$ denotes a finite mapping of field names to types. To include inherited fields, we define $\mathit{fields}\, C = \mathit{dfields}\, C \cup \mathit{fields}\, D$ and assume $\overline{f}$ is disjoint from the names in $\mathit{fields}\, D$. The distinguished class **Object** has no methods, $\mathit{fields}(\textbf{Object})$ is the empty list, and $\mathit{super}(\textbf{Object})$ is undefined.

A *typing context* Γ is a finite mapping from variable and parameter names to data types, such that $\mathsf{self} \in \mathit{dom}\, \Gamma$. Whereas the Java format $\mathsf{T \; x}$ is used in code to give x type T, it is written $\mathsf{x{:}T}$ in typing contexts. Typing of commands for methods declared in class C is expressed using judgements $\Gamma \vdash S$ where $\Gamma\, \mathsf{self} = C$. Moreover, if $\mathit{mtype}(m, C) = \overline{T} {\rightarrow} T$ and $\mathit{pars}(m, C) = \overline{x}$ then $\Gamma\, \overline{x} = \overline{T}$ and $\Gamma\, \mathsf{result} = T$. We sometimes say "command" rather than the more precise "command in context" to refer to a derivable judgement $\Gamma \vdash S$. The judgement $\Gamma \vdash e : T$ says that expression e has type T.

***Definition* 4.2 (subtyping , $\leq$)** The class table determines a subtyping relation $\leq$, where $T \leq U$ means T is a subtype of U, as follows. If T or U is **bool** or **unit** then define $T \leq U$ iff $T = U$. For class types C and D, define $C \leq D$ iff either $C = D$ or $\mathit{super}\, C \leq D$. □

The definition of well formed class table requires that $\leq$ is acyclic and as a consequence we have $C \leq \textbf{Object}$ for all C.

Subsumption is built into the rules for specific constructs. For example, the assignment rule allows $x : D, y : E, \mathsf{self} : C \vdash x := y$ provided that $E \leq D$.

[5]But recursion cannot be coded by Landin's trick of recursion through the heap, because methods are statically bound to classes.

***Definition* 4.3 (well formed class table)** A class table is well formed provided it satisfies the following conditions.

—Each class declaration **class** C **extends** D $\{\ \overline{T}\ \overline{f};\ \overline{M}\ \}$ is well formed, that is, each method declaration M in $\overline{M}$ is well formed, according to the rules to follow.

—If C occurs as the type of a field, parameter, or local variable in some class then $CT(C)$ is defined. No field or method has multiple declarations in a class.

—The subclass relation $\leq$ is antisymmetric. □

The rules are straightforward renderings of the typing rules for Java, for private fields, public methods and public classes [Arnold and Gosling 1998].

Typing of method declarations

$$\frac{\begin{array}{c}\overline{x}:\overline{T}, \mathsf{self}:C, \mathsf{result}:T \vdash S \\ \mathit{mtype}(m, \mathit{super}\,C) \text{ is undefined or equals } \overline{T}\rightarrow T \\ \mathit{pars}(m, \mathit{super}\,C) \text{ is undefined or equals } \overline{x}\end{array}}{C \vdash T\ m(\overline{T}\ \overline{x})\{S\}}$$

In this method rule, the condition on *mtype* is the standard invariance restriction on method types, as in Java [Arnold and Gosling 1998; Abadi and Cardelli 1996]. The last antecedent in the rule, concerning $\mathit{pars}(m, D)$, ensures that all declarations of a method use the same parameter names. This loses no generality and slightly streamlines the formalization of the semantic domains and couplings.

Typing of expressions

$$\Gamma \vdash x : \Gamma x \quad \Gamma \vdash \mathbf{null} : B \quad \Gamma \vdash \mathbf{it} : \mathbf{unit} \quad \Gamma \vdash \mathbf{true} : \mathbf{bool} \quad \Gamma \vdash \mathbf{false} : \mathbf{bool}$$

$$\frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash e_1 = e_2 : \mathbf{bool}} \qquad \frac{\Gamma \vdash e : (\Gamma\,\mathsf{self}) \quad (f : T) \in \mathit{dfields}(\Gamma\,\mathsf{self})}{\Gamma \vdash e.f : T}$$

$$\frac{\Gamma \vdash e : D \quad B \leq D}{\Gamma \vdash (B)\ e : B} \qquad \frac{\Gamma \vdash e : D \quad B \leq D}{\Gamma \vdash e\ \mathbf{is}\ B : \mathbf{bool}}$$

The rule for equality test allows comparison of arbitrary data types, and is reference equality in the case of class types. (If e_1 and e_2 have types not related by $\leq$, the test $e_1 = e_2$ is false except when both are null.) The rule for field access enforces private visibility: Only a method declaration in class C can access fields declared in $CT(C)$ —it can access those fields on any object of its type. To access its own fields the expression is $\mathsf{self}.f$. The rule for cast is standard.[6]

[6]It is not adequate for expressions that arise through substitutions used in program logic (see Cavalcanti and Naumann [1999]) and in small-step semantics (see Igarashi et al. [2001]).

Typing of commands

$$\frac{\Gamma \vdash e : T \quad T \leq \Gamma\, x \quad x \neq \mathsf{self}}{\Gamma \vdash x := e} \qquad \frac{\Gamma \vdash e_1 : (\Gamma\, \mathsf{self}) \quad \Gamma \vdash e_2 : U \quad U \leq T \quad (f : T) \in \mathit{dfields}(\Gamma\, \mathsf{self})}{\Gamma \vdash e_1.f := e_2} \qquad \frac{\Gamma \vdash S_1 \quad \Gamma \vdash S_2}{\Gamma \vdash S_1;\ S_2}$$

$$\frac{\Gamma \vdash e : D \quad \mathit{mtype}(m, D) = \overline{T} \rightarrow T \quad \Gamma \vdash \overline{e} : \overline{U} \quad \overline{U} \leq \overline{T} \quad T \leq \Gamma\, x \quad x \neq \mathsf{self}}{\Gamma \vdash x := e.m(\overline{e})} \qquad \frac{B \leq \Gamma x \quad B \neq \mathbf{Object} \quad x \neq \mathsf{self}}{\Gamma \vdash x := \mathbf{new}\ B}$$

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma \vdash S_1 \quad \Gamma \vdash S_2}{\Gamma \vdash \mathbf{if}\ e\ \mathbf{then}\ S_1\ \mathbf{else}\ S_2\ \mathbf{fi}} \qquad \frac{\Gamma \vdash e : U \quad U \leq T \quad (\Gamma, x : T) \vdash S \quad x \neq \mathsf{self} \quad x \notin \mathit{dom}\, \Gamma}{\Gamma \vdash T\ x := e\ \mathbf{in}\ S}$$

In some of the command rules, the hypothesis involves a partial function which must be defined for the hypothesis to be satisfied. For example, in the rule for method call, $\mathit{mtype}(m, D)$ must be defined and equal to $\overline{T} \rightarrow T$.

Each expression and command construct is the conclusion of exactly one typing rule, and there are no other rules. Thus we have the following.

Lemma 4.4 A typing $\Gamma \vdash S$ or $\Gamma \vdash e : T$ has at most one derivation.[7] □

***Definition* 4.5 (inheritance)** Method m is *inherited in C from B* if $C \leq B$, there is a declaration for m in B, and there is no declaration for m in any D such that $C \leq D < B$. To make the class table explicit, we also say m is inherited from B in $CT(C)$. □

Because the language has single inheritance, the subtyping relation $\leq$ is a tree: if $D \leq B$ and $D \leq C$ then $B \leq C$ or $C \leq B$. If $\mathit{mtype}(m, C)$ is defined for some C then it is defined for all subclasses of C and there is a unique ancestor class declaring m that is least with respect to $\leq$.

Lemma 4.4 justifies proofs and definitions by structural induction on typings. The following notion facilitates induction on inheritance chains.

***Definition* 4.6 (method depth)** For any m and C such that $\mathit{mtype}(m, C)$ is defined, the *method depth of C for m in CT* is defined by $\mathit{depth}(m, C) = 1 + \mathit{depth}(m, \mathit{super}\, C)$ if $\mathit{mtype}(m, \mathit{super}\, C)$ is defined; otherwise, $\mathit{depth}(m, C) = 0$. □

An immediate consequence is that if $\mathit{mtype}(m, C)$ is defined and $\mathit{depth}(m, C) = 0$ then $CT(C)$ has a declaration for m.

5. SEMANTICS

This section defines the semantic domains, then the semantics of expressions and commands, and finally the semantics of well formed class tables.

[7] Strictly speaking this is not quite true, because in a context where **null** is typed as C it can also be typed as some subtype of C. But this has no bearing on semantics. There are several straightforward solutions to the problem and we leave it to the interested reader.

Because methods are associated with classes rather than with instances, the semantic domains are rather simple. There are no recursive domain equations to be solved: subclassing ($\leq$) is acyclic and the cycle of recursive references via class fields is broken via the heap. Mutually recursive method invocations can arise through direct calls on a single object and also through callbacks between reachable objects, as for example in the observer pattern. We impose no restrictions on such calls. A fixpoint construction is used for the method environment which comprises the semantics of the class table.

Often we write = between expressions involving partial functions such as those used in typing. Unless otherwise indicated, it means strong equality: both sides are defined and equal.

5.1 Semantic domains

The state of a method in execution is comprised of a *heap* h, which is a finite partial function from locations to object states, and a *store* η, which assigns locations and primitive values to the local variables and parameters given by a typing context Γ.[8] An *object state* is a mapping from field names to values. Function application associates to the left, so $h\,\ell\,f$ is the value of field f of the object $h\,\ell$ at location ℓ.

A command denotes a function mapping each initial state (h, η) either to a final state (h_0, η_0) or to the distinguished value $\bot$. We use the term *global state* for (h, η), to distinguish it from object states. The improper value $\bot$ represents non-termination as well as runtime errors: attempts to dereference *nil* or cast a location to a type it does not have.

In some languages it is a runtime error to dereference a dangling pointer, i.e., one not in the domain of the heap. In Java dangling pointers cannot arise: there is no command for deallocation and a correct garbage collector never deallocates reachable objects. For our purposes, garbage collection need not be modeled. Commands act on heaps and stores that are closed in the sense that all locations that occur are in the domain of the heap. The following paragraphs formalize our assumptions about locations and then define the semantic domains.

For locations, we assume that a countable set *Loc* is given, along with a distinguished value *nil* not in *Loc*. To track each object's class we assume given a function $loctype : Loc \to ClassNames$ such that for each C there are infinitely many locations ℓ with $loctype\ \ell = C$. We use the term *heap* for any partial function h such that $dom\,h \subseteq_{fin} Loc$ and each $h\,\ell$ is an object state of type $loctype\ \ell$. Object states are formalized later. Because the domain of a heap is finite, the assumption about *loctype* ensures an adequate supply of fresh locations.

We write $locs\ C$ for $\{\ell \in Loc \mid loctype\,\ell = C\}$, and $locs(C{\downarrow})$ for $\{\ell \mid loctype\,\ell \leq C\}$. There is no independent meaning for the notation $C{\downarrow}$.

***Definition* 5.1 (allocator, parametric)** An *allocator* is a location-valued function *fresh* such that $loctype(fresh(C, h)) = C$ and $fresh(C, h) \notin dom\,h$, for all C, h. An allocator is *parametric* if $dom\,h_1 \cap locs\,C = dom\,h_2 \cap locs\,C$ implies $fresh(C, h_1) = fresh(C, h_2)$. $\Box$

[8] "η" is mnemonic for "environment", the term we used in [Banerjee and Naumann 2002] to avoid the irrelevant connotations of "stack", before we adopted "store" following Reynolds [2001].

For example, taking $Loc = \mathbb{N}$ and *loctype* to be arbitrary, a parametric allocator is given by the function $\mathit{fresh}(C,h) = \mathit{min}\{\ell \mid \mathit{loctype}\,\ell = C \wedge \ell \notin \mathit{dom}\,h\}$.

In implementations, the object class is usually encoded as part of its state. One could uncurry this representation of heaps and take *Loc* to be $\mathbb{N} \times ClassNames$. Then $\mathit{fresh}(C,h)$ could return (n,C) where $n = \mathit{min}\{k \mid \neg(\exists B \cdot (k,B) \in \mathit{dom}\,h)\}$. This is an allocator that is not parametric, because the presence of objects of one class affects the availability of memory for objects of other classes.

We define the semantics in terms of an arbitrary allocator *fresh*. The assumption of parametricity is stated explicitly where it is needed, namely for the first abstraction theorem (Sect. 7) but not the second (Sect. 10). Parametricity of the allocator is a reasonable assumption for some applications but not all. The assumption streamlines the proof of the abstraction theorem, allowing us to highlight other issues. For the second abstraction theorem, we drop parametricity and complicate the definitions of coupling and simulation by adding a bijective renaming of locations.

In addition to heaps, it is convenient to name a number of other semantic categories that are explained in due course. In the following, θ ranges over suggestive notations and each θ has an associated set $[\![\theta]\!]$ of values.

Semantic categories

$$\theta \ ::= T \mid \Gamma \mid \mathit{state}\,C \mid \mathit{Heap} \mid \mathit{Heap} \otimes \Gamma \mid \mathit{Heap} \otimes T \mid \theta_\perp \mid C,\, \overline{x},\, \overline{T}{\to}T \mid \mathit{Menv}$$

The funny symbol in $\mathit{Heap} \otimes \Gamma$ is meant to indicate that the set $[\![\mathit{Heap} \otimes \Gamma]\!]$ is a dependent form of cartesian product, specifically a heap paired with a store closed in the heap. Stores are among the simpler semantic domains.

Semantics of types, object states, and stores

$$\begin{array}{lcl}
[\![\mathbf{bool}]\!] & = & \{\mathit{true}, \mathit{false}\} \\
[\![\mathbf{unit}]\!] & = & \{\mathit{it}\} \\
[\![C]\!] & = & \{\mathit{nil}\} \cup \mathit{locs}(C{\downarrow}) \\
[\![\mathit{state}\,C]\!] & = & \{s \mid \mathit{dom}\,s = \mathit{dom}(\mathit{fields}\,C) \wedge \forall (f:T) \in \mathit{fields}\,C \cdot s\,f \in [\![T]\!]\} \\
[\![\Gamma]\!] & = & \{\eta \mid \mathit{dom}\,\eta = \mathit{dom}\,\Gamma \wedge \eta\,\mathsf{self} \neq \mathit{nil} \wedge \forall x \in \mathit{dom}\,\eta \cdot \eta\,x \in [\![\Gamma\,x]\!]\}
\end{array}$$

***Definition* 5.2 (closed heap and store)** A heap h is *closed*, written $\mathit{closed}\,h$, iff $\mathit{rng}(h\,\ell) \cap \mathit{Loc} \subseteq \mathit{dom}\,h$, for all $\ell \in \mathit{dom}\,h$. A store $\eta \in [\![\Gamma]\!]$ is *closed in heap h*, written $\mathit{closed}(h,\eta)$, iff $\mathit{rng}\,\eta \cap \mathit{Loc} \subseteq \mathit{dom}\,h$. □

Note that $\mathit{rng}(h\,\ell)$ is the set of values in fields of the object state $h\,\ell$.

Recall that fresh locations should occur nowhere in the global state. For a closed store and heap, this follows from the requirement that $\mathit{fresh}(C,h) \notin \mathit{dom}\,h$.

Semantics of global states and methods

$$
\begin{array}{lcl}
[\![Heap]\!] & = & \{h \mid dom\,h \subseteq_{fin} Loc \wedge closed\,h \\
 & & \quad \wedge \forall \ell \in dom\,h \cdot h\ell \in [\![state\,(loctype\ \ell)]\!]\} \\
[\![Heap \otimes \Gamma]\!] & = & \{(h,\eta) \mid h \in [\![Heap]\!] \wedge \eta \in [\![\Gamma]\!] \wedge closed(h,\eta)\} \\
[\![Heap \otimes T]\!] & = & \{(h,d) \mid h \in [\![Heap]\!] \wedge d \in [\![T]\!] \wedge (d \in Loc \Rightarrow d \in dom\,h)\} \\
[\![\theta_\bot]\!] & = & [\![\theta]\!] \cup \{\bot\} \quad \text{(where } \bot \text{ is some fresh value not in } [\![\theta]\!]) \\
[\![C,\, \overline{x},\, \overline{T}{\rightarrow}T]\!] & = & [\![Heap \otimes (\overline{x}:\overline{T}, \mathsf{self}:C)]\!] \rightarrow [\![(Heap \otimes T)_\bot]\!] \\
[\![Menv]\!] & = & \{\mu \mid \forall C, m \cdot \mu Cm \text{ is defined iff } mtype(m,C) \text{ is defined,} \\
 & & \quad \text{and } \mu Cm \in [\![C, pars(m,C), mtype(m,C)]\!] \text{ if } \mu Cm \text{ defined }\}
\end{array}
$$

Just as a class declaration $CT(C)$ gives a collection of method declarations, the semantics of a class table is a *method environment* that assigns to each class C a method meaning $\mu\,C\,m$ for each m declared or inherited in C.

For the fixpoint construction of the method environment denoted by a class table, we need to impose order on the semantic domains. We use the term *complete partial order* for a poset with least upper bounds of countable ascending chains [Davey and Priestley 1990]. The degenerate case is ordering by equality, which is the order we use for the semantics of T, Γ, $state\,C$, $Heap$, $(Heap \otimes \Gamma)$, and $(Heap \otimes T)$. Then $[\![(Heap \otimes \Gamma)_\bot]\!]$ and $[\![(Heap \otimes T)_\bot]\!]$ are complete partial orders with the "flat" order: $\bot$ is below anything and other comparable elements are equal. The set $[\![C,\, \overline{x},\, \overline{T}{\rightarrow}T]\!]$ is defined to be the space of total functions $[\![Heap \otimes (\overline{x}:\overline{T}, \mathsf{self}:C)]\!] \rightarrow [\![(Heap \otimes T)_\bot]\!]$, all of which are continuous because $Heap \otimes (\overline{x}:\overline{T}, \mathsf{self}:C)$ is ordered by equality. The function space itself is ordered pointwise, making it a complete partial order with minimum element $\lambda(h,\eta) \cdot \bot$. Finally, we order $[\![Menv]\!]$ pointwise. All method environments μ in $[\![Menv]\!]$ have the same domain, determined by CT, so this is also a complete partial order, taken pointwise. It has a minimum element, namely $\lambda C \cdot \lambda m \cdot \lambda(h,\eta) \cdot \bot$.

Whereas $[\![state\ C]\!]$ consists of the states for objects of exactly class C, the set $[\![C]\!]$ is downward closed. For data types T_1, T_2 we have $T_1 \leq T_2 \Rightarrow [\![T_1]\!] \subseteq [\![T_2]\!]$.

***Definition* 5.3 (incomparable, $\lesseqgtr\!\!\!\!/$)** We write $C \not\lesseqgtr B$ for $C \not\leq B \wedge C \not\geq B$. For a list $\overline{C}$, $\overline{C} \not\lesseqgtr B$ means $C \not\lesseqgtr B$ for all C in $\overline{C}$. □

Lemma 5.4 For classes C, B, if $C \not\lesseqgtr B$ then $[\![C]\!] \cap [\![B]\!] = \{nil\}$. For primitive T we have $[\![T]\!] \cap [\![B]\!] = \varnothing$. □

The result is a direct consequence of the definitions. We often use the contrapositive: if there is a non-*nil* location in both $[\![B]\!]$ and $[\![C]\!]$ then $B \leq C$ or $C \leq B$.

5.2 Semantics of expressions, commands and methods

For expressions and commands, the semantics is defined by induction on typing derivations. As a consequence of uniqueness of typing derivations, Lemma 4.4, the semantics is a function of typings. The meaning of a command $\Gamma \vdash S$ will be defined

to be a function

$$[\![\Gamma \vdash S]\!] \in [\![Menv]\!] \to [\![Heap \otimes \Gamma]\!] \to [\![(Heap \otimes \Gamma)_\bot]\!] \ .$$

The meaning of an expression $\Gamma \vdash e : T$ will be defined to be a function

$$[\![\Gamma \vdash e : T]\!] \in [\![Heap \otimes \Gamma]\!] \to [\![T_\bot]\!]$$

such that the result value is always in the domain of the heap if it is a location. This is part of Lemma 5.7.

The command and expression constructs are strict in $\bot$, except, as usual, for the then- and else-commands in **if** $-$ **fi**. To streamline the treatment of $\bot$ in the semantic definitions, the metalanguage construct

$$\mathsf{let}\ d = E_1\ \mathsf{in}\ E_2$$

denotes $\bot$ if the value of E_1 is $\bot$; otherwise, its value is the value of E_2 with d bound to the value of E_1.

We let $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ in the following definitions. Note that identifiers in the defining equations are as in the corresponding typing rules. For semantic values we use the identifier d, but sometimes ℓ for elements of the sets $[\![C]\!]$. For expressions the semantics is straightforward; we choose the Java semantics for casts and tests.

Semantics of expressions

$$
\begin{array}{lcl}
[\![\Gamma \vdash x : T]\!](h,\eta) & = & \eta x \\
[\![\Gamma \vdash \mathbf{null} : B]\!](h,\eta) & = & nil \\
[\![\Gamma \vdash \mathbf{it} : \mathbf{unit}]\!](h,\eta) & = & it \\
[\![\Gamma \vdash \mathbf{true} : \mathbf{bool}]\!](h,\eta) & = & true \\
[\![\Gamma \vdash \mathbf{false} : \mathbf{bool}]\!](h,\eta) & = & false \\
[\![\Gamma \vdash e_1 = e_2 : \mathbf{bool}]\!](h,\eta) & = & \mathsf{let}\ d_1 = [\![\Gamma \vdash e_1 : T_1]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{let}\ d_2 = [\![\Gamma \vdash e_2 : T_2]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ d_1 = d_2\ \mathsf{then}\ true\ \mathsf{else}\ false \\
[\![\Gamma \vdash e.f : T]\!](h,\eta) & = & \mathsf{let}\ \ell = [\![\Gamma \vdash e : (\Gamma\,\mathsf{self})]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ \ell = nil\ \mathsf{then}\ \bot\ \mathsf{else}\ h\,\ell\,f \\
[\![\Gamma \vdash (B)\ e : B]\!](h,\eta) & = & \mathsf{let}\ \ell = [\![\Gamma \vdash e : D]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ \ell = nil \lor loctype\ \ell \leq B\ \mathsf{then}\ \ell\ \mathsf{else}\ \bot \\
[\![\Gamma \vdash e\ \mathbf{is}\ B : \mathbf{bool}]\!](h,\eta) & = & \mathsf{let}\ \ell = [\![\Gamma \vdash e : D]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ \ell \neq nil \land loctype\ \ell \leq B\ \mathsf{then}\ true\ \mathsf{else}\ false
\end{array}
$$

The semantics of commands is defined by structural induction on commands. In the semantics of commands, we write $[\mathit{fields}\,B \mapsto \mathit{defaults}]$ as an abbreviation for the function sending each $f \in dom(\mathit{fields}\,B)$ to the default value for $type(f, B)$. The defaults are *false* for **bool**, *it* for **unit**, and *nil* for classes. Function update or extension is written like $[\eta \mid x \mapsto d]$. We write $\downharpoonright$ for domain restriction: if x is in the domain of η then $\eta \downharpoonright x$ is the function like η but with x dropped from its domain.

Method calls of the form $x := e.m(\overline{e})$ are dynamically bound: the method meaning is determined by $loctype\,\ell$ in the semantic definition, where ℓ is the value of e. By typing, $loctype\,\ell \leq D$ and $pars(m, loctype\,\ell) = pars(m, D)$.

Semantics of commands

$$
\begin{array}{lcl}
[\![\Gamma \vdash x := e]\!]\mu(h,\eta) & = & \mathsf{let}\ d = [\![\Gamma \vdash e : T]\!](h,\eta)\ \mathsf{in}\ (h, [\eta \mid x \mapsto d]) \\
[\![\Gamma \vdash e_1.f := e_2]\!]\mu(h,\eta) & = & \mathsf{let}\ \ell = [\![\Gamma \vdash e_1 : (\Gamma\,\mathsf{self})]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ \ell = nil\ \mathsf{then}\ \bot\ \mathsf{else} \\
 & & \mathsf{let}\ d = [\![\Gamma \vdash e_2 : U]\!](h,\eta)\ \mathsf{in} \\
 & & ([h \mid \ell \mapsto [h\ell \mid f \mapsto d]], \eta) \\
[\![\Gamma \vdash x := \mathbf{new}\ B]\!]\mu(h,\eta) & = & \mathsf{let}\ \ell = fresh(B,h)\ \mathsf{in} \\
 & & \mathsf{let}\ h_0 = [h \mid \ell \mapsto [fields\,B \mapsto defaults]]\ \mathsf{in} \\
 & & (h_0, [\eta \mid x \mapsto \ell]) \\
[\![\Gamma \vdash x := e.m(\overline{e})]\!]\mu(h,\eta) & = & \mathsf{let}\ \ell = [\![\Gamma \vdash e : D]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ \ell = nil\ \mathsf{then}\ \bot\ \mathsf{else} \\
 & & \mathsf{let}\ \overline{x} = pars(m,D)\ \mathsf{in} \\
 & & \mathsf{let}\ \overline{d} = [\![\Gamma \vdash \overline{e} : \overline{U}]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{let}\ \eta_1 = [\overline{x} \mapsto \overline{d}, \mathsf{self} \mapsto \ell]\ \mathsf{in} \\
 & & \mathsf{let}\ (h_1, d_1) = \mu(loctype\,\ell)m(h,\eta_1)\ \mathsf{in} \\
 & & (h_1, [\eta \mid x \mapsto d_1]) \\
[\![\Gamma \vdash S_1;\ S_2]\!]\mu(h,\eta) & = & \mathsf{let}\ (h_1,\eta_1) = [\![\Gamma \vdash S_1]\!]\mu(h,\eta)\ \mathsf{in} \\
 & & [\![\Gamma \vdash S_2]\!]\mu(h_1,\eta_1) \\
[\![\Gamma \vdash \mathbf{if}\ e\ \mathbf{then}\ S_1\ \mathbf{else}\ S_2\ \mathbf{fi}]\!]\mu(h,\eta) & = & \mathsf{let}\ b = [\![\Gamma \vdash e : \mathbf{bool}]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{if}\ b\ \mathsf{then}\ [\![\Gamma \vdash S_1]\!]\mu(h,\eta)\ \mathsf{else}\ [\![\Gamma \vdash S_2]\!]\mu(h,\eta) \\
[\![\Gamma \vdash T\ x := e\ \mathbf{in}\ S]\!]\mu(h,\eta) & = & \mathsf{let}\ d = [\![\Gamma \vdash e : U]\!](h,\eta)\ \mathsf{in} \\
 & & \mathsf{let}\ \eta_1 = [\eta \mid x \mapsto d]\ \mathsf{in} \\
 & & \mathsf{let}\ (h_1,\eta_2) = [\![(\Gamma, x:T) \vdash S]\!]\mu(h,\eta_1)\ \mathsf{in} \\
 & & (h_1, (\eta_2 \upharpoonright x))
\end{array}
$$

Semantics of method declaration

Suppose M is a method declaration in $CT(C)$, with $M = T\ m(\overline{T}\ \overline{x})\{S\}$. Its meaning $[\![M]\!]$ is the total function $[\![Menv]\!] \to [\![C, \overline{x}, \overline{T} \to T]\!]$ defined by

$$
\begin{array}{rl}
[\![M]\!]\mu(h,\eta) = & \mathsf{let}\ \eta_1 = [\eta \mid \mathsf{result} \mapsto default]\ \mathsf{in} \\
 & \mathsf{let}\ (h_0,\eta_0) = [\![\overline{x} : \overline{T}, \mathsf{self} : C, \mathsf{result} : T \vdash S]\!]\mu(h,\eta_1)\ \mathsf{in} \\
 & (h_0, \eta_0\ \mathsf{result})
\end{array}
$$

For precision in the semantics of a method inherited in C from B we make an

explicit definition for the domain-restriction of a method meaning in $[\![B, \overline{x}, \overline{T} \to T]\!]$ to the global states (h, η) in $[\![Heap \otimes \overline{x} : \overline{T}, \mathsf{self} : C]\!]$.

***Definition* 5.5 (*restr*)** For $d \in [\![B, \overline{x}, \overline{T} \to T]\!]$ and $C \leq B$, define $restr(d, C)$, an element of $[\![C, \overline{x}, \overline{T} \to T]\!]$, by $restr(d, C)(h, \eta) = d(h, \eta)$. □

Semantics of class table and its approximation chain μ_j

The semantics of a well formed class table CT, written $[\![CT]\!]$, is the least upper bound of the ascending chain $\mu \in \mathbb{N} \to [\![Menv]\!]$ defined as follows.

$$\begin{array}{lll} \mu_0\, C\, m &= \lambda(h, \eta) \cdot \bot & \text{if } m \text{ is declared or inherited in } C \\ \mu_{j+1}\, C\, m &= [\![M]\!]\mu_j & \text{if } m \text{ is declared as } M \text{ in } C \\ \mu_{j+1}\, C\, m &= restr((\mu_{j+1}\, B\, m), C) & \text{if } m \text{ is inherited in } C \text{ from } B \end{array}$$

Remark* 5.6 (*On proofs) We give some proofs in considerable detail. To avoid repetition, we use the same identifiers as in the relevant semantic definition for each case —often different from those in the statement of the result being proved— taking care to avoid ambiguity. This saves explicit introduction of the identifiers or mention of the ranges and scopes of quantification. But it requires the reader to keep an eye on the semantic clauses. Often, without remark, we consider only the case where the outcome and various intermediate values are non-$\bot$.

Lemma 5.7 (semantics is well defined and typed) Let CT be well formed.

(1) If $C \leq B$ then for any Γ with $\mathsf{self} \notin dom\,\Gamma$ we have $[\![Heap \otimes \Gamma, \mathsf{self} : C]\!] \subseteq [\![Heap \otimes \Gamma, \mathsf{self} : B]\!]$.
(2) If $\Gamma \vdash e : T$ then $[\![\Gamma \vdash e : T]\!] \in [\![Heap \otimes \Gamma]\!] \to [\![T_\bot]\!]$.
(3) If $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ and $d = [\![\Gamma \vdash e : T]\!](h, \eta)$ with $d \neq \bot$ then $(h, d) \in [\![Heap \otimes T]\!]$.
(4) If $\Gamma \vdash S$ then $[\![\Gamma \vdash S]\!] \in [\![Menv]\!] \to [\![Heap \otimes \Gamma]\!] \to [\![(Heap \otimes \Gamma)_\bot]\!]$.
(5) $[\![CT]\!]$ is well defined and is an element of $[\![Menv]\!]$.

The proof is lengthy because the language is big and the semantic domains impose several invariants so that definedness amounts to type safety. The limit construction of $[\![CT]\!]$ can be justified using a straightforward characterization of the ordering on $[\![Menv]\!]$. The proof is sketched in the technical report [Banerjee and Naumann 2004a] and has been machine checked using PVS [Naumann 2005].

Garbage collection is not modeled in our semantics. As a consequence, the domain of the heap does not shrink and this fact simplifies some definitions in the sequel (mainly Def. 6.3 of partition extension and Def. 10.3 of coupling with bijection on allocated locations). To be precise, say that a method environment μ is *heap-extending* if $\mu\, C\, m\, (h, \eta) = (h_0, d)$ implies $dom\, h \subseteq dom\, h_0$ (for all C, m, h, η, h_0, d).

Lemma 5.8 (heap domain extended) (1) If $[\![\Gamma \vdash S]\!]\mu(h, \eta) = (h_0, \eta_0)$ and μ is heap-extending then $dom\, h \subseteq dom\, h_0$ (for all S, h, η, h_0, η_0).
(2) $[\![CT]\!]$ is heap-extending, as is each μ_j in its approximation chain.

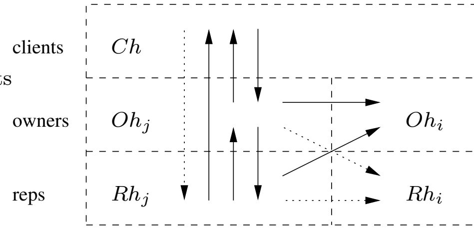

Fig. 6. Confinement scheme for island j. Dashed boxes are partition blocks. Solid lines indicate allowed references and dotted lines indicate prohibited ones. There is no restriction within blocks.

The proof of (1) is by induction on S; it amounts to checking that no command removes locations from the domain of the heap. The proof of (2) is by induction on the approximation chain, using (1) for declared methods.

6. CONFINEMENT RAMIFIED

This section formalizes the semantic notion of confinement discussed in Sect. 3.3.

6.1 Confinement of states

As discussed in Sect. 3.2 we assume that class names *Own* and *Rep* are given, such that $Own \not\lesseqgtr Rep$ and thus $[\![Own]\!] \cap [\![Rep]\!] = \{nil\}$.

We say heaps h_1 and h_2 are *disjoint* if $dom\, h_1 \cap dom\, h_2 = \varnothing$. Let $h_1 * h_2$ be the union of h_1 and h_2 if they are disjoint, and undefined otherwise.

We shall partition the heap as $h = Ch * \ldots$ where Ch contains client objects and the rest is partitioned into islands of the form $Oh * Rh$ consisting of a singleton heap Oh with an owner object and a heap Rh of its representation objects. In such a partition, the heaps Ch, Oh, and Rh need not be closed. An example is Fig. 5 in Sect. 3.3; the general scheme is depicted in Fig. 6. Our use of the word "partition" is slightly non-standard: we allow the blocks Rh_i and Ch to be empty.

***Definition* 6.1 (admissible partition)** An *admissible partition* of heap h is a set of pairwise disjoint heaps $Ch, Oh_1, Rh_1, \ldots, Oh_k, Rh_k$, for $k \geq 0$, with

$$h = Ch * Oh_1 * Rh_1 * \ldots * Oh_k * Rh_k$$

and for all i $(1 \leq i \leq k)$

—$dom\ Oh_i \subseteq locs(Own{\downarrow})$ and $size(dom\ Oh_i) = 1$ (owner blocks)
—$dom\ Rh_i \subseteq locs(Rep{\downarrow})$ (rep blocks)
—$dom\ Ch \cap locs(Own{\downarrow}) = \varnothing$ and $dom\ Ch \cap locs(Rep{\downarrow}) = \varnothing$ (client blocks)

***Definition* 6.2 (confined heap, confining partition, $\not\hookrightarrow$, $\not\hookrightarrow^{\overline{f}}$)** To say that no object in h_1 contains a reference to an object in h_2, we define $\not\hookrightarrow$ by

$$h_1 \not\hookrightarrow h_2 \ \Leftrightarrow\ \forall \ell \in dom\, h_1 \cdot rng(h_1\, \ell) \cap dom\, h_2 = \varnothing \ .$$

To say that no object in h_1 contains a reference to an object in h_2 *except via a field in* $\overline{f}$, we define $\not\hookrightarrow^{\overline{f}}$ by

$$h_1 \not\hookrightarrow^{\overline{f}} h_2 \ \Leftrightarrow\ \forall \ell \in dom\, h_1 \cdot rng((h_1\, \ell){\lfloor}\overline{f}) \cap dom\, h_2 = \varnothing \ .$$

A heap h *is confined*, written *conf* h, iff it has a confining partition. A *confining partition* is an admissible partition such that for all j, i with $j \neq i$ we have

(1) $Ch \not\mapsto Rh_j$ (clients do not point to reps)

(2) $Oh_j \not\mapsto Rh_i$ (owners do not share reps)

(3) $Oh_j \not\mapsto^{\overline{g}} Rh_j$ where $\overline{g} = dom(dfields(Own))$ (reps are private to Own)

(4) $Rh_j \not\mapsto Rh_i$ (reps are confined to their islands)

A heap may have several admissible partitions, because there is no inherent order on islands and because unreachable reps can be put in any island. The definitions and results do not depend on choice of partition. We have not found a workable formulation that determines unique partitions. To describe the effect of confined commands on partitions we use the following.

***Definition* 6.3 (extension of confining partition, $\trianglelefteq$)** Define $h \trianglelefteq h_0$ iff h is confined and for any confining partition of h,

$$h = Ch * Oh_1 * Rh_1 * \ldots * Oh_k * Rh_k \quad (k \geq 0),$$

there is a confining partition of h_0,

$$h_0 = Ch^0 * Oh_1^0 * Rh_1^0 * \ldots * Oh_n^0 * Rh_n^0 \ ,$$

that is an extension in the sense that it satisfies the following:

—$n \geq k$

—$dom(Ch) \subseteq dom(Ch^0)$

—$dom(Oh_j) \subseteq dom(Oh_j^0)$ for all $j \leq k$

—$dom(Rh_j) \subseteq dom(Rh_j^0)$ for all $j \leq k$ □

Note that $h \trianglelefteq h_0$ implies $dom\, h \subseteq dom\, h_0$.

Confinement of a store depends on the class in which it may occur. For owners and reps it depends on the domain of the heap as well.

***Definition* 6.4 (confined store, global state)** Let h be a confined heap and η be a store in $[\![\Gamma, \mathsf{self}: C]\!]$ for some Γ. We say η is *confined in h for C* iff

(1) $C \not\leq Rep \wedge C \not\leq Own \Rightarrow rng\, \eta \cap locs(Rep{\downarrow}) = \varnothing$

(2) $C \leq Own \Rightarrow rng\, \eta \cap locs(Rep{\downarrow}) \subseteq dom(Rh_j)$
for some confining partition and j with $\eta\, \mathsf{self} \in dom(Oh_j)$

(3) $C \leq Rep \Rightarrow rng\, \eta \cap locs(Rep{\downarrow}) \subseteq dom(Rh_j)$
for some confining partition and j with $\eta\, \mathsf{self} \in dom(Rh_j)$

A global state (h, η) is *confined*, written $conf\, C\,(h, \eta)$, iff h is confined and η is confined in h for C. □

Apropos the examples in Sect. 2.1, take Rep to be Bool and suppose the sequence z := **new** OBool; w := z.bad() occurs in a method of some client class. Executed in a confined initial state, the state after assignment of a new OBool to z is still confined. The assignment to w then yields a state where the heap is confined but the client's store is not.

6.2 Confinement of commands and methods

A confined command is one that preserves confinement of global states. Because command meanings depend on the method environment and expression meanings, confinement for those is formalized first. We need to ensure that a method call yields a heap confined for the caller. This is achieved using the condition $h \trianglelefteq h_0$ in the following Definition, together with Lemma 6.13 to follow.

***Definition* 6.5 (confined method environment)** Method environment μ is confined, written $\mathit{conf}\,\mu$, if and only if the following holds for all C and m with $\mathit{mtype}(m, C)$ defined. Let $\mathit{mtype}(m, C) = \overline{T}{\to}T$ and $\mathit{pars}(m, C) = \overline{x}$. For all $(h, \eta) \in [\![\mathit{Heap} \otimes \overline{x}:\overline{T}, \mathsf{self}:C]\!]$, if $\mathit{conf}\, C\,(h, \eta)$ and $\mu C m(h, \eta) \neq \bot$ then

(1) $C \not\leq \mathit{Rep} \Rightarrow h \trianglelefteq h_0 \wedge d \notin \mathit{locs}(\mathit{Rep}{\downarrow})$
(2) $C \leq \mathit{Rep} \Rightarrow h \trianglelefteq h_0 \wedge (d \in \mathit{locs}(\mathit{Rep}{\downarrow}) \Rightarrow d \in \mathit{dom}(\mathit{Rh}_j))$
for some confining partition $h_0 = \mathit{Ch} * \mathit{Oh}_1 * \mathit{Rh}_1 \ldots$
and j with $\eta\,\mathsf{self} \in \mathit{dom}(\mathit{Rh}_j)$

where $(h_0, d) = \mu C m(h, \eta)$. □

Note that the consequent $h \trianglelefteq h_0$ implies $\mathit{conf}\, h_0$, by definition of $\trianglelefteq$ using $\mathit{conf}\, h$ which follows from the antecedent $\mathit{conf}\, C\,(h, \eta)$. Also, a confined method environment is heap-extending in the sense defined just before Lemma 5.8. So in reasoning about commands we will not need to separately assume that the method environment is heap-extending.

Condition (1) in Def. 6.5 fails for method bad of the example in Sect. 2.2, regardless of whether the return type of bad is taken to be **Object** or Bool.

The conditions for confinement of expressions are like those for confined stores —after all, a store provides the meaning for the expression x. The conditions are somewhat different for confined method environments, because methods are public and can be called both by clients and from within an owner island. (In Sect. 9, Def. 6.5 is refined to allow module-scoped owner methods to return reps.) Also, confinement of commands does not explicitly require partition extension $h \trianglelefteq h_0$ like Def. 6.5 does, because it is a consequence of the other conditions (see Lemma 6.14).

***Definition* 6.6 (confined expression)** Let $C = \Gamma\,\mathsf{self}$. Expression $\Gamma \vdash e : T$ is confined iff for any (h, η), if $\mathit{conf}\, C\,(h, \eta)$ and $[\![\Gamma \vdash e : T]\!](h, \eta) \neq \bot$ then the following hold, where $d = [\![\Gamma \vdash e : T]\!](h, \eta)$.

(1) $C \not\leq \mathit{Rep} \wedge C \not\leq \mathit{Own} \Rightarrow d \notin \mathit{locs}(\mathit{Rep}{\downarrow})$
(2) $C \leq \mathit{Own} \Rightarrow (d \in \mathit{locs}(\mathit{Rep}{\downarrow}) \Rightarrow d \in \mathit{dom}(\mathit{Rh}_j))$
for some confining partition and j with $\eta\,\mathsf{self} \in \mathit{dom}(\mathit{Oh}_j)$
(3) $C \leq \mathit{Rep} \Rightarrow (d \in \mathit{locs}(\mathit{Rep}{\downarrow}) \Rightarrow d \in \mathit{dom}(\mathit{Rh}_j))$
for some confining partition and j with $\eta\,\mathsf{self} \in \mathit{dom}(\mathit{Rh}_j)$ □

***Definition* 6.7 (confined command)** Let $C = \Gamma\,\mathsf{self}$. Command $\Gamma \vdash S$ is confined iff

— $\mathit{conf}\,\mu \wedge \mathit{conf}\, C\,(h, \eta) \wedge [\![\Gamma \vdash S]\!]\mu(h, \eta) \neq \bot \Rightarrow \mathit{conf}\, C\,(h_0, \eta_0)$, for any μ and any (h, η), where $(h_0, \eta_0) = [\![\Gamma \vdash S]\!]\mu(h, \eta)$
—if S is a method call then it has confined arguments (see below). □

Confinement of arguments means that the store, η_1, passed in the semantics of method call is confined for the callee.

***Definition* 6.8 (confined arguments)** Let $C = \Gamma\,\mathsf{self}$. A call $\Gamma \vdash x := e.m(\overline{e})$ has *confined arguments* provided the following holds. Suppose $\overline{U}$ is the static type of $\overline{e}$ and D the static type of e. For any (h,η) with $conf\,C\,(h,\eta)$, let

$$\overline{d} = [\![\Gamma \vdash \overline{e} : \overline{U}]\!](h,\eta) \qquad \ell = [\![\Gamma \vdash e : D]\!](h,\eta) \qquad \eta_1 = [\overline{x} \mapsto \overline{d}, \mathsf{self} \mapsto \ell] \ .$$

If $\ell \neq \bot$, $\ell \neq nil$, and $\bot$ does not occur in $\overline{d}$ then $conf\,(loctype\,\ell)\,(h,\eta_1)$. □

A purely semantic formulation would call class table CT confined just if $[\![CT]\!]$ is a confined method environment. But under simple restrictions, confinement of $[\![CT]\!]$ follows from confinement of method bodies. Thus we choose the following.

***Definition* 6.9 (confined class table)** Class table CT is confined iff for every C and every m with $mtype(m,C) = \overline{T} \rightarrow T$ the following hold.

(1) If m is declared in C by $T\ m(\overline{T}\ \overline{x})\{S\}$ then S and all its constituents are confined.
(2) If $C \leq Own$ then $T \not\leq Rep$.
(3) If m is inherited in Own from some $B > Own$ then $\overline{T} \not\leq Rep$.
(4) No method m is inherited in Rep from any $B > Rep$. □

In Sect. 10 we add module-scoped methods on which condition (2) need not be imposed. This condition ensures that owner methods do not return reps, which is not ensured by confinement of the method body. Condition (4) is needed because confinement of a method inherited from $B > Rep$ depends on the arguments, including self, being confined at B where reps are disallowed. Invocation of such a method on an object of type Rep (or a subclass) would yield a store with self a rep. A more refined restriction is to disallow inheritance into Rep only for methods which leak self; see Sect. 12.2.

***Example* 6.10** Condition (2) precludes the bad method of Sect. 2.1, for both return types **Object** and Bool. Except for this, all examples in Sect. 2 yield confined class tables (e.g., the well formed class table obtained by combining Figs. 2 and 3). One way to prove confinement for these examples is to check that they are safe according to the static analysis of Sect. 11. □

6.3 Properties of confinement

We need a number of results about confinement. The most important is that the semantics of a confined class table is a confined method environment (Theorem 6.15). This depends on Lemma 6.14 which says that confined commands extend heap partitions, provided that method meanings have this property.

Lemma 6.11 If T is **bool** or **unit**, then every $\Gamma \vdash e : T$ is confined.

PROOF. Direct from the definitions: confinement only pertains to locations. □

Lemma 6.12 Suppose $rng\,\eta \cap locs(Rep{\downarrow}) = \varnothing$ and $C \leq B$. Then for any h and any $\eta \in [\![\Gamma, \mathsf{self} : C]\!]$ we have $conf\,C\,(h,\eta)$ iff $conf\,B\,(h,\eta)$.

PROOF. Straightforward case analysis. □

Lemma 6.13 If $conf\,C\,(h,\eta)$ and $h \trianglelefteq h_0$ then $conf\,C\,(h_0,\eta)$.

Proof. Straightforward from the definitions. □

Although confining partitions are not unique, a given confining partition of an initial state can be extended to one on the final state for any command.

Lemma 6.14 (extension by commands) Suppose $\Gamma \vdash S$ is confined and all its constituents are confined. Let $C = \Gamma\,\mathsf{self}$. For any μ, h, η with $conf\,\mu$ and $conf\,C\,(h,\eta)$, we have $[\![\Gamma \vdash S]\!]\mu(h,\eta) \neq \bot \Rightarrow h \trianglelefteq h_0$, where $(h_0,-) = [\![\Gamma \vdash S]\!]\mu(h,\eta)$.

Proof. By structural induction on S. Let $C = \Gamma\,\mathsf{self}$. We assume a confining partition $h = Ch * Oh_1 * Rh_1 * \ldots * Oh_k * Rh_k$ is given (k may be 0, i.e., there need not be any islands). We show how to construct confining partition $h_0 = Ch^0 * Oh^0_1 * Rh^0_1 * \ldots$ that extends the given one.

Case $\Gamma \vdash e_1.f := e_2$. From $[\![\Gamma \vdash e_1.f := e_2]\!]\mu(h,\eta) \neq \bot$ and Lemma 5.7(3) we have that $\ell \in dom\,h$ where $\ell = [\![\Gamma \vdash e_1 : C]\!](h,\eta)$. By semantics, $h_0 = [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]$. We partition h_0 using the given partition for h. That is, the domain for each block of the updated heap h_0 is the same as the corresponding block for h. Clearly this extends the partition for h. To show that this partition is confining for h_0, it suffices to show that the update of $h\ell f$ to d satisfies the confinement property for ℓ. We argue by cases on $loctype\,\ell$

- $loctype\,\ell \not\leq Own \wedge loctype\,\ell \not\leq Rep$. Then Def. 6.2(1) applies; it requires $d \notin locs(Rep{\downarrow})$. By typing, $loctype\,\ell \leq C$, so $C \not\leq Own \wedge C \not\leq Rep$. Thus by confinement of e_1 (a constituent of $e_1.f := e_2$ and therefore confined by hypothesis), we have by Def. 6.6(1) that $d \notin locs(Rep{\downarrow})$.
- $loctype\,\ell \leq Own$. Def. 6.2(2) and (3) apply here. Letting j be the index of the island with $\{\ell\} = dom(Oh^0_j) = dom(Oh_j)$, we must show both $Oh^0_j \not\rightsquigarrow Rh^0_i$ (for $i \neq j$) and $Oh^0_j \not\rightsquigarrow^{\overline{g}} Rh^0_j$. By typing, $loctype\,\ell \leq C$, so $C \leq Own$ or $Own \leq C$ by the tree property of $\leq$. We argue by cases on C.
  - $Own < C$. By $Own \not\geq Rep$, we have $C \not\leq Rep$ so confinement of e_2 at C yields $d \notin locs(Rep{\downarrow})$. Thus $Oh^0_j \not\rightsquigarrow Rh^0_i$ and $Oh^0_j \not\rightsquigarrow^{\overline{g}} Rh^0_j$.
  - $C \leq Own$. By confinement of e_2, if $d \in locs(Rep{\downarrow})$ then $d \in dom(Rh^0_j)$ so $Oh^0_j \not\rightsquigarrow Rh^0_i$ for $i \neq j$. If $C = Own$ then, by the typing rule for field update, f is in the private fields $\overline{g}$ of Own, so the update cannot violate $Oh^0_j \not\rightsquigarrow^{\overline{g}} Rh^0_j$. If $C < Own$ then $d \notin locs(Rep{\downarrow})$ because if d is a rep then there would be no confining partition, contradicting confinement of h_0 which holds by confinement of S.
- $loctype\,\ell \leq Rep$. Def. 6.2(4) applies in this case: we need to show $Rh^0_j \not\rightsquigarrow Rh^0_i$ where $i \neq j$ and j is the island for ℓ in the partition of h. By typing, $loctype\,\ell \leq C$, hence $C \leq Rep$ or $Rep < C$. But if $Rep < C$ then $C \not\geq Own$ and the confinement condition for e_1 (Def. 6.6(1)) at C contradicts $loctype\,\ell \leq Rep$, so we have $C \leq Rep$. Now confinement of e_2 yields $d \in locs(Rep{\downarrow}) \Rightarrow d \in dom(Rh_j)$. This proves $Rh^0_j \not\rightsquigarrow Rh^0_i$, because $dom(Rh_j) = dom(Rh^0_j)$.

Case $\Gamma \vdash x := \mathbf{new}\ B$. In the semantic definition, $h_0 = [h \mid \ell \mapsto [fields\,B \mapsto defaults]]$ where $\ell = fresh(B,h)$. Define $Bh = [\ell \mapsto [fields\,B \mapsto defaults]]$ so

$h_0 = h * Bh$. Next, we argue that $h \trianglelefteq h_0$. Because h is closed, ℓ is not in the range of any object state in h. To construct an extending partition it suffices to deal with the new object, as its addition cannot violate confinement of existing objects. (This would not be the case if dangling pointers were allowed, without further restrictions on *fresh*.) We define the extension and argue by cases on B.

—$B \not\leq \mathit{Own} \wedge B \not\leq \mathit{Rep}$. For a confining partition of h_0 we extend that for h by defining $Ch^0 = Ch * Bh$ and using the given partition of owner islands. Because *defaults* contains no locations, this is a confining partition.

—$B \leq \mathit{Own}$. We extend the partition by adding an island $Oh^0_{k+1} * Rh^0_{k+1}$ with $Oh^0_{k+1} = Bh$ and $Rh^0_{k+1} = \varnothing$. This is a confining partition because *defaults* has no locations.

—$B \leq \mathit{Rep}$. We can obtain a confining extension by adding Bh to any of the Rh_i, as *defaults* has no locations.

Case $\Gamma \vdash x := e.m(\overline{e})$. As $e.m(\overline{e})$ is confined, its argument values are confined. Thus we can obtain the result directly from the semantics of $e.m(\overline{e})$ and confinement of μ —which explicitly stipulates $h \trianglelefteq h_0$.

The remaining cases are straightforward. □

Theorem 6.15 Suppose that CT is confined. Then the semantics $[\![CT]\!]$ is confined, as is each μ_j in the approximation chain used to define it.

The proof uses fixpoint induction, which is only sound for admissible predicates [Mitchell 1996], i.e., those closed under limits of ascending chains. For confinement of method environments the definition is given pointwise, ultimately unfolding to the property that the semantics of each method body preserves confinement. This definition, as well as the one for the simulation $\mathcal{R}$ later, is in the usual form of logical relations. By the structure of the definition, and continuity of the semantics, the property is an admissible predicate.[9]

Proof. Confinement of $[\![CT]\!]$ follows by fixpoint induction from confinement of μ_i for all i, which we show by induction on i. The base case holds because $\mu_0 Cm = \lambda(h,\eta) \cdot \bot$, for any C, m, and this is confined by definition.

For the induction step, suppose *conf* μ_i, to show *conf* μ_{i+1}. Consider an arbitrary m. We argue for all C with $\mathit{mtype}(m, C)$ defined, by induction on method depth (Def. 4.6) of C for m. The base case is C such that $\mathit{depth}(m, C) = 0$. In this case, $CT(C)$ has a declaration

$$T\ m(\overline{T}\ \overline{x})\{S\}\ .$$

Suppose *conf* $C\,(h,\eta)$ and $\mu_{i+1}Cm(h,\eta) \neq \bot$. Let $(h_0, d) = \mu_{i+1}Cm(h,\eta)$, which by definition of μ_{i+1} is obtained as

$$\begin{aligned} \eta_1 &= [\eta \mid \mathsf{result} \mapsto \mathit{default}] \\ (h_0, \eta_0) &= [\![\overline{x} : \overline{T}, \mathsf{self} : C, \mathsf{result} : T \vdash S]\!]\mu_i(h, \eta_1) \\ d &= \eta_0\, \mathsf{result} \end{aligned}$$

[9]A rigorous proof can be given using a straightforward characterization of ascending chains in $[\![\mathit{Menv}]\!]$. For a similar notion of simulation and the same semantic model, this has been machine checked [Naumann 2005].

Default values do not violate confinement so $\mathit{conf}\, C\,(h, \eta_1)$. As CT is confined, S and its constituents are confined. By Lemma 6.14 we have $h \trianglelefteq h_0$, so by Lemma 6.13 we have $\mathit{conf}\, C\,(h_0, \eta)$. To show the confinement condition for $\mu_{i+1} C m$ it remains to deal with the result value d. We have $\mathit{conf}\, C\,(h_0, \eta_0)$ by confinement of S. We argue by cases on C.

—$C \not\leq \mathit{Own} \wedge C \not\leq \mathit{Rep}$. We need $d \notin \mathit{locs}(\mathit{Rep}{\downarrow})$, for Def. 6.5(1), and this follows from $\mathit{conf}\, C\,(h_0, \eta_0)$ by Def. 6.4(1).

—$C \leq \mathit{Own}$. We need $d \notin \mathit{locs}(\mathit{Rep}{\downarrow})$, and since by typing we have $d \in [\![T_\bot]\!]$, Def. 6.9(2) ensures $T \not\geq \mathit{Rep}$ and hence $d \notin \mathit{locs}(\mathit{Reps}{\downarrow})$. (Note that semantic confinement of η_0 at $C \leq \mathit{Own}$ allows reps, so it is not enough for this case).

—$C \leq \mathit{Rep}$. Then we need $d \in \mathit{locs}(\mathit{Rep}{\downarrow})$ to imply that d is in the domain of Rh_j for some partition and island j such that $\eta\,\mathsf{self} \in \mathit{dom}(Rh_j)$. This follows from $\mathit{conf}\, C\,(h_0, \eta_0)$ by Def. 6.4(3).

This concludes the base case of the induction on depth.

For the induction step, i.e., $\mathit{depth}(m, C) > 0$, m may be inherited or declared in C. If it is declared in C the argument is the same as for the case $\mathit{depth}(m, C) = 0$ above. Suppose m is inherited in C from B. Now $\mu_{i+1} C m = \mathit{restr}((\mu_{i+1} B m), C)$ by definition of μ_{i+1}. By induction on depth $\mu_{i+1} B m$ satisfies the confinement condition for m, B. To show the condition for $\mu_{i+1} C m$, suppose $\mathit{conf}\, C\,(h, \eta)$. We claim that $\mathit{conf}\, B\,(h, \eta)$. Using the claim, we argue as follows. If $\mu_{i+1} B m(h, \eta) \neq \bot$, let $(h_0, d) = \mu_{i+1} B m(h, \eta)$. By induction on depth we have $\mathit{conf}\, B\,(h_0, \eta)$ and $h \trianglelefteq h_0$. By Lemma 6.13 we obtain $\mathit{conf}\, C\,(h_0, \eta)$. It remains to show the confinement condition for d and to prove the claim. We argue by cases on C.

In the following non-rep cases, the claim holds by Lemma 6.12. To apply the Lemma, we just need to show that $\mathit{rng}\,\eta \cap \mathit{locs}(\mathit{Rep}{\downarrow}) = \varnothing$.

—$C \not\leq \mathit{Own} \wedge C \not\leq \mathit{Rep}$. In this case, we have $\mathit{rng}\,\eta \cap \mathit{locs}(\mathit{Rep}{\downarrow}) = \varnothing$ by confinement of η at C, Def. 6.4(1).

—$C \leq \mathit{Own} < B$. Then Own inherits m from $B > \mathit{Own}$, so by confinement of the class table, Def. 6.9(3), we have $\overline{T} \not\geq \mathit{Rep}$. Also, $\mathit{Own} \not\geq \mathit{Rep}$, so by Lemma 5.4 we have no reps in $\mathit{rng}\,\eta$.

In the preceding cases, the condition imposed on d by Def. 6.5(1) for class C is $d \notin \mathit{locs}(\mathit{Rep}{\downarrow})$. But this same condition is imposed for class B, and it holds by induction on depth. For the remaining cases we prove the claim $\mathit{conf}\, B\,(h, \eta)$ as follows.

—$C < B \leq \mathit{Own}$. Both B and C impose the same condition (Def. 6.4(2)).

—$C < B \leq \mathit{Rep}$. Both C and B impose the same conditions on η (Def. 6.4(3)).

In these two cases the requirement for d at C, Def. 6.5(2) or (1), is the same as for B, so it holds by induction on depth.

The case $C \leq \mathit{Rep} < B$ cannot occur in a confined class table. If m is inherited in $C \leq \mathit{Rep}$ from B then it is inherited in Rep from B, and this is explicitly disallowed in Def. 6.9(4). □

7. FIRST ABSTRACTION THEOREM

This section formulates and proves the central result of the paper. First, we make precise the idea of comparing two class tables that differ only in their implementation of class Own. Then we define local coupling: a relation between single instances of class Own for the two implementations. This induces the coupling relations for other data types, for heaps containing multiple instances of Own, and for method meanings. Related method meanings have the simulation property: if initial states are coupled, then so are outcomes. The main theorem says that if methods of Own have the simulation property then so do all methods of all classes.

7.1 Comparing class tables

We compare two implementations of a designated class Own. They can have completely different declarations, so long as methods of the same signatures are present —declared or inherited— in both. They can use different reps, distinguished by class name Rep for one implementation and Rep' for the other. We allow $Rep = Rep'$. For simplicity, we assume that both Rep and Rep' are in each of the two compared class tables.

***Definition* 7.1 (comparable class tables, non-rep class)** Suppose class names Own, Rep, Rep' are given, such that $Own \not\lesseqgtr Rep$ and $Own \not\lesseqgtr Rep'$. We say C is a *non-rep* class iff $C \not\leq Rep$ and $C \not\leq Rep'$. Well formed class tables CT and CT' are *comparable* provided the following hold.

(1) CT and CT' are identical except for their values on Own. (In particular, $CT(Rep) = CT'(Rep)$ and $CT(Rep') = CT'(Rep')$.)
We write $\vdash, \vdash'$ for the typing relations determined by CT, CT' respectively, and similarly for the auxiliary functions, such as $mtype, mtype'$. We also write $[\![-]\!], [\![-]\!]'$ for the respective semantics and assume that the same allocator, $fresh$, is used for both $[\![-]\!]$ and $[\![-]\!]'$.

(2) $super\, Own = super'\, Own$.

(3) For any m, either $mtype(m, Own)$ and $mtype'(m, Own)$ are both undefined or both are defined and equal. □

***Example* 7.2** Let CT be given by Figs. 2 and 3, with Node2 from Fig. 4. Let CT' be given by Figs. 4 and 3 together with Observer from Fig. 2. These are comparable, taking Rep to be Node and Rep' to be Node2. □

Note that the typing relations $\Gamma \vdash -$ and $\Gamma \vdash' -$ are the same except if $\Gamma\,\mathsf{self} = Own$. Similarly, $dfields\, C = dfields'\, C$ unless $C = Own$.

Instead of condition (3), one could require that $CT(Own)$ and $CT'(Own)$ declare the same methods. But that would disallow some situations that occur in practice. Suppose class C extends B by adding a method m implemented using calls to methods inherited from B. This might be the easiest way to achieve desired functionality for m, but there could be an alternative data structure that is more efficient for m and for the methods of B. An alternative implementation of C could add that data structure and override the methods of B to use it. One can argue that the program is poorly designed, e.g., because space for attributes of B is wasted in C objects. Better designs are possible. Nonetheless, such examples do arise

in practice; allowing them complicates the proof of Theorem 7.20 but none of the other results. The main consequence we need from condition (3) is the following.

Lemma 7.3 If $\mathit{mtype}(m, C)$ is defined then $\mathit{depth}(m, C) = \mathit{depth}'(m, C)$.

PROOF. Straightforward. □

One can imagine a theory in which an owner subclass has different declarations in CT and CT'. But we are concerned with an abstraction provided by a single class rather than by a collection of classes, so we require $CT(C) = CT'(C)$ even for $C < Own$. In Sect. 7.3 we impose a restriction on owner subclasses that is needed for the first abstraction theorem. The issue is explored in Sect. 9 and the restriction lifted in Sect. 10.

7.2 Coupling relations and simulation

The definitions are organized as follows. A *local coupling* $\mathcal{L}$ is a suitable relation on islands. This induces a family of *coupling relations*, $\mathcal{R}\,\theta$ for each semantic category θ. Then comes the definition of *simulation*, a coupling that is preserved by all methods of Own and established by the constructor.

***Definition* 7.4 (local coupling, $\mathcal{L}$)** Given comparable class tables, a *local coupling* is a binary relation $\mathcal{L}$ on heaps —not necessarily closed— such that the following holds: For any h, h', if $\mathcal{L}\,h\,h'$ then there is a location ℓ with $\mathit{loctype}\,\ell \leq Own$ and partitions $h = Oh * Rh$ and $h' = Oh' * Rh'$ such that

(1) $\mathit{dom}\ Oh = \{\ell\} = \mathit{dom}\ Oh'$
(2) $\mathit{dom}(Rh) \subseteq \mathit{locs}(Rep{\downarrow})$ and $\mathit{dom}(Rh') \subseteq \mathit{locs}(Rep'{\downarrow})$
(3) $h\ell f = h'\ell f$ for all $f \in \mathit{dom}(\mathit{fields}(\mathit{loctype}\,\ell))$ with $f \notin \overline{g}$ and $f \notin \overline{g}'$, where $\overline{g} = \mathit{dom}(\mathit{dfields}(Own))$ and $\overline{g}' = \mathit{dom}(\mathit{dfields}'(Own))$ □

Example 7.7 below shows why we allow $\mathcal{L}$ to act on heaps that are not closed.

Although $\mathcal{L}$ is unconstrained for the private fields and reps, condition (3) determines it for fields of proper subclasses of Own (while allowing $\mathcal{L}$ to depend on these fields). Once we have defined the induced relation $\mathcal{R}$, item (3) will be equivalent to the condition $\mathcal{R}\ (\mathit{type}(f, \mathit{loctype}\,\ell))\ (h\ell f)\ (h'\ell f)$. Because CT and CT' are well formed, the declared field names $\overline{g}$ and $\overline{g}'$ of $CT(Own)$ and $CT'(Own)$ are not declared in proper subclasses or superclasses of Own. So f in (3) ranges over fields including those declared in proper superclasses and subclasses of Own.

***Example* 7.5** Sect. 2.2 discusses this coupling relation:

$$\mathsf{o.g} \neq \mathit{nil} \neq \mathsf{o'.g} \land \mathsf{o.g.f} = \neg(\mathsf{o'.g.f}) \ .$$

For this example we take both Rep and Rep' to be Bool, and Own to be OBool. The displayed formula can be interpreted as local coupling relation $\mathcal{L}$ which relates h to h' just if

$$h = [\ell_1 \mapsto [g \mapsto \ell_2], \ell_2 \mapsto [f \mapsto d]] \quad \text{and} \quad h' = [\ell_1 \mapsto [g \mapsto \ell_3], \ell_3 \mapsto [f \mapsto \neg d]]$$

for some boolean d and locations ℓ_1 in $\mathit{locs}(\mathsf{OBool})$ and ℓ_2, ℓ_3 in $\mathit{locs}(\mathsf{Bool})$. We assume that the class table contains only Bool, OBool, and some client classes. If OBool had subclasses, the relation on their fields would be determined by condition (3) above. □

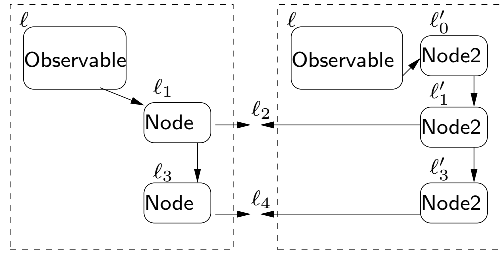

Fig. 7. Local coupling example: a related pair of islands. Labels indicate locations as described in Example 7.7. Note dangling pointers ℓ_2 and ℓ_4 and sentinel node ℓ'_0.

***Example* 7.6** Sect. 3.1 uses the formula $\mathsf{o.g} = \mathsf{o'.g} \wedge \mathsf{o.g}\ \mathbf{mod}\ 2 = 0$. This can be interpreted as the local coupling that relates h to h' just if there is some ℓ with $loctype\ \ell \leq \mathsf{A}$, h and h' have domain $\{\ell\}$, and $h\,\ell\,g = h'\,\ell\,g = 2 \times m$ for some integer $m \geq 0$. □

***Example* 7.7** The Observer examples show why we allow a local coupling $\mathcal{L}$ to relate non-closed heaps. Consider the version in Fig. 2. Here *Rep* is Node, *Own* is Observable, and there is a client class Observer. Fig. 5 illustrates two instances of this simple data structure. Fig. 4 gives code for an alternative version which uses an extra node as sentinel for the list. The sentinel does not point to an Observer. Fig. 7 depicts a corresponding pair of heaps for the two alternatives, using arrows without destination objects to indicate dangling pointers. Upon initialization of an Observable, there are no installed Observers, so for the version of Fig. 2 we should have $\mathsf{fst} = nil$. But in the alternative version, this should correspond to snt holding the location of a Node2 with $\mathsf{ob} = nil$ and $\mathsf{nxt} = nil$. This is established by the constructor in Fig. 4. An attempt at formalizing the correspondence is as follows:

$$(\mathsf{o.fst} = nil = \mathsf{o'.snt.nxt}) \vee (\mathsf{o.fst} \neq nil \neq \mathsf{o'.snt.nxt} \wedge \alpha(\mathsf{o.fst}) = \alpha'(\mathsf{o'.snt.nxt}))$$

where α, α' are functions that yield the list of locations in the ob fields of successive nodes. But how should this formula be interpreted if, say, $\mathsf{o'.snt} = nil$ or there is sharing such as a chain with cyclic tail? Separation logic [Reynolds 2001] offers a precise way to formulate such definitions but its development is at an early stage. We simply sketch the coupling in terms of semantics: $\mathcal{L}\,h\,h'$ iff either h and h' have the form

$$\begin{aligned} h &= [\,\ell \mapsto [\mathsf{fst} \mapsto nil]] \\ h' &= [\,\ell \mapsto [\mathsf{snt} \mapsto \ell'_0],\ \ell'_0 \mapsto [\mathsf{ob} \mapsto nil, \mathsf{nxt} \mapsto nil]] \end{aligned}$$

or they have the form

$$\begin{aligned} h &= [\,\ell \mapsto [\mathsf{fst} \mapsto \ell_1],\ \ell_1 \mapsto [\mathsf{ob} \mapsto \ell_2, \mathsf{nxt} \mapsto \ell_3],\ \ell_3 \mapsto [\mathsf{ob} \mapsto \ell_4, \mathsf{nxt} \mapsto \ldots], \ldots] \\ h' &= [\,\ell \mapsto [\mathsf{snt} \mapsto \ell'_0],\ \ell'_0 \mapsto [\mathsf{ob} \mapsto nil, \mathsf{nxt} \mapsto \ell'_1], \\ &\qquad \ell'_1 \mapsto [\mathsf{ob} \mapsto \ell_2, \mathsf{nxt} \mapsto \ell'_3],\ \ell'_3 \mapsto [\mathsf{ob} \mapsto \ell_4, \mathsf{nxt} \mapsto \ldots], \ldots] \end{aligned}$$

for some locations ℓ in $locs(\mathsf{Observable})$, $\ell_1, \ell_3, \ldots$ in $locs(\mathsf{Node})$, $\ell'_0, \ell'_1, \ell'_3, \ldots$ in $locs(\mathsf{Node2})$, and $\ell_2, \ell_4, \ldots$ in $locs(\mathsf{Observer}{\downarrow})$. Note that the owners are at the same location, ℓ, as are the referenced Observer objects at $\ell_2, \ell_4, \ldots$. No correspondence is required between locations $\ell_1, \ell_3, \ldots$ and $\ell'_0, \ell'_1, \ell'_3, \ldots$ of reps. □

A local coupling induces a relation $\mathcal{R}\,Heap$ on arbitrary heaps by the requirement that they have confining partitions such that islands can be put in correspondence

so that pairs are related by $\mathcal{L}$. The formal definition uses the induced relation $\mathcal{R}\ (state\, C)$ for object states of non-rep classes $C \not\leq Own$, and this in turn is defined in terms of $\mathcal{R}\, C$ for non-rep classes $C \not\leq Own$. For uniformity, we give the definition of $\mathcal{R}$ for all θ, but forcing the case for $\theta = state\, Own$ to be false. Aside from the ramifications of heap confinement, the definition is induced in the standard way for logical relations.

***Definition* 7.8 (coupling relation, $\mathcal{R}\,\theta$)** Given a local coupling relation $\mathcal{L}$, we define for each θ a relation $\mathcal{R}\,\theta \subseteq [\![\theta]\!] \times [\![\theta]\!]'$ as follows.

For heaps h, h', we define $\mathcal{R}\ Heap\ h\ h'$ iff there are confining partitions of h, h', with the same number n of owner islands, such that

—$\mathcal{L}\ (Oh_i * Rh_i)\ (Oh'_i * Rh'_i)$ for all i in $1..n$
—$dom(Ch) = dom(Ch')$
—$\mathcal{R}\ (state\ (loctype\ \ell))\ (h\ell)\ (h'\ell)$ for all $\ell \in dom(Ch)$

For other categories θ we define $\mathcal{R}\,\theta$ as follows.

$$
\begin{array}{ll}
\mathcal{R}\ \mathbf{bool}\ d\ d' & \Leftrightarrow\ d = d' \\
\mathcal{R}\ \mathbf{unit}\ d\ d' & \Leftrightarrow\ d = d' \\
\mathcal{R}\ C\ d\ d' & \Leftrightarrow\ d = d' \\
\mathcal{R}\ \Gamma\ \eta\ \eta' & \Leftrightarrow\ \forall x \in dom\,\Gamma \cdot \mathcal{R}\ (\Gamma x)\ (\eta x)\ (\eta' x) \\
\mathcal{R}\ (state\, C)\ s\ s' & \Leftrightarrow \\
\quad C \not\leq Own \wedge \forall f \in dom(fields\, C) \cdot \mathcal{R}\ (type(f, C))\ (s\ f)\ (s'\ f) & \\
\mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta') & \Leftrightarrow\ \mathcal{R}\ Heap\ h\ h' \wedge \mathcal{R}\ \Gamma\ \eta\ \eta' \\
\mathcal{R}\ (Heap \otimes T)\ (h, d)\ (h', d') & \Leftrightarrow\ \mathcal{R}\ Heap\ h\ h' \wedge \mathcal{R}\ T\ d\ d' \\
\mathcal{R}\ (\theta_\bot)\ \alpha\ \alpha' & \Leftrightarrow\ (\alpha = \bot = \alpha') \vee (\alpha \neq \bot \neq \alpha' \wedge \mathcal{R}\ \theta\ \alpha\ \alpha') \\
\mathcal{R}\ (C, \overline{x}, \overline{T} \to T)\ d\ d' & \Leftrightarrow\ \forall (h, \eta) \in [\![Heap \otimes \Gamma]\!], (h', \eta') \in [\![Heap \otimes \Gamma]\!]' \ \cdot \\
\quad \mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta') \wedge conf\ C\ (h, \eta) \wedge conf\ C\ (h', \eta') & \\
\quad \Rightarrow \mathcal{R}\ (Heap \otimes T)_\bot\ (d(h, \eta))\ (d'(h', \eta')) & \\
\quad \text{where } \Gamma = [\overline{x} \mapsto \overline{T}, \mathsf{self} \mapsto C] & \\
\mathcal{R}\ Menv\ \mu\ \mu' & \Leftrightarrow\ \forall C, m\ \cdot (C \text{ is non-rep}) \wedge (mtype(m, C) \text{ is defined}) \\
\quad \Rightarrow \mathcal{R}\ (C, pars(m, C), mtype(m, C))\ (\mu\, C\, m)\ (\mu'\, C\, m) & \qquad \Box
\end{array}
$$

The gist of the abstraction theorem is that if the methods of Own are related by $\mathcal{R}$ then all methods are. We can now express this conclusion as $\mathcal{R}\ Menv\ [\![CT]\!]\ [\![CT']\!]'$. To express the antecedent, note that the relation applicable to a method m of Own is $\mathcal{R}\ (Own, \overline{x}, \overline{T} \to T)$ where $mtype(m, Own) = \overline{T} \to T$ and $pars(m, Own) = \overline{x}$. The definition of $\mathcal{R}\,(C, \overline{x}, \overline{T} \to T)$ quantifies over confined initial states but does not require confinement of outcomes.[10] The antecedent will also take into account that methods may be declared or inherited.

Although the definition is technically intricate, the core idea is the extension of a local coupling, for a single owner instance, to a heap containing potentially many

[10]One might think that $\mathcal{R}\ Heap$ could be defined in terms of admissible partitions without the assumption of confinement. But because partitions are not unique this leads to difficulties: a heap could be confined with respect to one partition but related with respect to another.

owners. This idea is given straightforward expression using heap partitions. By contrast, sharing of representations between owners would require a more complicated form of extension (see Sect. 12).

***Definition* 7.9 (simulation)** A simulation is a coupling $\mathcal{R}$, based on a local coupling $\mathcal{L}$, such that the following hold.

(1) (construction of Own establishes $\mathcal{L}$) For any $\ell \in locs(Own\downarrow)$, any h, h' with $\mathcal{R}\ Heap\ h\ h'$, and any μ, μ', let

$$h_0 = [h \mid \ell \mapsto [fields(loctype\ \ell) \mapsto defaults]]$$
$$h'_0 = [h' \mid \ell' \mapsto [fields(loctype\ \ell') \mapsto defaults']]$$

Then $\mathcal{L}\ h_0\ h'_0$.

(2) (methods of Own preserve $\mathcal{R}$) Let $\mu \in \mathbb{N} \to [\![Menv]\!]$ (resp. $\mu' \in \mathbb{N} \to [\![Menv]\!]'$) be the approximation chain in the definition of $[\![CT]\!]$ (resp. $[\![CT']\!]'$). For every m with $mtype(m, Own)$ defined, the following implications hold for every i, where $\overline{x} = pars(m, Own)$ and $\overline{T} \to T = mtype(m, Own)$.
   (a) $\mathcal{R}\ Menv\ \mu_i\ \mu'_i \ \Rightarrow\ \mathcal{R}\ (Own, \overline{x}, \overline{T} \to T)\ ([\![M]\!]\mu_i)\ ([\![M']\!]'\mu'_i)$
   if m has declaration M in $CT(Own)$ and M' in $CT'(Own)$
   (b) $\mathcal{R}\ Menv\ \mu_i\ \mu'_i \ \Rightarrow\ \mathcal{R}\ (Own, \overline{x}, \overline{T} \to T)\ ([\![M]\!]\mu_i)\ (restr([\![M_B]\!]'\mu'_i, Own))$
   if m has declaration M in $CT(Own)$ and is inherited from B in $CT'(Own)$, with M_B the declaration of m in B
   (c) the condition symmetric to (2b), if m is inherited in $CT(Own)$ but declared in $CT'(Own)$ □

To handle constructors, condition (1) would say that $\mathcal{L}$ holds upon termination of the constructor. For constructors without method calls this can be formalized easily (see [Banerjee and Naumann 2004a]).

The following properties are straightforward consequences of the definition.

Lemma 7.10 For all h, h' and all locations $\ell \notin locs(Rep\downarrow, Rep'\downarrow)$, if $\mathcal{R}\ Heap\ h\ h'$ then $\ell \in dom\, h \Leftrightarrow \ell \in dom\, h'$. □

Lemma 7.11 $[\![T]\!] = [\![T]\!]'$ for all T, and $[\![\Gamma]\!] = [\![\Gamma]\!]'$ for all Γ. □

Lemma 7.12 For any data type T, $\mathcal{R}\ T$ is the identity relation on $[\![T]\!]$ and $\mathcal{R}\ T_\perp$ is the identity relation on $[\![T_\perp]\!]$. □

Lemma 7.13 If $\overline{U} \leq \overline{T}$ and $\mathcal{R}\ \overline{U}\ \overline{d}\ \overline{d'}$ then $\mathcal{R}\ \overline{T}\ \overline{d}\ \overline{d'}$. □

7.3 Restricting reps in owner subclasses

The preceding properties express a strong connection between locations for related heaps. To ensure that this connection is preserved by object construction, we shall assume the allocator is parametric. But it is not reasonable to require that related heaps have the same rep locations, so parametricity cannot be exploited for reps. As a result, the present form of simulation is not adequate for construction of reps in subclasses of Own, although such construction is allowed by confinement. The first abstraction theorem depends on an assumption expressed in the following terms.

***Definition* 7.14 (new rep in sub-owner)** We say CT *has a new rep in a sub-owner* if, for some $B \leq Rep$ or $B \leq Rep'$, an object construction **new** B occurs in some method declaration in a class $C < Own$.

If CT has no new reps in sub-owners then neither does a comparable CT' (and vice versa). The examples in Sects. 2 and 3 have no new reps in sub-owners; examples which do are given in Sect. 9.

In the rest of Sect. 7 we make the following assumption. It is used in the proof of Lemma 7.22 on which the first abstraction theorem depends. For the second abstraction theorem the second sentence of the assumption will be dropped.

Assumption 7.15 First, CT and CT' are confined class tables for which a simulation $\mathcal{R}$ is given. Second, CT has no new reps in sub-owners and the allocator is parametric in the sense of Def. 5.1.

7.4 Identity extension

In our theory, $\mathcal{R}\ T$ is the identity for every data type T (Lemma 7.12), but that is only because the interesting data is in the heap. In general, $[\![state\, Own]\!] \neq [\![state\, Own]\!]'$ and $\mathcal{R}(state\, Own)$ is not the identity. Related heaps can contain owner objects with different states that may point to completely different rep objects. But consider executing a method on an object o from whose fields no Own objects are reachable, i.e., Own objects are not part of the representation of o. The resulting heap may contain Own objects that were assigned to local variables, but if the method is confined then those objects are unreachable in the final state.

***Definition* 7.16 (garbage collection, Own-free)** For a set or list $\overline{d}$ of values, define the heap $gc(\overline{d}, h)$ to be the restriction of h to cells reachable from $\overline{d}$. For $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$, define $collect(h, \eta) = (gc(rng\, \eta, h), \eta)$. Extend $collect$ to $[\![(Heap \otimes \Gamma)_\bot]\!]$ by $collect\, \bot = \bot$.

Say h is *Own-free* just if $dom\, h \cap locs(Own{\downarrow}) = \varnothing$ and η is *Own-free* just if $rng\, \eta \cap locs(Own{\downarrow}) = \varnothing$. State (h, η) is *Own-free* just if both h and η are. □

Lemma 7.17 (identity extension) Suppose $\mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta')$ and $\Gamma\, \mathsf{self}$ is non-rep. Let (h, η) and (h', η') be confined at $\Gamma\, \mathsf{self}$. If $collect(h, \eta)$ and $collect(h', \eta')$ are Own-free then $collect(h, \eta) = collect(h', \eta')$.

Proof. In confined heaps, reps are only reachable from owners. The argument is a straightforward induction using the definition of $\mathcal{R}$. □

Lemma 7.18 (diagonal) For any $\mathcal{R}$ given by Def. 7.8 from a local coupling, if $h \in [\![Heap]\!]$ is Own-free then $\mathcal{R}\ Heap\ h\ h$. If, in addition, $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ then $\mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h, \eta)$.

Proof. If h is Own-free and confined then it has no reps; its admissible partition is a single block, the clients. For such a heap it is immediate from the definition of $\mathcal{R}$ that $\mathcal{R}\ Heap\ h\ h$. If $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ then $rng\, \Gamma$ is Own-free and $\mathcal{R}\ \Gamma\ \eta\ \eta$ is direct from the definition. □

7.5 Abstraction theorem

The proof depends on a lemma for commands, which is given following the theorem. The other main ingredient for the proof is the following connection between $\mathcal{R}$ and

the semantics of inherited methods.

Lemma 7.19 Suppose C and all class names in $\overline{T}$ are non-rep, and $B < C$. If $\mathcal{R}\ (C, \overline{x}, \overline{T}{\rightarrow}T)\ d\ d'$ then $\mathcal{R}\ (B, \overline{x}, \overline{T}{\rightarrow}T)\ (\mathit{restr}(d, B))\ (\mathit{restr}(d', B))$ where *restr* is the restriction to global states of B (see Def. 5.5).

PROOF. Straightforward, using Lemma 5.7(1). $\square$

Theorem 7.20 (abstraction) If CT and CT' are confined and $\mathcal{R}$ is a simulation (as per Assumption 7.15), then $\mathcal{R}\ \mathit{Menv}\ [\![CT]\!]\ [\![CT']\!]'$.

PROOF. We show that the relation holds for each step in the approximation chain in the semantics of class tables (see the definition of μ_i following Def. 5.5). That is, we show by induction on i that

$$\mathcal{R}\ \mathit{Menv}\ \mu_i\ \mu'_i \quad \text{for every } i \in \mathbb{N}\ .$$

The result $\mathcal{R}\ \mathit{Menv}\ [\![CT]\!]\ [\![CT']\!]'$ then follows by fixpoint induction, as $[\![CT]\!]$ and $[\![CT']\!]'$ are defined to be the fixpoints of these ascending chains. Admissibility of fixpoint induction is discussed preceding the proof of Theorem 6.15.[11]

For the base case, we have $\mathcal{R}\ (C, \mathit{pars}(m, C), \mathit{mtype}(m, C))\ (\mu_0\, C\, m)\ (\mu'_0\, C\, m)$ for every m, C because $\lambda(h, \eta) \cdot \bot$ relates to itself.

For the induction step, suppose

$$\mathcal{R}\ \mathit{Menv}\ \mu_i\ \mu'_i \quad . \tag{$*$}$$

We must show $\mathcal{R}\ \mathit{Menv}\ \mu_{i+1}\ \mu'_{i+1}$, that is, for every non-rep C and every m with $\mathit{mtype}(m, C)$ defined:

$$\mathcal{R}\ (C, \overline{x}, \overline{T}{\rightarrow}T)\ (\mu_{i+1}\, C\, m)\ (\mu'_{i+1}\, C\, m) \tag{$\dagger$}$$

where $\overline{x} = \mathit{pars}(m, C)$ and $\overline{T}{\rightarrow}T = \mathit{mtype}(m, C)$. For arbitrary m we show ($\dagger$) for all C with $\mathit{mtype}(m, C)$ defined, using a secondary induction on $\mathit{depth}(m, C)$. We have $\mathit{depth}'(m, C) = \mathit{depth}(m, C)$ from Lemma 7.3.

The base case is the unique C with $\mathit{depth}(m, C) = 0$; here m is declared in both $CT(C)$ and $CT'(C)$. We go by cases on C. If $C = \mathit{Own}$, we get ($\dagger$) from the assumption that $\mathcal{R}$ is a simulation. In detail: Using ($*$) and Def. 7.9(2a) we get

$$\mathcal{R}\ (\mathit{Own}, \overline{x}, \overline{T}{\rightarrow}T)\ ([\![M]\!]\mu_i)\ ([\![M']\!]'\mu'_i)\ ,$$

whence ($\dagger$) by definition of μ_{i+1} and μ'_{i+1}. The other case is C a non-rep class different from Own. Then by Def. 7.1(1) of comparable class tables we have $CT(C) = CT'(C)$ and in particular both class tables have the same declaration

$$T\ m(\overline{T}\ \overline{x})\ \{S\}\ .$$

To show ($\dagger$), suppose $\mathit{conf}\ C\,(h, \eta)$, $\mathit{conf}\ C\,(h', \eta')$, $\mathcal{R}\ \mathit{Heap}\ h\ h'$, and $\mathcal{R}\ \Gamma\ \eta\ \eta'$, where $\Gamma = \overline{x}{:}\overline{T}, \mathsf{self}{:}C$. Then by Lemma 7.22 below, using $\mathcal{R}\ \mathit{Menv}\ \mu_i\ \mu'_i$, the results from S are related. That is, either $[\![\Gamma \vdash S]\!]\mu_i(h, \eta) = \bot = [\![\Gamma \vdash' S]\!]'\mu'_i(h', \eta')$

[11] Readers familiar with Reynolds [1984] may expect that, as our language has fixpoints, the result only holds for couplings that are $\bot$-strict and join-complete. But our local couplings have this property, trivially, because heaps are ordered by equality. The induced coupling is strict and join-complete by construction.

or neither is $\bot$. In the latter case, (h_0, η_0) is related to (h'_0, η'_0) where $(h_0, \eta_0) = [\![\Gamma \vdash S]\!]\mu_i(h, \eta)$ and $(h'_0, \eta'_0) = [\![\Gamma \vdash' S]\!]'\mu'_i(h', \eta')$. Then, by definition of $\mathcal{R}\,\Gamma$, $\mathcal{R}\ \Gamma\ \eta_0\ \eta'_0$ implies $\mathcal{R}\ T\ (\eta_0\,\mathsf{result})\ (\eta'_0\,\mathsf{result})$. Thus (†) holds by definition of μ_{i+1} and μ'_{i+1}. This concludes the base case of the secondary induction. The appeal to Lemma 7.22 depends on $conf\ \mu_i$ and $conf\ \mu'_i$ which holds by Assumption 7.15 and Theorem 6.15.

For the induction step, suppose $depth(m, C) > 0$. Using the definition of $depth$, the induction hypothesis is

$$\mathcal{R}\ (C, \overline{x}, \overline{T}{\rightarrow}T)\ (\mu_{i+1}\,(super\,C)\,m)\ (\mu'_{i+1}\,(super\,C)\,m) \quad . \qquad (‡)$$

If m is declared in both $CT(C)$ and $CT'(C)$ then the argument is the same as in the base case of the secondary induction. If m is inherited in both $CT(C)$ and $CT'(C)$ then (†) follows from (‡) because the semantics defines $\mu_{i+1}\,C\,m$ by restriction from $\mu'_{i+1}\,(super\,C)\,m$ and restriction preserves simulation. (This is Lemma 7.19, which is applicable because if $B > Own$ and m is inherited in Own from B then $\overline{T} \not\geq Rep$ and $\overline{T} \not\geq Rep'$ by confinement of CT, CT', Def. 6.9(3).) The remaining possibility is that m is declared in $CT(C)$ and inherited in $CT'(C)$ from some B (or the other way around). Then $C = Own$, by comparability of CT and CT'. Using Def. 7.9(2b) and $(*)$ we get

$$\mathcal{R}\ (Own, \overline{x}, \overline{T}{\rightarrow}T)\ ([\![M]\!]\mu_i)\ (restr([\![M_B]\!]\mu'_i, Own))$$

and thus (†) by definition of μ_{i+1} and μ'_{i+1}. □

Lemma 7.21 (preservation by expressions) Consider any non-rep class $C \neq Own$ and any $\Gamma \vdash e : T$ with $\Gamma\,\mathsf{self} = C$. If $\Gamma \vdash e : T$ is confined and all constituents of e are confined then the following holds: For all $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ and $(h', \eta') \in [\![Heap \otimes \Gamma]\!]'$, if $conf\,C\,(h, \eta)$, $conf\,C\,(h', \eta')$, and $\mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta')$ then

$$\mathcal{R}\ (T_\bot)\ ([\![\Gamma \vdash e : T]\!](h, \eta))\ ([\![\Gamma \vdash' e : T]\!]'(h', \eta')) \ .$$

Proof. By induction on the derivation of $\Gamma \vdash e : T$. For each case of e, we give an argument assuming that $\Gamma, C, T, \eta, \eta', h, h'$ satisfy the hypotheses of the Lemma.

Case $\Gamma \vdash (B)\ e : B$. Induction on e yields that $\mathcal{R}\ D_\bot\ \ell\ \ell'$ (or else both denotations of e are $\bot$). By confinement of e, as $C \neq Own$ and C is non-rep, we have $\ell \notin locs(Rep{\downarrow})$ and $\ell' \notin locs(Rep'{\downarrow})$. Thus, $\ell' = \ell$ by Lemma 7.12. Hence either both semantics yield ℓ, whence $\mathcal{R}\ B_\bot\ \ell\ \ell$, or both yield $\bot$ and again $\mathcal{R}\ B_\bot\ \bot\ \bot$.

Case $\Gamma \vdash e$ **is** $B :$ **bool**. The argument is similar to that for type cast.

Case $\Gamma \vdash e.f : T$. By induction on e we have $\mathcal{R}\ C_\bot\ \ell\ \ell'$, hence $\ell = \ell'$ by Lemma 7.12. In the non-$\bot$ case, $\ell \neq nil$. By closure of the heaps, $\ell \in dom\,h$ and $\ell \in dom\,h'$. We consider cases on whether $C < Own$. Consider confining partitions $(Ch * Oh_1 * Rh_1 \ldots) = h$ and $(Ch' * Oh'_1 * Rh'_1 \ldots) = h'$ that have corresponding islands as in the definition of $\mathcal{R}\ Heap$. In the case $C < Own$, we have $\ell \in locs(Own{\downarrow})$ and hence ℓ in some $dom(Oh_i)$. From $\mathcal{R}\ Heap\ h\ h'$ we have

$$\mathcal{L}\ (Oh_i * Rh_i)\ (Oh'_i * Rh'_i)$$

and thus $\ell \in dom(Oh'_i)$ by local coupling Def. 7.4(1). Since $C \neq Own$, we know by visibility that f is not in the private fields $\overline{g}$ of Own. Thus, as $type(f, loctype\,\ell)) = T$, we have $\mathcal{R}\ T\ (h\ell f)\ (h'\ell f)$ by Def. 7.4(3) and Lemma 7.12.

Finally, in the case $C \not\leq Own$ (recall that C is non-rep and $C \neq Own$ by hypothesis, we have $\ell \in dom(Ch)$ and hence $\ell \in dom(Ch')$ by definition $\mathcal{R}$ *Heap*. Hence $\mathcal{R}$ $(state\,(loctype\,\ell))$ $(h\ell)$ $(h'\ell)$ and thus $\mathcal{R}$ T $(h\ell f)$ $(h'\ell f)$ by definition of $\mathcal{R}$ $(state\,(loctype\,\ell))$.

The remaining cases are straightforward. □

Lemma 7.22 (preservation by commands) Suppose $\mathcal{R}$ is a simulation, and moreover μ and μ' are confined method environments such that $\mathcal{R}$ *Menv* μ μ'. Consider any non-rep class $C \neq Own$ and any $\Gamma \vdash S$ with $\Gamma\,\mathsf{self} = C$. If $\Gamma \vdash S$ is confined and all constituents of S are confined then the following holds: For all (h, η) and (h', η'), if $conf\,C\,(h, \eta)$, $conf\,C\,(h', \eta')$, and $\mathcal{R}$ $(Heap \otimes \Gamma)$ (h, η) (h', η') then

$$\mathcal{R}\ (Heap \otimes \Gamma)_{\perp}\ ([\![\Gamma \vdash S]\!]\mu(h, \eta))\ ([\![\Gamma \vdash' S]\!]'\mu'(h', \eta'))\ .$$

Proof. For any C, the proof is by induction on the derivation of $\Gamma \vdash S$.

Case $\Gamma \vdash x := e$. By confinement of e and Lemma 7.21 we have $\mathcal{R}$ $T_{\perp}$ d d'. Hence, by $\mathcal{R}$ Γ η η' and definition of $\mathcal{R}$ Γ, we have $\mathcal{R}$ Γ $[\eta \mid x \mapsto d]$ $[\eta' \mid x \mapsto d']$ whence the result.

Case $\Gamma \vdash e_1.f := e_2$. By Lemma 7.21 for e_1 we have $\mathcal{R}$ C ℓ ℓ', hence $\ell = \ell'$ by Lemma 7.12. By Lemma 7.21 for e_2 we have $\mathcal{R}$ U d d' and hence $\mathcal{R}$ T d d' by Lemma 7.13, where $(f : T) \in dfields\,C$ as in the typing rule. To conclude the argument it suffices to show

$$\mathcal{R}\ Heap\ [h \mid \ell \mapsto [h\ell \mid f \mapsto d]]\ [h' \mid \ell \mapsto [h'\ell \mid f \mapsto d']]\quad . \qquad (*)$$

Consider confining partitions $(Ch * Oh_1 * Rh_1 \ldots) = h$ and $(Ch' * Oh'_1 * Rh'_1 \ldots) = h'$ that correspond as in the definition of $\mathcal{R}$ *Heap* h h'. We argue by cases on C.

—$C < Own$: Then $loctype\,\ell \leq C < Own$. From typing we have $e_1 : C$ and hence there is some i with $\{\ell\} = dom(Oh_i)$ and by $\mathcal{R}$ *Heap* h h' we get

$$\mathcal{L}\ (Oh_i * Rh_i)\ (Oh'_i * Rh'_i)$$

and so $\{\ell\} = dom(Oh'_i)$. By typing and $C \neq Own$, field f is not in the private fields $\overline{g}$ of Own. So $(*)$ follows from $\mathcal{R}$ *Heap* h h' and $\mathcal{R}$ T d d'.

—$C \not\leq Own$: As C is non-rep, we have $\ell \in dom\,Ch$ and thus $\ell \in dom\,Ch'$ by hypothesis $\mathcal{R}$ *Heap* h h'. Moreover, $\mathcal{R}$ $(state\,(loctype\,\ell))$ $(h\ell)$ $(h'\ell)$ and so by $\mathcal{R}$ T d d' we get $\mathcal{R}$ $(state\,(loctype\,\ell))$ $[h\ell \mid f \mapsto d]$ $[h'\ell \mid f \mapsto d']$. Hence $(*)$.

Case $\Gamma \vdash x := \mathbf{new}\ B$. By confinement of $x := \mathbf{new}\ B$ the final states are confined: $conf\,C\,(h_0, \eta_0)$ and $conf\,C\,(h'_0, \eta'_0)$. We have $C \not\leq Rep$ and $C \neq Own$. In the case $C \not< Own$ confinement of η_0 and η'_0 implies $rng\,\eta_0 \cap locs(Rep{\downarrow}) = \varnothing = rng\,\eta'_0 \cap locs(Rep'{\downarrow})$. So $\ell \notin locs(Rep{\downarrow})$ and $\ell' \notin locs(Rep'{\downarrow})$, hence by typing B is non-rep. In the case $C < Own$, we have B non-rep by Assumption 7.15 (no reps in sub-owners). Either way, B is non-rep so Lemma 7.10 applies, to yield $dom\,h \cap locs\,B = dom\,h' \cap locs\,B$. Thus by parametricity of *fresh* we have $\ell = fresh(B, h) = fresh(B, h') = \ell'$. So, by Lemma 7.12 and $\mathcal{R}$ Γ η η' we have $\mathcal{R}\,\Gamma\,\eta_0\,\eta'_0$.

It remains to show $\mathcal{R}$ $Heap_{\perp}$ h_0 h'_0 in order to get the final result $\mathcal{R}$ $(Heap \otimes \Gamma)_{\perp}$ (h_0, η_0) (h'_0, η'_0). We argue by cases on B.

—$B \not\leq Own$: Writing fields' for the fields given by CT', we have $\mathit{fields}\,B = \mathit{fields}'\,B$ and thus $\mathcal{R}\ (\mathit{state}\,B)\ [\mathit{fields}\,B \mapsto \mathit{defaults}]\ [\mathit{fields}'\,B \mapsto \mathit{defaults}]$. So, as B is non-rep and $B \neq Own$, we can add ℓ to Ch and Ch' to get partitions that witness $\mathcal{R}\ \mathit{Heap}\ h_0\ h'_0$. Combining this with what was shown above we have $\mathcal{R}\ (\mathit{Heap} \otimes \Gamma)_\perp\ (h_0, \eta_0)\ (h'_0, \eta'_0)$.

—$B \leq Own$: By simulation Def. 7.9(1), we have $\mathcal{R}\ \mathit{Heap}_\perp\ h_0\ h'_0$.

Case $\Gamma \vdash x := e.m(\overline{e})$. By Lemma 7.21 for e we have $\mathcal{R}\ D_\perp\ \ell\ \ell'$, hence $\ell = \ell'$ by Lemma 7.12. Let $\eta_1 = [\mathsf{self} \mapsto \ell, \overline{x} \mapsto \overline{d}]$ and $\eta'_1 = [\mathsf{self} \mapsto \ell, \overline{x} \mapsto \overline{d}']$. By confinement of $x := e.m(\overline{e})$ (Def. 6.7) we have confined arguments, i.e., $\mathit{conf}\,(\mathit{loctype}\,\ell)\,(h, \eta_1)$ and $\mathit{conf}\,(\mathit{loctype}\,\ell)\,(h', \eta'_1)$. By Lemma 7.21 for $\overline{e}$ we have $\mathcal{R}\ \overline{U}_\perp\ \overline{d}\ \overline{d}'$ and hence $\mathcal{R}\ \overline{U}\ \overline{d}\ \overline{d}'$ as we are considering the non-$\perp$ case. Thus $\mathcal{R}\ [\overline{x}:\overline{T}, \mathsf{self}: \mathit{loctype}\,\ell]\ \eta_1\ \eta'_1$ by Lemma 7.13. From $\mathcal{R}\ \mathit{Menv}\ \mu\ \mu'$ we get

$$\mathcal{R}\ (\mathit{loctype}\,\ell, \overline{x}, \overline{T} \to T)\ (\mu(\mathit{loctype}\,\ell)m)\ (\mu'(\mathit{loctype}\,\ell)m)$$

hence, as h, h', η_1, η'_1 are confined and related, $\mathcal{R}\ (\mathit{Heap} \otimes T)_\perp\ (h_1, d_1)\ (h'_1, d'_1)$, where $(h_1, d_1) = \mu(\mathit{loctype}\,\ell)m(h, \eta)$ and $(h'_1, d'_1) = \mu'(\mathit{loctype}\,\ell)m(h', \eta')$. Thus $\mathcal{R}\ T\ d_1\ d'_1$ and $\mathcal{R}\ \mathit{Heap}\ h_1\ h'_1$. It remains to show that the updated stores $[\eta \mid x \mapsto d_1]$ and $[\eta' \mid x \mapsto d'_1]$ are related for Γ. This follows from $\mathcal{R}\ T\ d_1\ d'_1$ and $T \leq \Gamma\,x$, using Lemma 7.13.

The remaining cases are similar. □

8. APPLICATIONS

In this section we define program equivalence and then use the abstraction theorem to show some program equivalences for the examples discussed in Sections 2 and 3.

To establish the hypothesis of the abstraction theorem for the examples we use the couplings given as examples in Sect. 7.2. Both the theorem and these couplings are defined in terms of the semantics. To show that the couplings are simulations we argue directly in terms of the semantics. For practical purposes in program verification, the abstraction theorem would be expressed syntactically as a proof rule and rules for program constructs would be used to establish the simulation property [Reynolds 1981a; Jones 1986; de Roever and Engelhardt 1998].

Program equivalence. We take *program* to mean a well formed class table CT together with a command $\Gamma \vdash S$. We consider the object states reachable from variables of Γ to be the inputs and outputs of the program. For example, if S is the body of method main in Sect. 2.1 then Γ is self:Main and what can be reached is self and the string self.inout.

We restrict attention to confined programs, by which we mean that CT and $\Gamma \vdash S$ are confined (as well as the constituent parts of S). Thus, by Theorem 6.15 the method environment $[\![CT]\!]$ is confined. To prove program equivalence using the abstraction theorem, we need to both introduce and eliminate a simulation. Elimination is by identity extension Lemma 7.17 and introduction is by the diagonal Lemma 7.18.

We compare programs only for class tables CT, CT' that are comparable in the sense of Def. 7.1, and with commands in the same context Γ. As commands denote functions on global states, the obvious notion of equivalence is that $[\![\Gamma \vdash S]\!]$ and

$[\![\Gamma \vdash' S']\!]'$ are equal as functions. By Lemma 7.11, $[\![\Gamma]\!] = [\![\Gamma]\!]'$ for any Γ, but in general the semantic domains differ for owner object states which may have different private fields. A global state $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$ for CT need not be an element of $[\![Heap \otimes \Gamma]\!]'$ for CT'. However, an Own-free heap in $[\![Heap]\!]$ is also an element of $[\![Heap]\!]'$. We compare command meanings on the Own-free states, defined using *collect* from Def. 7.16.

***Definition* 8.1 (client program equivalence)** Suppose programs $CT, (\Gamma \vdash S)$ and $CT', (\Gamma \vdash' S')$ are such that CT, CT' are comparable and confined. The programs are *equivalent* iff

$$collect([\![\Gamma \vdash S]\!]\hat{\mu}(h, \eta)) = collect([\![\Gamma \vdash' S']\!]'\hat{\mu}'(h, \eta))$$

for all confined and Own-free $(h, \eta) \in [\![Heap \otimes \Gamma]\!]$, where $\hat{\mu} = [\![CT]\!]$ and $\hat{\mu}' = [\![CT']\!]'$. □

If $\Gamma\,\mathsf{self} \leq Own$ then η cannot be Own-free. The resulting vacuous quantification makes the definition equate all commands for such Γ. But we are only interested in using the definition for clients. Simulation is the relation of interest between owners.

The main corollary of the abstraction theorem is that if there is a simulation for CT, CT' then $CT, (\Gamma \vdash S)$ is equivalent to $CT', (\Gamma \vdash S)$ for any suitable S. For S to be suitable means it is typable in both CT and CT' and moreover its final states are Own-free. Rather than formalize the corollary and its proof we illustrate the reasoning in a series of examples.

Examples. The static analysis for confinement given in Sect. 11 can be used to show that each of the following examples is confined for the appropriate Own and Rep.

***Example* 8.2** Consider the command S comprising the body of method main of class Main in Sect. 2.1 and take $\Gamma = (\mathsf{self}\!:\!\mathsf{Main})$. As CT we take the declarations of Main, Bool, and the first version of OBool. For CT' we use the second version of OBool. Let Rep and Rep' be Bool and Own be OBool.

To show that $CT, (\Gamma \vdash S)$ is equivalent to $CT', (\Gamma \vdash S)$, recall the local coupling of Example 7.5 and let $\mathcal{R}$ be the induced coupling. Let (h, η) be any confined state for Γ, noting that $\mathsf{Main} \not\leq Own$ so η is Own-free. Let $\hat{\mu} = [\![CT]\!]$ and $\hat{\mu}' = [\![CT']\!]'$. To show

$$collect([\![\Gamma \vdash S]\!]\hat{\mu}(h, \eta)) = collect([\![\Gamma \vdash' S']\!]'\hat{\mu}'(h, \eta)) \quad , \tag{$*$}$$

note first that $\mathcal{R}\ (Heap \otimes \Gamma)\ (h, \eta)\ (h, \eta)$ by Lemma 7.18. It is straightforward to show that $\mathcal{R}$ is established initially and preserved by the methods of OBool; thus $\mathcal{R}$ is a simulation. The abstraction theorem yields $\mathcal{R}\ Menv\ \hat{\mu}\ \hat{\mu}'$. This in turn justifies application of the preservation Lemma 7.22 to command S, as its context Main is a non-rep class and $\mathsf{Main} \neq Own$. Thus the outcomes $[\![\Gamma \vdash S]\!]\hat{\mu}(h, \eta)$ and $[\![\Gamma \vdash' S']\!]'\hat{\mu}'(h, \eta)$ are related by $\mathcal{R}$. By definition of $\mathcal{R}$, either both outcomes are $\perp$, in which case $(*)$ holds by definition of *collect*, or the outcomes are non-$\perp$ states (h_0, η_0) and (h_0', η_0') with $\mathcal{R}\ (Heap \otimes \Gamma)\ (h_0, \eta_0)\ (h_0', \eta_0')$. Note that h_0 and h_0' each contains at least one owner, the one constructed in S. But $\mathsf{Main} \not\leq Own$, so $rng\,\eta_0$ and $rng\,\eta_0'$ are Own-free. Moreover, the owners were reached only by variable z

which is local in S; they are not reachable via fields of the objects $h_0(\eta\,\mathsf{self})$ or $h'_0(\eta'\,\mathsf{self})$. That is, both $collect(h_0, \eta_0)$ and $collect(h'_0, \eta'_0)$ are Own-free. Thus by identity extension Lemma 7.17 we have $collect(h_0, \eta_0) = collect(h'_0, \eta'_0)$ which concludes the proof of $(*)$. □

This proof depends on parametricity of the allocator, because that is needed for the abstraction theorem. The same argument will go through, however, for the second abstraction theorem in the sequel which drops parametricity of the allocator.

***Example* 8.3** Recall the Meyer-Sieber-O'Hearn example from Sect. 3.1, and in particular the command

$$\mathsf{C\ y} := \textbf{new}\ \mathsf{C}\ \textbf{in}\ \mathsf{A\ x} := \textbf{new}\ \mathsf{A}\ \textbf{in}\ \mathsf{x.callP(y)} \qquad (\ddagger)$$

Take (‡) to be the body of method main in

class Main **extends Object** { **unit** main(){ ... } }

To be very precise we need to include a class

class Rep **extends Object** { }

so we can take Rep and Rep' to be Rep which is not comparable to the classes C and A of interest. Let Own be A. Let CT consist of the declarations of A, Rep, Main, and an arbitrary class

class C **extends Object** { **unit** P(A z){ ... } ... }

such that methods of C satisfy the confinement conditions. Then CT and CT' are confined, because no reps are constructed or manipulated. We use the local coupling of Example 7.6. To appeal to the abstraction theorem, we must argue that $\mathcal{R}$ is a simulation. The constructors are **skip** and the default value 0 for field g establishes the relation. Preservation by inc is straightforward because both versions have the same code and it makes no method calls. We give the details for preservation by callP. The relevant condition is Def. 7.9(2a). To show it for callP, suppose $i \geq 0$ and $\mathcal{R}\ Menv\ \mu_i\ \mu'_i$. Note that μ_i and μ'_i are confined, by Theorem 6.15. Suppose that $\mathcal{R}\ (Heap \otimes y\!:\!C, \mathsf{self}\!:\!A)\ (h,\eta)\ (h',\eta')$ with $conf\ A\,(h,\eta)$ and $conf\ A\,(h',\eta')$. In both versions of callP, the body is a sequence and the first command is y.P(self). Let $\eta_1 = [z \mapsto \eta\,\mathsf{self}, \mathsf{self} \mapsto \eta\, y]$ and $\eta'_1 = [z \mapsto \eta'\,\mathsf{self}, \mathsf{self} \mapsto \eta'\, y]$ be the environments for semantics of this call. By definition of $\mathcal{R}$ we get $\mathcal{R}\ (Heap \otimes z\!:\!A, \mathsf{self}\!:\!C)\ (h,\eta_1)\ (h',\eta'_1)$. From the hypothesis $conf\ A\,(h,\eta)$ we get $conf\ C\,(h,\eta_1)$ and likewise $conf\ C\,(h',\eta'_1)$. Applying the hypothesis $\mathcal{R}\ Menv\ \mu_i\ \mu'_i$ to these environments we get that either $\mu_i CP(h,\eta_1) = \bot = \mu'_i CP(h,\eta_1)$ or neither are $\bot$ and $\mathcal{R}\ (Heap \otimes \textbf{unit})\ (h_0, it)\ (h'_0, it)$ where $(h_0, it) = \mu_i CP(h,\eta_1)$ and $(h'_0, it) = \mu'_i CP(h',\eta'_1)$. The call is desugared to an assignment of the result value to a local variable but the value is discarded for both versions, so the states following the calls are (h_0,η) and (h'_0,η') and we have $\mathcal{R}\ (Heap \otimes y\!:\!C, \mathsf{self}\!:\!A)\ (h_0,\eta)\ (h'_0,\eta')$. In these states we have $h_0\ell g = h'_0\ell g \wedge h_0\ell g\ \textbf{mod}\ 2 = 0$. So the command

if self.g **mod** 2 = 0 **then abort else skip fi**

aborts, as does its counterpart which is simply **abort**. This concludes the argument that the bodies of callP are related.

Having established the antecedents of the abstraction theorem, we conclude that the command (‡) preserves $\mathcal{R}$. By semantics of the second version of A we know callP aborts, so both interpretations of (‡) abort. The programs are equivalent. □

This example is handled without using the identity extension Lemma 7.17, but that is only because the example uses abortion. In subsequent examples the proof needs all the steps from Sect. 2.1 just as in Example 8.2, though we only spell out the interesting bits.

One can imagine a variation of Def. 7.9(2a) which requires that $\mathcal{R}\ Menv\ \mu\ \mu'$ implies $\mathcal{R}\ (Own, \overline{x}, \overline{T} \to T)\ ([\![M]\!]\mu)\ ([\![M']\!]'\mu')$ for any confined method environments μ, μ', and not just those in the approximation chain. In fact, this suffices to prove the abstraction theorem but it makes the theorem impossible to apply in interesting cases. The argument in Example 8.3 would not go through because we would not know that the semantics of a called method is at least an approximation of its declaration.

***Example* 8.4** We consider the observer pattern, taking Own to be Observable. Let CT be given by the first version, Fig. 2, together with the client given in Fig. 3. Let CT' be given by the sentinel version of Fig. 4 together with Fig. 3. We consider equivalence for the command self:Main,ob:AnObserver $\vdash S$ where S is the body of Main.main. Because obl is local to S, no owners are reachable in the final state.

Taking Rep, Rep' to be Node,Node2, we use the coupling relation of Example 7.7. Clearly the constructors establish the relation (cf. [Banerjee and Naumann 2004a] where constructors are formalized). To show that method add preserves it, note that the bodies of these methods are both sequential compositions; both construct a new node and then set its ob field to the value passed as a parameter. The next step is to add it to the beginning of the list; the difference between the two versions is that self.snt.nxt is assigned in Fig. 4 whereas self.fst is assigned in Fig. 2. Both versions of add then invoke methods on the new node n, so we have to show that the results of these invocations are related. To give a precise argument in terms of the semantics, we consider cases on i. For $i = 0$, both μ_i and μ'_i make every method abort, in which case the body of add aborts due to method calls n.setOb(...). As the methods in class Node and class Node2 are not recursive, their semantics is already completely defined for $i = 1$, so for $i > 0$ the behavior of add in μ_i and μ'_i is to insert nodes at the head of the list, maintaining the relation.

The remaining owner method is notifyAll. Again, the two versions are similar except for skipping over the sentinel node. To argue that the calls to getNext act correctly one considers cases as in the proof for add. For the calls to notify on the Observer objects, recall that by the relation, the related lists contain the same Observer pointers in the same order. The two versions thus make the same series of invocations of notify. Each of those calls preserves the relation by hypothesis $\mathcal{R}\ Menv\ \mu_i\ \mu'_i$. □

The last step of the argument, concerning invocations of notify, is like reasoning about invocations of P in Example 8.3. This example has the additional complication of calls to objects within the owner island. The case distinction between $i = 0$ and $i > 0$ is needed because our argument is purely in semantic terms. In a practical proof system, one would reason only in terms of the actual semantics of

the methods involved rather than its approximants.

***Example* 8.5** Suppose we change the client of Fig. 3 to use the following.

class AnObserver **extends** Observer { **unit** notify(){ **skip** } }

Then in Fig. 4 we can replace the body of Observable.notifyAll by **skip** and still have equivalence with the implementation of Fig. 2. What changes with respect to Example 8.4 is that the two implementations do not make the corresponding calls to notify. But because AnObserver.notify is **skip**, calling it has the same effect as not calling it; in particular, the relation is preserved.

The argument here is not modular: by contrast with the preceding example, here we reason directly in terms of the client code. □

9. OWNER SUBCLASSING: THE PROTECTED INTERFACE

This section considers examples involving subclasses of the owner class. Rather than formalizing the "protected" scoping construct of Java, we consider the "sealed package", to address interaction between owners and reps as well as sub-owners. We model just a single module that includes Own and some or all of the subclasses of Own and of Rep. Some methods can be designated as having module scope so they cannot be called from client classes; for these the confinement conditions are relaxed. Our treatment is illustrative, not comprehensive; e.g., fields are still considered to be private. Additional features such as protected fields should be a straightforward addition.

Methods with protected scope can be modeled by taking the module to include all subclasses of Own and none of Rep. We can also model the situation where a method of Own is private to Own, or used only by reps, in which case it can be present in $CT(Own)$ but absent from $CT'(Own)$; even if present in both, it need not have the simulation property. On the other hand, if a module-scoped method is invoked in a subclass of Own then for well-formedness it must be present in both $CT(Own)$ and $CT'(Own)$; moreover, it must have the simulation property in order for the abstraction theorem to hold. We slightly abuse the term "protected" to refer to a module-scoped method of Own that is called in some subclass of Own.

9.1 Owner subclassing and module scope

Fig. 8 is a variation on the observer pattern in which class Observable has subclass ObservableAcc. For accounting purposes it keeps track of the number of times each observer has been notified. To this end, the rep class NodeAcc overrides method notifyAll of the client class Node4. Such examples have led to our treatment of owner subclasses: They are distinguished from clients in that their methods may manipulate reps, but unlike Own they cannot store reps in fields.

Method addn has been added to Observable, so that ObservableAcc can construct reps of the subtype NodeAcc and install them in the list despite that fst is a private field not visible in ObservableAcc. Method Observable.getFirst is also added for this purpose. But getFirst leaks a rep; it cannot be allowed in the public interface. One possibility is to treat getFirst and addn as visible only in subclasses of Observable. Instead, we give them module scope, meaning that calls to getFirst and addn are allowed in subclasses of both Observable and Node4.

```
class Node4 extends Object { // rep for Observable
  Observer ob;
  Node4 nxt;
  unit setOb(Observer o){ self.ob := o }
  unit setNext(Node4 n){ self.nxt := n }
  Observer getOb(){ result := self.ob }
  Node4 getNext(){ result := self.nxt }
  unit notifyAll(){ self.ob.notify(); if self.nxt ≠ null then self.nxt.notifyAll() else skip fi } }
class NodeAcc extends Node4 {
  int notifs;
  unit notifyAll(){ self.notifs := self.notifs+1; super.notifyAll() }
  int notifications(Observer o){
    if o = self.getOb() then result := notifs
    else if self.getNext() ≠ null then result := (NodeAcc)(self.getNext()).notifications(o)
    else result := 0; fi }}
class ObservableSup extends Object { // superclass of owner; "abstract" class
  unit add(Observer ob){ abort }
  unit notifyAll(){ abort }
class Observable extends ObservableSup { // owner
  Node4 fst; // first node of list
  Node4 getFirst(){ result := self.fst } // module scope
  unit add(Observer ob){ Node4 n := new Node4; self.addn(ob,n) }
  unit addn(Observer ob, Node4 n){ n.setNext(self.fst); n.setOb(ob); self.fst := n } // module scope
  unit notifyAll(){ self.fst.notifyAll() }
class ObservableAcc extends Observable {
  unit add(Observer ob){ Node4 n := new NodeAcc; self.addn(ob,n) }
  int notifications(Observer ob){ result := ((NodeAcc)(self.getFirst())).notifications(ob) } }
```

Fig. 8. Version with owner and rep subclasses and super-call. The owner also has a superclass.

Method add in class ObservableAcc constructs a rep, violating the condition "no reps in sub-owners" in Assumption 7.15. That assumption is needed for the first abstraction theorem because methods of an owner subclass are like clients in that they must preserve the induced relation. That means in particular that they manipulate related —equal!— rep locations. (By contrast, methods of *Own* preserve the local coupling which need not impose a correspondence on rep locations, cf. Example 8.4.) But if we compare two versions, one with sentinel node and one without, the parametricity condition for *fresh* will not apply and the new objects in ObservableAcc.add will be at different locations. The solution, given in Sect. 10, is to relax equality to bijection.

This relaxation is needed anyway, to avoid unobservable distinctions. As an example, suppose we add to class Observable in Fig. 2 the following method:

```
String version(){ result := new String("vsn 0") }
```

Consider an alternative that is identical in every way except for the following:

```
String version(){ result := new String("trash"); result := new String("vsn 0") }
```

According to Def. 7.9, the induced relation for locations of type String is equality. But, even if the allocator is parametric, the locations returned by these two methods are not equal. (So condition (2a) fails in Def. 7.9 of simulation.) Yet the two cannot be distinguished in any program context. This claim is justified by the generalized

```
class Observable extends ObservableSup {
  Node4 fst;
  Node4 getFirst(){ result := self.fst } // module scope
  Node4 makeNode(){ result := new Node4 } // module scope
  unit add(Observer ob){ Node4 n := makeNode(); n.setNext(self.fst); n.setOb(ob); self.fst := n }
  unit notifyAll(){ self.fst.notifyAll() } }
class ObservableAcc extends Observable {
  Node4 makeNode(){ result := new NodeAcc } // module scope
  int notifications(Observer ob){ result := ((NodeAcc)(self.getFirst())).notifications(ob) } }
```

Fig. 9. Variation on Fig. 8 using factory pattern. Node4 and NodeAcc are as in Fig. 8.

theory of Sect. 10, where the induced relation allows an arbitrary bijection between locations of client types like String. For this example, the bijection would get extended to relate the returned results from the two versions.

Returning to the example in Fig. 8, the interface betweeen Observable and its subclass ObservableAcc is awkwardly designed. An improvement is to use the factory pattern [Gamma et al. 1995] so that add itself can be inherited. In Fig. 9, we add method makeNode, which should have module scope, and remove addn.

These examples show subclasses of reps and owners. There is inheritance into the owner but not into the rep. Inheritance into reps is disallowed by our definition of confined class table, because to handle it requires a more sophisticated analysis to prevent leaks via self; a suitable analysis of "anonymous methods" is discussed in Sect. 12.2. Inheritance into owners also needs restriction; we have chosen a simple restriction that nonetheless allows the preceding examples.

Finally, let us consider an alternative version of Fig. 9 to illustrate the consequences of allowing the owner class, but not its subclasses, to differ in comparable class tables. In Fig. 9 the subclass ObservableAcc manipulates reps, both constructing a new NodeAcc and invoking method notifications declared in NodeAcc. Although an alternative version of Observable could use an entirely different type of nodes internally, it has to provide method getFirst with return type Node4. Because clients can manipulate objects of class ObservableAcc, methods of that class must preserve the relation and this only holds if methods they invoke preserve the relation. So coupling must be preserved not only by public methods of Observable but also by those module scope methods that are invoked in ObservableAcc.

9.2 Formalization of module-scoped methods

We assume that a class table designates the class names Own and Rep and is equipped with a predicate $mscope$ with the interpretation that $mscope(m, C)$ means this method has module scope. The following changes are made to the definitions of preceding sections.

(1) For a well formed class table, $mscope$ must satify conditions that reflect what in concrete syntax would be achieved by declaring Rep, Own, and some of their subclasses inside the module: If $mscope(m, C)$ then
 —$mtype(m, C)$ is defined,
 —$C \leq Own$ or $C \leq Rep$,
 —$mtype(m, B)$ is undefined for all $B > Own$ and $B > Rep$, and

—$B \leq C$ or $C \leq B$ implies $mscope(m, B)$ for all B with $mtype(m, B)$ defined

(2) The typing rule for method call has an added restriction that module-scoped methods are visible only within the module:

$$\frac{\Gamma \vdash e : D \quad \Gamma \vdash \overline{e} : \overline{U} \quad mtype(m, D) = \overline{T} \rightarrow T \quad \overline{U} \leq \overline{T} \quad T \leq \Gamma x \quad x \neq \mathsf{self} \quad mscope(m, D) \Rightarrow \Gamma\,\mathsf{self} \leq Own \vee \Gamma\,\mathsf{self} \leq Rep}{\Gamma \rhd x := e.m(\overline{e})}$$

(3) For method environments, the confinement condition of Def. 6.5(1) is replaced by the following:

—$C \leq Own \wedge mscope(m, C) \Rightarrow conf\, C\, (h_0, \eta) \wedge h \trianglelefteq h_0 \wedge (d \in locs(Rep{\downarrow}) \Rightarrow d \in dom(Rh_j))$ for some confining partition and j with $\eta\,\mathsf{self} \in dom(Oh_j)$

—$C \not\leq Rep \wedge (C \not\leq Own \vee \neg mscope(m, C)) \Rightarrow conf\, C\, (h_0, \eta) \wedge h \trianglelefteq h_0 \wedge d \notin locs(Rep{\downarrow})$

(4) For confinement of class tables, the restriction of Def. 6.9(2) is not imposed on methods such that $mscope(m, C)$.

(5) For simulation, Def. 10.10 in the sequel revises Def. 7.9(2) to require preservation of the relation only for public methods, that is, if $\neg(mscope(m, Own))$. But those module-scoped methods that are called in sub-owners must also preserve the relation.
To formalize this, we define $prot(m, C)$ just if $C \leq Own$, $mscope(m, Own)$, and there is a call to m in some subclass of Own.

(6) Comparable class tables must agree on not only the public but also the "protected" (in the sense of *prot* above) methods of Own: Def. 7.1(1) is extended to require that $mscope(m, C) = mscope'(m, C)$ for all $C \neq Own$. Moreover, if $mtype(m, Own)$ is defined then the following hold (and *mutatis mutandis* for $mtype'$):

— $\neg mscope(m, Own)$ implies $mtype'(m, Own) = mtype(m, Own)$ and moreover $\neg mscope'(m, Own)$, and

—$prot(m, Own)$ implies $mtype'(m, Own) = mtype(m, Own)$ and $mscope'(m, Own)$ (which in turn implies $prot'(m, Own)$).

***Example* 9.1** Method getFirst of Observable in Fig. 8 is called in subclass ObservableAcc, so *prot*(getFirst, Observable) holds and getFirst must be present in a comparable class table (and be simulated).

On the other hand, consider the loop in notifyAll in Observable of Fig. 2. One could desugar the loop to this code using a tail recursive helper method doNotif. The helper could be given module scope and need not be called in owner subclasses; in which case we set *prot*(doNotif, Observable) false. □

Results of Sections 5 and 6 hold for the extended language; the only proof affected by the changes is that of Theorem 6.15 which says that $[\![CT]\!]$ is confined if CT is confined. The result holds for the revised definitions.[12]

[12]The necessary revisions for the proof are as follows:

In the base case of the induction on depth, the argument proving confinement of $\mu_{i+1} Cm$ for the result value d goes by cases on C. The argument for the case $C \leq Own$ still holds for m with $\neg mscope(m, C)$. For the case $C \leq Own$ and $mscope(m, C)$, the revised definition requires the result value d to satisfy $d \in locs(Rep{\downarrow}) \Rightarrow d \in dom(Rh_j)$ for some confining partition and j

10. SECOND ABSTRACTION THEOREM

This section improves the first abstraction theorem in two ways. First, the result applies to the language extended with module-scoped methods (see Sect. 9.2). The module-scoped methods of the two versions of Own can be different unless they are used in subclasses of Own. The second improvement is that parametricity of the allocator is no longer required (cf. Sect. 7.3). To compare behaviors of two versions of a program we use a bijection between locations rather than equality. This can be seen as expressing that the language is parametric in locations, which would fail if the language had pointer arithmetic. As discussed in Sect. 9.1, bijections handle the problem with new reps in sub-owners that necessitates Assumption 7.15. Moreover, it allows coarsening of the notion of equivalence for commands and method meanings so that, for example, the bodies of the two versions of method version in Sect. 9.1 are equivalent.

***Definition* 10.1 (typed bijection)** A *typed bijection* is finite bijective function σ from $Locs$ to $Locs$ such that $\sigma\,\ell = \ell'$ implies $loctype\,\ell = loctype\,\ell'$. □

Throughout the section we let σ range over typed bijections and sometimes omit the word "typed". To express how bijections cut down to bijections on partition blocks, we use the notation $\sigma(X)$ for the direct image of X through σ.

***Definition* 10.2 (local coupling, $\mathcal{L}$)** Given comparable class tables, a local coupling is a function $\mathcal{L}$ that assigns to each typed bijection a binary relation $\mathcal{L}\,\sigma$ on heaps (not necessarily closed heaps) that satisfies the following. For any σ, h, h', if $\mathcal{L}\,\sigma\,h\,h'$ then there are partitions $h = Oh * Rh$ and $h' = Oh' * Rh'$ and locations ℓ and ℓ' in $locs(Own{\downarrow})$ such that

(1) $\sigma\,\ell = \ell'$ and $\{\ell\} = dom\;Oh$ and $\{\ell'\} = dom\;Oh'$
(2) $dom(Rh) \subseteq locs(Rep{\downarrow})$ and $dom(Rh') \subseteq locs(Rep'{\downarrow})$
(3) $\mathcal{R}\;\sigma\;(type(f, loctype\,\ell))\;(h\ell f)\;(h'\ell' f)$ for all $(f : T) \in dom(fields(loctype\,\ell))$ with $f \notin \overline{g} = dom(dfields(Own))$ and $f \notin \overline{g}' = dom(dfields'(Own))$. □

Item (3) uses the induced coupling $\mathcal{R}$ defined below; it is a harmless forward reference because the definition of $\mathcal{R}$ for data types does not depend on $\mathcal{R}$ (or $\mathcal{L}$) for heaps. Note that we do not require $dom\,\sigma$ to include the reps, nor do we disallow that it includes some of them.

***Definition* 10.3 (coupling relation, $\mathcal{R}$)** Suppose $\mathcal{L}$ is a local coupling. For each typed bijection σ we define relation $\mathcal{R}\;\sigma\;\theta \subseteq [\![\theta]\!] \times [\![\theta]\!]'$ as follows. In the case of method meanings and method environments the relation $\mathcal{R}$ is not parameterized on a bijection; rather it quantifies over all σ.

For heaps h, h', we define $\mathcal{R}\;\sigma\;Heap\;h\;h'$ iff there exist confining partitions of h, h', with the same number n of owner islands, such that

with $\eta\,\mathsf{self} \in dom(Oh_j)$. This follows by definition from $conf\;C\,(h_0, \eta_0)$.

In the step of the induction on depth, there is case analysis on C and B, proving claim $conf\;B\,(h, \eta)$ and confinement of the result value d. For the case $C \leq Own < B$, the argument still holds, noting that $\neg mscope(m, C)$ because in a well formed class table module-scoped methods do not occur outside owner and rep classes. For the cases $C < B \leq Own$ and $C < B \leq Rep$, the arguments still hold, noting that the restrictions on $mscope$ ensure $mscope(m, B) = mscope(m, C)$ so the relevant conditions are the same.

—$dom\,\sigma \subseteq dom\,h$ and $rng\,\sigma \subseteq dom\,h'$

—$\mathcal{L}\ \sigma\ (Oh_i * Rh_i)\ (Oh'_i * Rh'_i)$ for all i in $1..n$

—$\sigma(dom(Ch)) = dom(Ch')$, i.e., σ restricts to a bijection between $dom(Ch)$ and $dom(Ch')$

—$\mathcal{R}\ \sigma\ (state\,(loctype\ \ell))\ (h\ell)\ (h'\ell')$ for all ℓ, ℓ' with $\ell \in dom(Ch)$ and $\sigma\,\ell\,\ell'$

For other categories θ we define $\mathcal{R}\ \sigma\ \theta$ as follows.

$$
\begin{array}{ll}
\mathcal{R}\ \sigma\ \mathbf{bool}\ d\ d' & \Leftrightarrow\ d = d' \\
\mathcal{R}\ \sigma\ \mathbf{unit}\ d\ d' & \Leftrightarrow\ d = d' \\
\mathcal{R}\ \sigma\ C\ d\ d' & \Leftrightarrow\ \sigma\,d = d' \vee d = nil = d' \\
\mathcal{R}\ \sigma\ \Gamma\ \eta\ \eta' & \Leftrightarrow\ \forall x \in dom\,\Gamma \cdot \mathcal{R}\ \sigma\ (\Gamma x)\ (\eta x)\ (\eta' x) \\
\mathcal{R}\ \sigma\ (state\,C)\ s\ s' & \Leftrightarrow \\
\quad C \not\leq Own \wedge \forall f \in dom(fields\,C) \cdot \mathcal{R}\ \sigma\ (type(f, C))\ (s\ f)\ (s'\ f) & \\
\mathcal{R}\ \sigma\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta') & \Leftrightarrow\ \mathcal{R}\ \sigma\ Heap\ h\ h' \wedge \mathcal{R}\ \sigma\ \Gamma\ \eta\ \eta' \\
\mathcal{R}\ \sigma\ (Heap \otimes T)\ (h, d)\ (h', d') & \Leftrightarrow\ \mathcal{R}\ \sigma\ Heap\ h\ h' \wedge \mathcal{R}\ \sigma\ T\ d\ d' \\
\mathcal{R}\ \sigma\ (\theta_\bot)\ \alpha\ \alpha' & \Leftrightarrow\ (\alpha = \bot = \alpha') \vee (\alpha \neq \bot \neq \alpha' \wedge \mathcal{R}\ \sigma\ \theta\ \alpha\ \alpha') \\
\mathcal{R}\ (C, \overline{x}, \overline{T} \rightarrow T)\ d\ d' & \Leftrightarrow\ \forall \sigma,\ (h, \eta) \in [\![Heap \otimes \Gamma]\!], (h', \eta') \in [\![Heap \otimes \Gamma]\!]' \ \cdot \\
\quad \mathcal{R}\ \sigma\ (Heap \otimes \Gamma)\ (h, \eta)\ (h', \eta') \wedge conf\ C\ (h, \eta) \wedge conf\ C\ (h', \eta') & \\
\quad \Rightarrow \exists \sigma_0 \supseteq \sigma \cdot \mathcal{R}\ \sigma_0\ (Heap \otimes T)_\bot\ (d(h, \eta))\ (d'(h', \eta')) & \\
\quad \text{where } \Gamma = [\overline{x} \mapsto \overline{T}, \mathsf{self} \mapsto C] & \\
\mathcal{R}\ Menv\ \mu\ \mu' & \Leftrightarrow\ \forall C, m\ \cdot \\
\quad (\neg mscope(m, C) \vee prot(m, C)) \wedge (C \text{ is non-rep}) \wedge (mtype(m, C) \text{ is defined}) & \\
\quad \Rightarrow \mathcal{R}\ \ (C, pars(m, C), mtype(m, C))\ (\mu\,C\,m)\ (\mu'\,C\,m) \quad \Box &
\end{array}
$$

(Recall that *prot* is defined in (5) of Sect. 9.2.)

As an example, the body of makeNode in ObservableAcc (Fig. 9) returns a new rep. Consider a coupling with a version using a sentinel. Given a bijection σ and related heaps h, h', the location $\ell = fresh(\mathsf{Node4}, h)$ may be different from $\ell' = fresh(\mathsf{Node4}, h')$ even if *fresh* is parametric, because h' has extra reps, the sentinels. But σ can be extended with the pair (ℓ, ℓ').

The following facts are straightforward consequences of the definition. The first says that if h and h' are related by $\mathcal{R}$ at σ, then σ is a bijection between the domains of h and h' except for reps.

Lemma 10.4 For all σ, h, h' and all ℓ, ℓ' not in $locs(Rep\!\downarrow, Rep'\!\downarrow)$, if $\mathcal{R}\ \sigma\ Heap\ h\ h'$ then $\sigma((dom\,h) \downharpoonright (locs(Rep\!\downarrow, Rep'\!\downarrow))) = (dom\,h') \downharpoonright (locs(Rep\!\downarrow, Rep'\!\downarrow))$. $\Box$

Lemma 10.5 If $\overline{U} \leq \overline{T}$ and $\mathcal{R}\ \sigma\ \overline{U}\ \overline{d}\ \overline{d'}$ then $\mathcal{R}\ \sigma\ \overline{T}\ \overline{d}\ \overline{d'}$. $\Box$

For equivalence of values and states, we define a family of relations indexed on categories θ. To streamline the notation, we say "$x \sim_\sigma x'$ in $[\![\theta]\!]$" here, and simply use the symbol $\sim_\sigma$ later.

***Definition* 10.6 (value equivalence)** For any σ, we define a relation $\sim_\sigma$ for data

values, object states, heaps, and stores, as follows.

$$
\begin{array}{llll}
\ell \sim_\sigma \ell' & \text{in } [\![C]\!] & \Leftrightarrow & \sigma\,\ell = \ell' \vee \ell = nil = \ell' \\
d \sim_\sigma d' & \text{in } [\![T]\!] & \Leftrightarrow & d = d' \quad \text{for primitive types } T \\
s \sim_\sigma s' & \text{in } [\![state\,C]\!] & \Leftrightarrow & \forall f \in \mathit{fields}\,C \cdot s\,f \sim_\sigma s'\,f \\
\eta \sim_\sigma \eta' & \text{in } [\![\Gamma]\!] & \Leftrightarrow & \forall x \in dom\,\Gamma \cdot \eta\,x \sim_\sigma \eta'\,x \\
h \sim_\sigma h' & \text{in } [\![Heap]\!] & \Leftrightarrow & dom\,\sigma \subseteq dom\,h \wedge rng\,\sigma \subseteq dom\,h' \\
 & & & \wedge \forall \ell \in dom\,h \cdot h\,\ell \sim_\sigma h'(\sigma\,\ell) \\
(h,\eta) \sim_\sigma (h',\eta') & \text{in } [\![Heap \otimes \Gamma]\!] & \Leftrightarrow & h \sim_\sigma h' \wedge \eta \sim_\sigma \eta' \\
d \sim_\sigma d' & \text{in } [\![\theta_\perp]\!] & \Leftrightarrow & d = \perp = d' \vee (d \neq \perp \neq d' \wedge d \sim_\sigma d' \text{ in } [\![\theta]\!])
\end{array}
$$

Lemma 10.7 (identity extension) Suppose $\mathcal{R}\ \sigma\ (Heap \otimes \Gamma)\ (h,\eta)\ (h',\eta')$ and $\Gamma\,\mathsf{self}$ is non-rep. Let (h,η) and (h',η') be confined at $\Gamma\,\mathsf{self}$. If both $collect(\eta,h)$ and $collect(\eta',h')$ are Own-free then $collect(\eta,h) \sim_\sigma collect(\eta',h')$. □

In the case that σ is equality, the relations $\mathcal{R}\,\sigma\,\theta$ coincide with $\mathcal{R}\,\theta$ and $\sim_\sigma$ is just equality. This yields the analog of the Diagonal Lemma 7.18 as the reader may check.

***Definition* 10.8 (client program equivalence)** Suppose programs $CT, (\Gamma \vdash S)$ and $CT', (\Gamma \vdash' S')$ are such that CT, CT' are comparable and confined, and moreover S (resp. S') occurs in CT (resp. CT'). The programs are equivalent iff for all confined, Own-free (h,η) and (h',η') in $[\![Heap \otimes \Gamma]\!]$ and all σ with $(h,\eta) \sim_\sigma (h',\eta')$, there is some $\sigma_0 \supseteq \sigma$ with

$$collect([\![\Gamma \vdash S]\!]\hat{\mu}(h,\eta)) \sim_{\sigma_0} collect([\![\Gamma \vdash' S']\!]'\hat{\mu}'(h',\eta'))\ ,$$

where $\hat{\mu} = [\![CT]\!]$ and $\hat{\mu}' = [\![CT']\!]'$. □

Lemma 10.9 Suppose B, C and all class names in $\overline{T}$ are non-rep and moreover $B < C$. If $\mathcal{R}\ \ (C,\overline{x},\overline{T}{\to}T)\ d\ d'$ then $\mathcal{R}\ \ (B,\overline{x},\overline{T}{\to}T)\ (restr(d,B))\ (restr(d',B))$ where $restr$ is the restriction to global states of B (see Def. 5.5). □

As discussed in Sect. 9, the relation must be preserved not only by public methods but also by any module scope methods that are called by methods declared in subclasses of Own.

***Definition* 10.10 (simulation)** A simulation is a coupling relation $\mathcal{R}$ such that

(1) (construction of Own establishes $\mathcal{L}$) For any μ, μ', any ℓ, ℓ' in $locs(Own{\downarrow})$ with $\sigma\,\ell = \ell'$, and any h, h' with $\mathcal{R}\ \sigma\ Heap\ h\ h'$, let

$$
\begin{array}{l}
h_0 = [h \mid \ell \mapsto [\mathit{fields}(loctype\,\ell) \mapsto \mathit{defaults}]] \\
h'_0 = [h' \mid \ell' \mapsto [\mathit{fields}'(loctype\,\ell') \mapsto \mathit{defaults}]]
\end{array}
$$

Then there is $\sigma_0 \supseteq \sigma$ such that $\mathcal{L}\ \sigma\ h_0\ h'_0$.

(2) (methods of Own preserve $\mathcal{R}$) Let $\mu \in \mathbb{N} \to [\![Menv]\!]$ (resp. $\mu' \in \mathbb{N} \to [\![Menv]\!]'$) be the approximation chain in the definition of $[\![CT]\!]$ (resp. $[\![CT']\!]'$). For every m with $mtype(m, Own)$ defined and $\neg mscope(m, Own)$ or $prot(m, Own)$, the following implications hold for every i, where $\overline{x} = pars(m, Own)$ and $\overline{T}{\to}T = mtype(m, Own)$.

 (a) $\mathcal{R}\ Menv\ \mu_i\ \mu'_i\ \Rightarrow\ \mathcal{R}\ (Own,\overline{x},\overline{T}{\to}T)\ ([\![M]\!]\mu_i)\ ([\![M']\!]'\mu'_i)$
 if m has declaration M in $CT(Own)$ and M' in $CT'(Own)$

(b) $\mathcal{R}\ \mathit{Menv}\ \mu_i\ \mu'_i\ \Rightarrow \mathcal{R}\ (Own, \overline{x}, \overline{T}{\to}T)\ ([\![M]\!]\mu_i)\ (\mathit{restr}([\![M_B]\!]'\mu'_i, Own))$
if m has declaration M in $CT(Own)$ and is inherited from B in $CT'(Own)$, with M_B the declaration of m in B

(c) the condition symmetric to (2b), if m is inherited in $CT(Own)$ but declared in $CT'(Own)$ $\square$

Instead of Assumption 7.15 we need only the following.

Assumption 10.11 CT and CT' are confined class tables for which a (generalized) simulation $\mathcal{R}$ is given.

Theorem 10.12 (abstraction) $\mathcal{R}\ \mathit{Menv}\ [\![CT]\!]\ [\![CT']\!]'$.

The proof is essentially the same as the proof of Theorem 7.20. The definition of $\mathcal{R}\ \mathit{Menv}$ requires the relation to be preserved by those module-scoped methods that are called by subowners, and this is ensured by Def. 10.10(2) of simulation. The lemmas used in the proof are as follows.

Lemma 10.13 (preservation by expressions) Consider any non-rep class $C \neq Own$ and any $\Gamma \vdash e : T$ with $\Gamma\, \mathsf{self} = C$. If $\Gamma \vdash e : T$ is confined and all constituents of e are confined then the following holds: For all σ and all $(h, \eta) \in [\![\mathit{Heap} \otimes \Gamma]\!]$ and $(h', \eta') \in [\![\mathit{Heap} \otimes \Gamma]\!]'$, if $\mathcal{R}\ \sigma\ (\mathit{Heap} \otimes \Gamma)\ (h, \eta)\ (h', \eta')$ then

$$\mathcal{R}\ \sigma\ (T_\perp)\ ([\![\Gamma \vdash e : T]\!](h, \eta))\ ([\![\Gamma \vdash' e : T]\!]'(h', \eta'))\ .$$

Proof. Similar to the proof of Lemma 7.21. $\square$

Lemma 10.14 (preservation by commands) Suppose that μ and μ' are confined method environments and $\mathcal{R}\ \mathit{Menv}\ \mu\ \mu'$. Consider any non-rep class $C \neq Own$ and any $\Gamma \vdash S$ with $\Gamma\, \mathsf{self} = C$. If $\Gamma \vdash S$ is confined and all constituents of S are confined then the following holds: For any σ and any $(h, \eta) \in [\![\mathit{Heap} \otimes \Gamma]\!]$ and $(h', \eta') \in [\![\mathit{Heap} \otimes \Gamma]\!]'$, if $\mathit{conf}\ C\ (h, \eta)$, $\mathit{conf}\ C\ (h', \eta')$, and $\mathcal{R}\ \sigma\ (\mathit{Heap}{\otimes}\Gamma)\ (h, \eta)\ (h', \eta')$ then there is $\sigma_0 \supseteq \sigma$ such that

$$\mathcal{R}\ \sigma_0\ (\mathit{Heap} \otimes \Gamma)_\perp\ ([\![\Gamma \vdash S]\!]\mu(h, \eta))\ ([\![\Gamma \vdash' S]\!]'\mu'(h', \eta'))\ .$$

Proof. The proof is very similar to the proof of the corresponding Lemma 7.22 except in the cases of method call, field update, and most interestingly **new**. We no longer have the assumption of parametricity of the allocator, and we must consider construction of reps in sub-owners.

Case $\Gamma \vdash x := e.m(\overline{e})$. This goes through as before except for the case where $C < Own$. In that case, the called method may have module scope and this is why such methods (designated by *prot*) are included in the definition of $\mathcal{R}\ \mathit{Menv}$.

Case $\Gamma \vdash x := \mathbf{new}\ B$. By confinement of CT, this command is confined and hence the final states are confined: $\mathit{conf}\ C\ (h_0, \eta_0)$ and $\mathit{conf}\ C\ (h'_0, \eta'_0)$. We have $C \not\leq \mathit{Rep}$ and $C \neq Own$. Let $\ell = \mathit{fresh}(B, h)$ and $\ell' = \mathit{fresh}(B, h')$. Define $\sigma_0 = \sigma \cup \{(\ell, \ell')\}$. This makes σ_0 bijective because ℓ, ℓ' are fresh and $\mathcal{R}\ \sigma\ \mathit{Heap}\ h\ h'$ implies, by definition, that $\mathit{dom}\,\sigma \subseteq \mathit{dom}\,h$ and $\mathit{rng}\,\sigma \subseteq \mathit{dom}\,h'$.

By $\mathcal{R}\ \sigma\ \Gamma\ \eta\ \eta'$ and definition of σ_0 we have $\mathcal{R}\ \sigma_0\ \Gamma\ \eta_0\ \eta'_0$. We proceed to show $\mathcal{R}\ \sigma_0\ \mathit{Heap}\ h_0\ h'_0$, by cases on B. Let $h_1 = [\ell \mapsto [\mathit{fields}\, B \mapsto \mathit{defaults}]]$ and $h'_1 = [\ell' \mapsto [\mathit{fields}'\, B \mapsto \mathit{defaults}]]$.

—$B \not\leq Own \wedge B \not\leq Rep$: We have $\mathit{fields}\, B = \mathit{fields}'\, B$ and thus

$$\mathcal{R}\ \sigma_0\ (\mathit{state}\ B)\ [\mathit{fields}\, B \mapsto \mathit{defaults}]\ [\mathit{fields}'\, B \mapsto \mathit{defaults}]\ .$$

So, as B is non-rep and $B \neq Own$, we can add ℓ to Ch and ℓ' to Ch' to get partitions that witness $\mathcal{R}\ \sigma_0\ Heap\ h_0\ h'_0$.

—$B \leq Own$: By local coupling, Def. 10.2, we get σ_0 with $\mathcal{L}\, \sigma_0\, h_1\, h'_1$. Moreover, h_1 and h'_1 are owner islands and the confining partitions for h, h' extend to ones for $h * h_1$. and $h' * h'_1$ with σ_0. Finally, by definition of $\mathcal{R}$ we get $\mathcal{R}\ \sigma_0\ Heap\ h_0\ h'_0$ as $h_0 = h * h_1$ and $h'_0 = h' * h'_1$.

—$B \leq Rep$: Here, $C \leq Own$ or $C \leq Rep$, as otherwise the command would not be confined. Let j be such that $\eta\, \mathsf{self} \in dom(Oh_j * Rh_j)$. Add ℓ to Rh_j and ℓ' to Rh'_j. This yields $\mathcal{R}\ \sigma_0\ Heap\ h_0\ h'_0$ with $h_0 = h * h_1$ and $h'_0 = h' * h'_1$. $\square$

11. STATIC ANALYSIS

This section gives a syntax directed static analysis for a property we call safety which is shown to imply confinement.

The input is a well formed class table and designated class names *Own* and *Rep*. The analysis is given for the language of Sect. 9.2. We do not claim that the analysis is definitive, but it has the following pleasant characteristics. The restrictions are weak enough to admit interesting programs including all the examples discussed in the paper. The analysis is modular in the sense that, with one exception, only rep and owner code (including subclasses) is constrained. The exception is for **new**: a client cannot construct a new rep. For practical application, this can be ensured in a modular way: *Rep* and its subclasses would be declared with module scope.

***Definition* 11.1 (*safe*)** Class table CT is *safe* iff for every C and every m with $mtype(m, C) = \overline{T} \rightarrow T$ the following hold.

(1) If m is declared in C by $T\ m(\overline{T}\ \overline{x})\{S\}$ then $\overline{x} : \overline{T}, \mathsf{self} : C, \mathsf{result} : T \rhd S$ where $\rhd$ is the safety relation defined in the sequel.
(2) If $C \leq Own$ and $\neg mscope(m, C)$ then $T \not\geq Rep$.
(3) If m is inherited in Own from some $B > Own$ then $\overline{T} \not\geq Rep$.
(4) No m is inherited in Rep from any $B > Rep$.

The safety relation $\rhd$ is defined by the following rules. There is no restriction on field declarations per se. A client can have a *Rep* type field, but can assign only **null** to it.

Safety for expressions

$$\Gamma \rhd x : \Gamma x \quad \Gamma \rhd \mathbf{null} : B \quad \Gamma \rhd \mathbf{unit} : \mathbf{unit} \quad \Gamma \rhd \mathbf{true} : \mathbf{bool} \quad \Gamma \rhd \mathbf{false} : \mathbf{bool}$$

$$\frac{C = (\Gamma\, \mathsf{self}) \quad \Gamma \rhd e : C \quad (f : T) \in \mathit{dfields}\, C \quad (C = Own \wedge e \neq \mathsf{self} \Rightarrow T \not\geq Rep)}{\Gamma \rhd e.f : T}$$

$$\frac{\Gamma \rhd e_1 : T \quad \Gamma \rhd e_2 : T}{\Gamma \rhd e_1 = e_2 : \mathbf{bool}} \qquad \frac{\Gamma \rhd e : D \quad B \leq D}{\Gamma \rhd (B)\ e : B} \qquad \frac{\Gamma \rhd e : D \quad B \leq D}{\Gamma \rhd e\ \mathbf{is}\ B : \mathbf{bool}}$$

For expressions, the analysis imposes restrictions on field accesses and nothing else. If $e.f$ appears in the body of an owner method, then a *Rep* can be accessed only via the private fields of *Own*; this requires e to be self (instance-based visibility).[13] If $e.f$ appears in a sub-owner, then the private fields of *Own* cannot be accessed, hence the result cannot be a *Rep*.

For commands, the rules impose restrictions on **new**, field update, and method call. The conditions on field update are analogous to those for field access. Object construction $x := \mathbf{new}\ B$ in the body of a client method cannot create a new rep.

For method call $x := e.m(\overline{e})$, the condition labelled (a) says that if m is a client method called from a subclass of *Own* or *Rep*, then m cannot be passed reps as parameters. Condition (b) considers method calls from an owner class or its subclasses: it says that if m's type is comparable to *Own* then reps can be passed as parameters only if e is self. Moreover, if e is not self, then the method cannot return reps as result.

Safety for commands

$$\frac{C = (\Gamma\,\mathsf{self}) \quad B \neq \mathbf{Object} \qquad x \neq \mathsf{self} \quad B \leq \Gamma x \qquad C \not\leq \mathit{Rep} \wedge C \not\leq \mathit{Own} \Rightarrow B \not\leq \mathit{Rep}}{\Gamma \rhd x := \mathbf{new}\ B}$$

$$\frac{C = (\Gamma\,\mathsf{self}) \quad (f:T) \in \mathit{dfields}\,C \qquad \Gamma \rhd e_1 : C \quad \Gamma \rhd e_2 : U \quad U \leq T \qquad C = \mathit{Own} \wedge e_1 \neq \mathsf{self} \Rightarrow U \not\lesseqgtr \mathit{Rep} \qquad C < \mathit{Own} \Rightarrow U \not\lesseqgtr \mathit{Rep}}{\Gamma \rhd e_1.f := e_2}$$

$$\frac{\begin{array}{l} \Gamma \rhd e : D \quad \Gamma \rhd \overline{e} : \overline{U} \quad \mathit{mtype}(m,D) = \overline{T} \rightarrow T \quad \overline{U} \leq \overline{T} \quad T \leq \Gamma\,x \quad x \neq \mathsf{self} \\ C = (\Gamma\,\mathsf{self}) \qquad \mathit{mscope}(m,D) \Rightarrow C \leq \mathit{Own} \vee C \leq \mathit{Rep} \\ (a) \quad (C \leq \mathit{Own} \vee C \leq \mathit{Rep}) \wedge (D \not\leq \mathit{Rep} \wedge (D \not\leq \mathit{Own} \vee \neg(\mathit{mscope}(m,D)))) \\ \qquad \Rightarrow \overline{T} \not\lesseqgtr \mathit{Rep} \\ (b) \quad C \leq \mathit{Own} \Rightarrow D \not\lesseqgtr \mathit{Own} \vee (e = \mathsf{self}) \vee (\overline{T} \not\lesseqgtr \mathit{Rep} \wedge T \not\lesseqgtr \mathit{Rep}) \end{array}}{\Gamma \rhd x := e.m(\overline{e})}$$

$$\frac{x \neq \mathsf{self} \quad \Gamma \rhd e : T \quad T \leq \Gamma\,x}{\Gamma \rhd x := e} \qquad \frac{\Gamma \rhd S_1 \quad \Gamma \rhd S_2}{\Gamma \rhd S_1;\ S_2}$$

$$\frac{\Gamma \rhd e : \mathbf{bool} \quad \Gamma \rhd S_1 \quad \Gamma \rhd S_2}{\Gamma \rhd \mathbf{if}\ e\ \mathbf{then}\ S_1\ \mathbf{else}\ S_2\ \mathbf{fi}} \qquad \frac{\Gamma \rhd e : U \quad U \leq T \quad (\Gamma, x:T) \rhd S}{\Gamma \rhd T\ x := e\ \mathbf{in}\ S}$$

Theorem 11.2 (soundness) If CT is safe then it is confined.

PROOF. Items (2)–(4) in the definition of safety are the same as items (2)–(4) in the definition of confinement for class tables. For item (1), the confinement of method bodies follows from safety thereof, by Lemmas 11.3, 11.4, and 11.5 to follow. □

The proofs of the following lemmas are routine and appear in the companion technical report.

[13] An expression like self.fst.nxt.nxt is not allowed in our language because we consider only private fields; but if module-scoped fields were added, this expression would be allowed by the analysis.

Lemma 11.3 (argument values confined) Suppose $\Gamma \vdash e : D$ and $\Gamma \vdash \overline{e} : \overline{U}$ are confined. If $\Gamma \rhd x := e.m(\overline{e})$ then $\Gamma \vdash x := e.m(\overline{e})$ has confined arguments.

Lemma 11.4 (soundness for expressions) If $\Gamma \rhd e : T$ then $\Gamma \vdash e : T$ is confined.

Lemma 11.5 (soundness for commands) If $\Gamma \rhd S$ then $\Gamma \vdash S$ is confined.

12. DISCUSSION AND RELATED WORK

Programmers draw pictures of pointers in heap-based data structures and often manage to get things right as far as the presence of pointers goes. For example, lists don't get disconnected. The absence of pointers is harder to picture and many bugs are due to unexpected aliasing. Expectations are raised through use of encapsulation constructs such as private fields and modules, but heap structure is not entirely manifested in language constructs. Simulation relations are often used for reasoning about abstractions and here too aliasing presents a challenge: Multiple instances of an abstraction may reference a shared client object or be shared by multiple clients —but client references to representation objects can violate encapsulation. Various notions of ownership confinement have been proposed for encapsulation of objects. We have formalized one and shown that clients are independent from confined representations. Independence is formalized by an abstraction theorem that licenses reasoning about equivalence of class implementations using simulation relations. Confinement is formalized by drawing boundaries that signify the absence of pointers.

12.1 Related work

Representation independence. The main proof technique for representation independence is so fundamental that it has appeared in many places, with a variety of names, e.g., bisimulation, logical relations, abstraction mappings, relational parametricity (e.g., [Plotkin 1973; Reynolds 1984; Lynch and Vaandrager 1995; de Roever and Engelhardt 1998]). Among the many uses of simulations are program transformations and justification of logics for reasoning about data abstraction and modification of encapsulated state.

Representation independence results are known for general transition systems [Milner 1971; Lynch and Vaandrager 1995], first order imperative languages [He et al. 1986; de Roever and Engelhardt 1998], higher order functional [Reynolds 1984; Mitchell 1986; 1991; 1996; Power and Robinson 2000; Pitts 2000] and higher order imperative languages [O'Hearn and Tennent 1995; Pitts 1997; Naumann 2002], and sequential object-oriented programs without heap allocation ([Cavalcanti and Naumann 2002] treats a language with class-based visibility and [Reddy 2002] treats one with instance-based visibility). As far as we know, our results are the first for shared references to mutable state, a ubiquitous feature in object-oriented and imperative programs. The lacuna is mentioned in [Grossman et al. 2000].

The combination of local state with higher order procedures makes it difficult to prove representation independence for Algol, where procedures can be passed as arguments but not assigned to state variables. The root problem for Algol semantics [Reynolds 1981b; O'Hearn and Tennent 1995] and proof rules [Olderog 1983; German et al. 1989] is the interaction between arbitrary nesting of variable and

procedure declarations and possibility of passing procedures as arguments. Representation independence has been proved, using denotational semantics based on possible worlds models [O'Hearn and Tennent 1995] and using operational semantics [Pitts 1997], on which we say more below.

In imperative languages like C and Modula-3, procedures can be passed as arguments and even stored in variables, but only if their free variables are in outermost scope. This restriction greatly simplifies implementation of the language and it makes it possible to use simple semantic models. Naumann [2002] uses such a model to prove an abstraction theorem and apply it to Meyer-Sieber examples.[14] The constructs of a Java-like language offer similar expressive power and also admits a simple model as we have shown in this paper.

To get an adequate induction hypothesis for an abstraction theorem, parametricity needs to be imposed on the latent effects of procedure abstractions, either as a property to be proved or as an intrinsic feature of the semantic model [Reynolds 1981b; O'Hearn and Tennent 1995]. It seems that these conditions are most easily expressed in terms of a denotational model, but if procedures can be stored in the heap on which they act, difficult domain equations must be solved.[15] Recursive data types also lead to nontrivial domain equations. Even if solutions can be found, they may be quite complex structures that are difficult to understand and work with. Nevertheless, a modern treatment of recursive domain equations offers some hope of progress [Reus and Streicher 2002; Reus 2003].[16]

Difficulties with denotational semantics led to considerable advances using operational semantics [Gordon and Pitts 1998]. For Idealized Algol, in which only integers can be stored in variables and there are no recursive types, Pitts [1997] formalizes logical relations using operational semantics and proves equivalences like the Meyer-Sieber example in our Sect. 3.1. Although complex domains are avoided, an operationally based notion of logical relation can be "far from straightforward" and "quite difficult" (in the words of Pitts [2005], discussing logical relations for existential types in a purely functional setting).[17]

One of the most relevant works using operational semantics is that of Grossman et al. [2000] where representation independence is approached using a dynamic notion of ownership by *principals* as in the security literature. To prove that clients are independent from the representation of an abstraction provided by a host program, a wrapper construct is used to tag code fragments with their owner (e.g., client or "host"), and to provide an opaque type for the client's view of the abstraction. This is a promising approach. However, the results so far only show "independence of evaluation" (reminiscent of nonintereference results in information flow security [Volpano et al. 1996; Abadi et al. 1999]) and do not provide a

[14]The simpler of their examples can be proved directly in the model without use of simulations [Naumann 2001].
[15]Recently Levy [2002] used functor categories to give a denotational model for a higher order language with pointers, but the model does not capture relational parametricity and the language has neither object-oriented features nor recursive types.
[16]The cited work also provides semantics that is compositional at the level of classes, whereas our semantics is given for a complete class table. Nonetheless, our result provides modular reasoning about a single class in the context of an arbitrary class table.
[17]For functional programs, see also Sumii and Pierce [2005].

general notion of simulation. Although Grossman et al. [2000] offer their work as a simpler alternative to domain theoretic semantics, the technical treatment is somewhat intricate by the time the language is extended to include references, recursive and polymorphic types.

Except for parametric polymorphism, we treat all these features, as well as others such as subclassing, dynamic binding, type tests and casts. Although Java syntax seems less elegant than, say, lambda calculus, it has several features that ease the difficulties. Owing to name-based type equivalence and subtyping, and the binding of methods to objects via their class, we can use a denotational model with quite simple domains and fixpoint definitions in the manner of Strachey [2000].

For applications in security and automated static checking, it is important to devise robust, comprehensible models that support not only the idealized languages of research studies but also the full languages used in practice. Denotational semantics has conceptual advantages, at least if the domains are simple enough to have a clear operational significance. However, we admit that our enthusiasm for the efficacy of denotational techniques has been tempered by the irritation of flushing out bugs in intricate definitions and induction hypotheses.

Our abstraction theorem and identity extension lemma can be used directly to prove equivalence of programs, where a program is a command in the context of a class table and designated class C. It would be reasonable to use a notion of equivalence based on field visibility: states would be equated if they are equal after hiding all fields except those visible in C. But this would beg the question whether hiding imposes encapsulation that is not intrinsic to the language. In this paper we use the finer equivalence on programs: for commands to be equivalent they must yield outcomes that are identical, up to renaming, after garbage collection. Thus encapsulation is formulated in terms of private fields and confined reps but the identity extension lemma is expressed, in effect, in terms of local variable blocks (in the style of, e.g., He et al. [1986]).

Besides the "client interface" provided by public methods and analogous to the interfaces studied in previous work on representation independence, a class also has a "protected" interface to its subclasses. The combination of protected and public interfaces is complicated, but a thorough treatment of representation independence for object-oriented programs must take it into account. For reasoning about the protected interface, work on behavioral subclassing has used simulations to connect a class with its subclass [Liskov and Wing 1994; Leavens and Dhara 2000] but a formal connection has not been made with the use of simulations to connect alternative representations. The PhD thesis of Stata [1997] considers other aspects of the protected interface.

Confinement. Quite a few confinement disciplines have been proposed, by Hogg [1991], Almeida [1997],Vitek and Bokowski [2001], Clarke et al. [2001], Müller and Poetzsch-Heffter [2000b], Boyland [2001], Lea [2000], Aldrich et al. [2002], and Clarke [2001] (the latter has a more comprehensive recent survey). Most proposals have significant shortcomings; they disallow important design patterns or are not efficiently checkable. Although the aim is to achieve encapsulation and thereby support modular reasoning in one form or another, few proposals have been formally justified in these terms —none in terms of representation independence.

Several works justify a syntactic discipline by proving that it ensures a confinement invariant [Müller and Poetzsch-Heffter 2000b; Clarke 2001; Aldrich et al. 2002]. Others go further and show some form of modular reasoning principle, as we discuss in detail below. Existing justifications involve disparate techniques and objectives, so that it is quite hard to assess and compare confinement disciplines. One of our contributions is to show how standard semantic techniques can be used for such assessments.

The fact that type names are semantically relevant lets us use them to formulate in semantic terms a condition similar to the ownership confinement notions of Müller [2002], Clarke et al. [2001] and their predecessors [Hogg 1991; Almeida 1997]. Whereas several papers emphasize reachability via paths, our formulation of confinement emphasizes partitioning of heap objects and the one-step points-to relation. In this we were inspired by the work of Reynolds [2001] that shows the efficacy of reasoning about partition blocks that may have dangling pointers.

Reasoning on the assumption of confinement is a separate concern from enforcement or checking of confinement. Semantic considerations led us to a flexible, syntax-directed static analysis, but other analysis techniques such as model checking or theorem proving for (an approximation of) the semantic confinement property could be interesting.

It is interesting to note that we get a strong reasoning principle on the basis of ownership confinement alone, in a form that can be checked without program annotations. By contrast, other works use annotations and combine ownership with uniqueness and effects (e.g., read-only) [Clarke and Drossopoulou 2002; Aldrich et al. 2002; Müller 2002]. Those works aim to record design decisions in all parts of a program, to support program understanding and reasoning. We are concerned with *replacing* one part; for this purpose it is not clear what annotations would be appropriate as a design record. What is clear is that one only needs a three-way distinction —owner, rep, and other— which can be checked rather flexibly without recourse to special annotations or types.

Confinement figures heavily in the verification logics of Müller and Poetzsch-Heffter [2000a] and in some work by the group of Nelson and Leino [Leino and Nelson 2002; Detlefs et al. 1998] where it is needed for sound reasoning about the "modifies clause" framing the scope of effects. Subsequent to the present work, Clarke and Drossopoulou [2002] state results on reasoning about effects, using a confinement discipline imposed using code annotations for confinement and effects. These works are concerned with delimiting the scope of effects, which is an important aspect of modular reasoning, but they do not address representation independence.

There has been much work on capturing encapsulation via visibility (lexical scope), using existential types and subsumption (see [Bruce et al. 1999; Bruce 2002; Pierce 2002] and references therein). None of these works addresses the problem of confinement; they are concerned with the complex typing issues for object oriented languages.

One of the main difficulties in designing safe and flexible type systems is due to the desire to eliminate or minimize the use of type testing and casting which are seen as loopholes that subvert type-based encapsulation. Indeed, parametric

polymorphism has been much pursued as a means to cope with generic patterns that, in current practice, are usually coded using subsumption, casts, and type **Object** (a recent reference is the textbook by Bruce [2002]). Although parametric polymorphism has obvious merit, our results show that casts and type tests are themselves relationally parametric. It is behavioral subclassing which is at risk in some uses of casts and tests. This does not contradict Reynolds [1984] because our language has a nominal type system [Pierce 2002]; it is the name of a type, not its set of values, that is involved with tests and casts.

12.2 Future challenges

Our aim is to deal with the rich languages currently in use, rather than to advance language design. The language for which our results are given encompasses many important features of object oriented languages. Two major features are missing and will require substantial additional work: concurrency and parametric polymorphism. The interaction between parametric and subtyping polymorphism can be non-trivial and there are a number of competing type systems. Some languages, e.g., C++, have parametric polymorphism but with significant limitations; for Java, parametric types are a late addition.[18]

Ownership confinement is appropriate for reasoning about many designs in practice and we have shown through a series of examples that our notion is applicable to widely used designs such as the observer and factory patterns. Two important issues are beyond the reach of our work (and much of the previous work on confinement). The first is multiple ownership. A canonical example is a collection class with iterators. The reps for the collection are nodes of a data structure. The collection object mediates additions and deletions. To allow enumeration of elements of the collection it is common to use iterator objects which need access to the nodes of the data structure. It would seem that either the collection and its iterators share joint ownership of the reps or the iterator is given a special status distinct from the owner and from clients. Either way, multiple client-visible objects collaborate to provide an abstraction (the iterable collection).

Another example that requires something like multiple ownership is the Observer role of the observer pattern (vs. the Subject role), which has not been the focus of attention in this paper. Often an observer is comprised of several objects; in particular, the callback object on which notify is invoked is an instance of an inner class and might be considered part of the representation of the observer, though it is exposed to the Observable.

Ownership type systems have been given that allow some form of multiple owners [Clarke 2001; Müller 2002; Aldrich et al. 2002; Aldrich and Chambers 2004]. Although our formalization of islands can be extended easily to encompass multiple owners, it is not as clear how to extend the notion of simulation in a useful way. Our result formalizes the notion that an owner instance provides an abstraction and this is easily expressed in terms of the class construct. The generalization can probably be expressed by grouping the related owners (e.g., the collection class and

[18]In [Banerjee and Naumann 2004b] we extend the semantics straightforwardly to encompass parameterized types in the form found in the C# language [Kennedy and Syme 2001]. We also give a representation independence theorem. There is no major difficulty.

the iterator class) in a module, with local couplings generalized to encompass the multiple iterators associated with a collection. Ownership type systems are under active development. Once a robust standard notion emerges, it could provide an appropriate general setting to which our work could be adapted.

The other challenging issue for confinement is ownership transfer. Consider a queue that owns objects representing tasks to be performed. For load balancing, tasks may be moved from one queue to another. In this case a task is owned by just one queue at a time and in a given state the system is confined according to the definition in this paper. A sequential program for transferring ownership from one queue might look as follows: q2.task := q1.task; q1.task := **null**. From a confined initial state this need not lead to a confined final state: there could be other references to task. But it does lead to a confined final state if q2.task is initially the only existing reference to the task. Unique references have been extensively studied so let us assume that a static analysis is given for uniqueness. Even with uniqueness, our theory fails to apply, for two reasons. The first reason is a small one: in the intermediate state two different owners reference the same task. This problem is well known and can be surmounted: It is easy to add to our language an atomic command with the effect of the above sequence [Minsky 1996] and to show, given uniqueness, that it is confined. For practical purposes one would use a static analysis to check that q1.task is a dead expression [Boyland 2001].

The second reason our theory does not apply is a technical one. To show that a method call is confined, we need that the caller's environment is confined in the final heap assuming it was confined in the initial one. We get this by using a condition stronger than confinement: from a confined state, a command or method yields a final heap that extends the initial one in the sense of Def. 6.3. All commands of our language yield heaps extended in this sense so all method meanings have this property. (See the proof of Theorem 6.15.) But, by definition of extension, $h \trianglelefteq h_0$ says that reps that exist in h have the same owners in h_0 as in h, disallowing ownership transfer.[19]

For static analysis there are some more modest issues worthy of investigation. The simple conditions of Def. 6.9 ensure suitable confinement of the class table but they are unnecessarily strong. Methods inherited into rep classes are not risky if they do not leak self; such "anonymous methods" can be statically checked as shown by Vitek and Bokowski [2001] and Grothoff et al. [2001] in work on module-based confinement.[20] The conditions of our static analysis may also admit useful variations.

Having shown that simulation is sound one might proceed to study completeness. It is not the case that our confinement conditions are necessary in general for simulations to be preserved. A trivial simulation might depend on no confinement at all. Also, a rep could be leaked but not exploited by any client. One can see confinement as a kind of simulation which happens to be a rectangular predicate: h relates to h' just if h and h' are confined, independent of each other. This

[19]In work subsequent to this paper, we give an abstraction theorem in a setting where ownership is encoded in auxiliary state and can be transferred freely [Banerjee and Naumann 2005].

[20]In fact the cited work is concerned with pragmatic aspects of the analysis and does not formalize a semantic property ensured by the analysis.

suggests folding the confinement condition into the simulation relation, an idea which has been studied by Reddy and Yang for a Pascal-like language.[21] For practical reasoning the benefits of treating confinement separately are clear: it accords with informal design practice, is amenable to static checking, and ensures soundness for a straightforward and modular notion of coupling.

The more practical question is how to express local couplings and prove the simulation property for owner methods. To formalize the couplings for the observer examples one needs a formalism for inductive predicates on recursive data structures; separation logic appears promising for this purpose [Reynolds 2002].

Representation independence licenses reasoning about equivalence of programs that are structurally similar [Banerjee et al. 2001; Riecke 1993]. This is quite adequate for uses of simulations such as static analyses and relating alternative interpretations for primitives, such as the lazy and eager access control implementations for Java [Banerjee and Naumann 2002]. But for abstraction in program development, typically called data refinement, it is not uncommon to consider significantly different program structures. To establish the hypothesis of the theorem in this case requires a full program logic. Indeed, the theorem would then provide one of the proof rules. For first-order imperative languages, several proof systems have been given for reasoning about two versions of an abstraction [de Roever and Engelhardt 1998]. Typically, relations (especially "abstraction functions") are used to derive from one version the specification of the other version, which is then proved correct in a program logic. Logics for imperative object-oriented languages are at an early stage of development [Abadi and Leino 1997; Cavalcanti and Naumann 1999; Poetzsch-Heffter and Müller 1999; Huisman and Jacobs 2000; Huisman 2002; Reynolds 2002].

ACKNOWLEDGMENTS

Our work benefited from discussions with a number of people. For particularly useful technical help and encouragement we thank Torben Amtoft, Steve Bloom, Paulo Borba, Sophia Drossopoulou, Nevin Heintze, Tony Hoare, Doug Lea, K. Rustan M. Leino, Peter Müller, Peter O'Hearn, Uday Reddy, John Reynolds, David Schmidt, and Hongseok Yang. The anonymous POPL and JACM reviewers provided many helpful expository suggestions as well as technical corrections and improvements. Peter O'Hearn and colleagues at Queen Mary provided a stimulating and congenial environment during August 2002 when we drafted the journal version.

REFERENCES

Abadi, M., Banerjee, A., Heintze, N., and Riecke, J. G. 1999. A core calculus of dependency. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM Press, 147–160.

[21] Their aim is to explicate the semantic structure of languages involving heap storage. Their approach should lead to a lucid account on par with parametricity models for other languages [Reynolds 1984; 1981b; Reddy 2002]. They have defined a parametricity semantics for a Pascal-like language [Reddy and Yang 2004] in which heap cells are tuples of pointers and integers rather than objects with scoped fields. Several challenges remain to be addresssed, if this approach is to provide a foundation for reasoning about instance-based abstractions in Java-like languages using a practical confinement discipline. For example, nominal types and class-based visibility (which is not modelled by naive use of existential types).

Abadi, M. and Cardelli, L. 1996. *A Theory of Objects.* Springer-Verlag.

Abadi, M. and Leino, K. R. M. 1997. A logic of object-oriented programs. In *Theory and Practice of Software Development (TAPSOFT)*. Springer-Verlag. Expanded in DEC SRC report 161.

Aldrich, J. and Chambers, C. 2004. Ownership domains: Separating aliasing policy from mechanism. In *European Conference on Object Oriented Programming (ECOOP)*. Lecture Notes in Computer Science. Springer-Verlag, 1–25.

Aldrich, J., Kostadinov, V., and Chambers, C. 2002. Alias annotations for program understanding. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*. ACM Press.

Almeida, P. S. 1997. Balloon types: Controlling sharing of state in data types. In *European Conference on Object Oriented Programming (ECOOP)*. Lecture Notes in Computer Science. Springer-Verlag, 32–59.

Arnold, K. and Gosling, J. 1998. *The Java Programming Language, second edition.* Addison-Wesley.

Banerjee, A., Heintze, N., and Riecke, J. G. 2001. Design and correctness of program transformations based on control-flow analysis. In *Intl. Symp. on Theoretical Aspects of Computer Software (TACS)*. Lecture Notes in Computer Science. Springer-Verlag, 420–447.

Banerjee, A. and Naumann, D. A. 2002. Representation independence, confinement and access control. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM Press, 166–177.

Banerjee, A. and Naumann, D. A. 2004a. Ownership confinement ensures representation independence for object-oriented programs. Tech. Rep. CS-2004-14, Stevens Institute of Technology. Also available as technical report KSU CIS-TR-2004-6, Kansas State University.

Banerjee, A. and Naumann, D. A. 2004b. State based encapsulation and generics. Tech. Rep. CS Report 2004-11, Stevens Institute of Technology.

Banerjee, A. and Naumann, D. A. 2005. State based ownership, reentrance, and encapsulation. In *European Conference on Object Oriented Programming (ECOOP)*. to appear.

Bhowmik, A. and Pugh, W. 1999. A secure implementation of Java inner classes. *PLDI* poster session, `http://www.cs.umd.edu/~pugh/java/SecureInnerClassesHandout.pdf`.

Boyland, J. 2001. Alias burying: Unique variables without destructive reads. *Software Practice and Experience 31,* 6, 533–553.

Bruce, K. B. 2002. *Foundations of Object-Oriented Programming Languages: Types and Semantics.* MIT Press.

Bruce, K. B., Cardelli, L., and Pierce, B. C. 1999. Comparing object encodings. *Information and Computation 155,* 1/2, 108–133.

Cavalcanti, A. L. C. and Naumann, D. A. 1999. A weakest precondition semantics for an object-oriented language of refinement. In *FM'99 - Formal Methods, Volume II.* Number 1709 in Lecture Notes in Computer Science. Springer-Verlag, 1439–1459.

Cavalcanti, A. L. C. and Naumann, D. A. 2002. Forward simulation for data refinement of classes. In *Formal Methods Europe*. Lecture Notes in Computer Science, vol. 2391. Springer-Verlag, 471–490.

Clarke, D. 2001. Object ownership and containment. Ph.D. thesis, University of New South Wales, Australia.

Clarke, D. and Drossopoulou, S. 2002. Ownership, encapsulation and the disjointness of type and effect. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*. ACM Press.

Clarke, D. G., Noble, J., and Potter, J. M. 2001. Simple ownership types for object containment. In *European Conference on Object Oriented Programming (ECOOP)*. Lecture Notes in Computer Science. Springer-Verlag.

Cormen, T. H., Leiserson, C. E., and Rivest, R. L. 1990. *Introduction to Algorithms.* MIT Press.

Cousot, P. and Cousot, R. 1977. Automatic synthesis of optimal invariant assertions: mathematical foundations. In *Proceedings of the ACM Symposium on Artificial Intelligence and Programming Languages, SIGPLAN Notices*. Vol. 12. ACM Press, 1–12.

Dahl, O.-J. and Nygaard, K. 1966. Simula: an Algol-based simulation language. *Communications of the ACM 9,* 9, 671–678.

Davey, B. and Priestley, H. 1990. *Introduction to Lattices and Order*. Cambridge University Press.

de Roever, W.-P. and Engelhardt, K. 1998. *Data Refinement: Model-Oriented Proof Methods and their Comparison*. Cambridge University Press.

Detlefs, D. L., Leino, K. R. M., and Nelson, G. 1998. Wrestling with rep exposure. Research Rep. 156, DEC Systems Research Center.

Dhara, K. K. and Leavens, G. T. 1996. Forcing behavioral subtyping through specification inheritance. In *Proceedings of the 18th International Conference on Software Engineering*. IEEE Computer Society Press, 258–267.

Donahue, J. E. 1979. On the semantics of "data type". *SIAM Journal of Computing 8,* 4, 546–560.

Gamma, E., Helm, R., Johnson, R., and Vlissides, J. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley.

German, S. M., Clarke, E. M., and Halpern, J. Y. 1989. Reasoning about procedures as parameters in the language L4. *Information and Computation 83*, 265–359.

Gong, L. 1999. *Inside Java 2 Platform Security*. Addison-Wesley.

Gordon, A. D. and Pitts, A. M., Eds. 1998. *Higher Order Operational Techniques in Semantics*. Cambridge University Press.

Grossman, D., Morrisett, G., and Zdancewic, S. 2000. Syntactic type abstraction. *ACM Trans. Program. Lang. Syst. 22,* 6, 1037–1080.

Grothoff, C., Palsberg, J., and Vitek, J. 2001. Encapsulating objects with confined types. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*. ACM Press.

Haynes, C. T. 1984. A theory of data type representation independence. In *International Symposium on Semantics of Data Types*, G. Kahn, D. B. MacQueen, and G. Plotkin, Eds. Lecture Notes in Computer Science, vol. 173. Springer-Verlag, 157–175.

He, J., Hoare, C. A. R., and Sanders, J. 1986. Data refinement refined (resumé). In *European Symposium on Programming*. Lecture Notes in Computer Science, vol. 213. Springer-Verlag.

Hoare, C. A. R. 1972. Proofs of correctness of data representations. *Acta Inf. 1*, 271–281.

Hogg, J. 1991. Islands: Aliasing protection in object-oriented languages. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*. ACM Press.

Hogg, J., Lea, D., Wills, A., deChampeaux, D., and Holt, R. 1992. The Geneva Convention on the treatment of object aliasing. *OOPS Messenger 3,* 2, 11–16.

Huisman, M. 2002. Verification of Java's AbstractCollection class: A case study. In *Mathematics of Program Construction*. Lecture Notes in Computer Science, vol. 2386. Springer-Verlag, 175–194.

Huisman, M. and Jacobs, B. 2000. Java program verification via a Hoare logic with abrupt termination. In *Fundamental Approaches to Software Engineering (FASE)*. Lecture Notes in Computer Science. Springer-Verlag, 284–303.

Igarashi, A., Pierce, B., and Wadler, P. 2001. Featherweight Java: A minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst. 23,* 3 (May), 396–459.

Jones, C. B. 1986. *Systematic software development using VDM*. International Series in Computer Science. Prentice-Hall.

Kennedy, A. and Syme, D. 2001. Design and implementation of generics for the .NET Common Language Runtime. In *PLDI*. 1–12.

Lea, D. 2000. *Concurrent Programming in Java*, Second ed. Addison-Wesley.

Leavens, G. T. and Dhara, K. K. 2000. Concepts of behavioral subtyping and a sketch of their extension to component-based systems. In *Foundations of Component-Based Systems*, G. T. Leavens and M. Sitaraman, Eds. Cambridge University Press, Chapter 6, 113–135.

LEINO, K. R. M. AND NELSON, G. 2002. Data abstraction and information hiding. *ACM Trans. Program. Lang. Syst. 24,* 5, 491–553.

LEVY, P. 2002. Possible world semantics for general storage in call-by-value. In *Computer Science Logic*. Number 2471 in Lecture Notes in Computer Science. Springer-Verlag.

LISKOV, B. AND GUTTAG, J. 1986. *Abstraction and Specification in Program Development.* MIT Press.

LISKOV, B. H. AND WING, J. M. 1994. A behavioral notion of subtyping. *ACM Trans. Program. Lang. Syst. 16,* 6, 1811–1841.

LYNCH, N. AND VAANDRAGER, F. 1995. Forward and backward simulations part I: Untimed systems. *Information and Computation 121,* 2, 214–233.

MEYER, A. R. AND SIEBER, K. 1988. Towards fully abstract semantics for local variables: Preliminary report. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM Press, 191–203.

MILNER, R. 1971. An algebraic definition of simulation between programs. In *Proceedings of Second Intl. Joint Conf. on Artificial Intelligence*. 481–489.

MINSKY, N. H. 1996. Towards alias-free pointers. In *European Conference on Object Oriented Programming (ECOOP)*. Lecture Notes in Computer Science, vol. 1098. 189–209.

MITCHELL, J. C. 1986. Representation independence and data abstraction. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM Press, 263–276.

MITCHELL, J. C. 1991. On the equivalence of data representations. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, V. Lifschitz, Ed. 305–330.

MITCHELL, J. C. 1996. *Foundations for Programming Languages.* MIT Press.

MÜLLER, P. 2002. *Modular Specification and Verification of Object-Oriented programs.* Lecture Notes in Computer Science, vol. 2262. Springer-Verlag.

MÜLLER, P. AND POETZSCH-HEFFTER, A. 2000a. Modular specification and verification techniques for object-oriented software components. In *Foundations of Component-Based Systems*, G. T. Leavens and M. Sitaraman, Eds. Cambridge University Press.

MÜLLER, P. AND POETZSCH-HEFFTER, A. 2000b. A type system for controlling representation exposure in Java. In *ECOOP Workshop on Formal Techniques for Java Programs.* Technical Report 269, Fernuniversität Hagen.

NAUMANN, D. A. 2001. Predicate transformer semantics of a higher order imperative language with record subtyping. *Sci. Comput. Program. 41,* 1, 1–51.

NAUMANN, D. A. 2002. Soundness of data refinement for a higher order imperative language. *Theor. Comput. Sci. 278,* 1–2, 271–301.

NAUMANN, D. A. 2005. Verifying a secure information flow analyzer. In *Theorem Proving in Higher Order Logics (TPHOLS).* to appear.

O'HEARN, P. W. AND TENNENT, R. D. 1995. Parametricity and local variables. *Journal of the ACM 42,* 3, 658–709.

OLDEROG, E.-R. 1983. Hoare's logic for programs with procedures — what has been achieved? In *Proceedings, Logics of Programs*, E. Clarke and D. Kozen, Eds. Lecture Notes in Computer Science, vol. 164. Springer-Verlag.

PIERCE, B. C. 2002. *Types and Programming Languages.* MIT Press.

PITTS, A. M. 1997. Reasoning about local variables with operationally-based logical relations. In *Algol-Like Languages*, P. W. O'Hearn and R. D. Tennent, Eds. Vol. 2. Birkhauser, Chapter 17, 173–193. Reprinted from *Proceedings Eleventh Annual IEEE Symposium on Logic in Computer Science*, Brunswick, NJ, July 1996, pp 152–163.

PITTS, A. M. 2000. Parametric polymorphism and operational equivalence. *Mathematical Structures in Computer Science 10*, 321–359.

PITTS, A. M. 2005. Typed operational reasoning. In *Advanced Topics in Types and Programming Languages*, B. C. Pierce, Ed. The MIT Press, Chapter 7, 245–289.

PLOTKIN, G. 1973. Lambda definability and logical relations. Tech. Rep. SAI-RM-4, University of Edinburgh, School of Artificial Intelligence.

Poetzsch-Heffter, A. and Müller, P. 1999. A programming logic for sequential Java. In *Programming Languages and Systems (ESOP)*, S. D. Swierstra, Ed. Lecture Notes in Computer Science, vol. 1576. Springer-Verlag, 162–176.

Power, A. J. and Robinson, E. P. 2000. Logical relations and data abstraction. In *Proceedings of Computer Science Logic (CSL)*, P. Clote and H. Schwichtenberg, Eds. Lecture Notes in Computer Science. Springer-Verlag, 497–511.

Reddy, U. S. 2002. Objects and classes in Algol-like languages. *Information and Computation 172,* 1 (Jan), 63–97.

Reddy, U. S. and Yang, H. 2004. Correctness of data representations involving heap data structures. *Sci. Comput. Program. 50,* 1–3, 129–160.

Reus, B. 2003. Modular semantics and logics of classes. In *Seventeenth International Workshop on Computer Science Logic (CSL)*. Lecture Notes in Computer Science, vol. 2803. Springer-Verlag, 456–469.

Reus, B. and Streicher, T. 2002. Semantics and logics of objects. In *IEEE Symposium on Logic in Computer Science (LICS)*. 113–124.

Reynolds, J. C. 1972. Definitional interpreters for higher-order programming languages. In *Proceedings of 25th ACM National Conference*. ACM Press, 717–740.

Reynolds, J. C. 1974. Towards a theory of type structure. In *Colloques sur la Programmation*. Lecture Notes in Computer Science, vol. 19. 408–425.

Reynolds, J. C. 1978. User-defined types and procedural data structures as complementary approaches to data abstraction. In *Programming Methodology*, D. Gries, Ed. Springer-Verlag, 309–317.

Reynolds, J. C. 1981a. *The Craft of Programming*. Prentice-Hall.

Reynolds, J. C. 1981b. The essence of Algol. In *Algorithmic Languages*, J. W. de Bakker and J. C. van Vliet, Eds. North-Holland.

Reynolds, J. C. 1984. Types, abstraction, and parametric polymorphism. In *Information Processing '83*, R. Mason, Ed. North-Holland, 513–523.

Reynolds, J. C. 2001. Intuitionistic reasoning about shared mutable data structure. In *Millenial Perspectives in Computer Science*. Palgrave.

Reynolds, J. C. 2002. Separation logic: a logic for shared mutable data structures. In *IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society Press.

Riecke, J. G. 1993. Fully abstract translations between functional languages. *Mathematical Structures in Computer Science 3,* 4, 387–415.

Stata, R. 1997. Modularity in the presence of subclassing. Research Report 145, DEC SRC, 130 Lytton Avenue Palo Alto, CA 94301.

Strachey, C. 2000. Fundamental concepts in programming languages. *Higher Order and Symbolic Computation 13,* 1, 11–49. Originally appeared in 1967 Lecture notes, International Summer School in Computer Programming, Copenhagen.

Sumii, E. and Pierce, B. C. 2005. A bisimulation for type abstraction and recursion. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM Press, 63–74.

Vitek, J. and Bokowski, B. 2001. Confined types in Java. *Software Practice and Experience 31,* 6, 507–532.

Volpano, D., Smith, G., and Irvine, C. 1996. A sound type system for secure flow analysis. *Journal of Computer Security 4,* 3, 167–187.

Wallach, D., Appel, A., and Felten, E. 2000. SAFKASI: a security mechanism for language-based systems. *ACM Transactions on Software Engineering and Methodology 9,* 4 (Oct.), 341–378.

Winskel, G. 1993. *The Formal Semantics of Programming Languages*. MIT Press.

Received Month Year; revised Month Year; accepted Month Year