

A Logical Account of Hoare’s Mismatch

Information Hiding via Second Order Framing in Region Logic

Anindya Banerjee

IMDEA Software, Madrid, Spain
anindya.banerjee@imdea.org

David A. Naumann

Stevens Institute of Technology, Hoboken NJ, USA
naumann@cs.stevens.edu

Abstract

We investigate information hiding in object-based programs and the associated mismatch. While client reasoning is in terms of interface specifications, the implementation of an interface is verified against different specifications that involve invariants about internal data structures. Soundness of this mismatched reasoning depends on encapsulation of internal data structures. The problem is that encapsulation is notoriously difficult to achieve in contemporary software in which shared mutable objects are ubiquitous. We account for the mismatch via proof rules that phrase the mismatch using explicit conditions that are imposed on client effects. Effects are tracked using ghost state and separation assertions in a style that has been used in a number of verification tools. Our approach permits the formulation of encapsulation disciplines (such as ownership, or package confinement) as part of the interface specification, in the form of a dynamic boundary, rather than as a discipline directly baked into a verifier. One implication is that we can provide a foundation for the axiomatic semantics of these verifiers. Our approach is flexible in that disciplines can be used on a per-module basis and then combined to achieve end to end soundness.

1. Introduction

Many programs manipulate shared, mutable objects, be they memory segments in C code, heap records in ML code, or objects in the sense of Java-like languages. Most software designs are based on abstraction at many scales. For example: A *Point* offers operations such as repositioning without revealing whether the internal representation uses polar or cartesian coordinates. A *Collection* offers addition and removal of elements without revealing whether the internal representation is a balanced binary tree or something else. An application framework (e.g., the Google Web Toolkit) offers complex concepts and facilities without revealing the still more complex underlying infrastructure. Abstraction is achieved by hiding irrelevant details, in particular internal data and invariants on which the implementation relies. Hiding of internal invariants means that they do not appear in specifications used by “clients” of the abstraction. This mismatch between what is verified about implementations and what is assumed by clients was articulated and justified by Hoare: invariants should depend only on data that is encapsulated to prevent clients from falsifying the invariants [17].

It is notoriously difficult to achieve encapsulation in the presence of shared, dynamically allocated mutable objects [20, 27]. Current tools for software verification either do not support hiding of invariants (e.g., Jahob [35], jStar [12]), do not treat object invariants soundly (e.g., ESC/Java [14], Eiffel) or at best offer soundness for restricted situations where hierarchical structure can be imposed on the heap (e.g. Spec# [4]) which is the best that is provided by current theory.

Scoping does provide some encapsulation. Class *Point* declares its fields as private, i.e., visible only to code within the class. Classes pertinent to collections — *Set*, *List*, iterators, various data structures— may be grouped together in a module. Scoping mechanisms are essential but insufficient for reasoning about dynamically allocated mutable objects.

Building on scope and type mechanisms it is possible to enforce module level alias control (as in Confined Types [15]) to justify hiding of invariants that depend on some or all instances of some classes declared in the module. For encapsulation at the granularity of instance-oriented “object invariants”, various ownership disciplines have been introduced [9, 7, 11, 25, 23]. Some have been deployed in verification systems, e.g., in JML tools [8] and in Spec# [4]. For example, suppose for simplicity that an instance of class *Set* is viewed as owning the nodes of a linked list and a type system enforces that neither clients nor other instances of *Set* can access these nodes. Code in *Set* relies on an invariant, such as the absence of duplicate elements, that can be hidden from clients. The inflexibility of various formalizations of ownership has led to quite a few competing and incompatible variations, each described in a way that requires its global imposition on all program components (e.g., [9, 7, 11]). Another problem is that, although it is easy to embody Hoare’s mismatch in the axiomatic semantics used by a verifier, the details are not often justified rigorously.

Separation Logic [28] (SL) offers elegant means to reason about the *footprint* of an invariant, i.e., the heap locations on which it depends, and the footprint of a command, i.e., the locations possibly written. On this basis, O’Hearn et al [27] give what we call a *second order frame* rule (SOF) that directly embodies Hoare’s mismatch. By contrast with ownership disciplines, the SOF rule pertains directly to hiding on the granularity of a module (i.e., a group of procedures sharing some encapsulated state). We show that this flexibly encompasses invariants such as those that involve cooperating clusters of objects such as a Subject and its Observers. Soundness of the SOF rule has been shown directly in terms of a standard, non-axiomatic program semantics [29]. The rule relies on two critical features of SL. The *separating conjunction* $P * Q$ expresses that formulas P and Q are both true, of disjoint parts of the heap. The *tight interpretation* of a correctness judgement $\{P\}C\{Q\}[\bar{x}]$ says that C is not only correct but neither reads nor writes outside the footprint of P . The Hoare triple is augmented by a *modifies clause* that lists the variables \bar{x} that may be written.

Our contribution is a SOF rule for a classical first-order assertion language that uses an ordinary interpretation of correctness judgements (for fault-avoiding partial correctness). This is achieved by making footprints explicit using *regions* (sets of references), including expressions for the image of a region under a field name and for disjointness of regions. State-dependent region expressions are used in the modifies clause and in a subsidiary judgement about footprints of formulas. Read and write footprints are expressed us-

ing (mutable) ghost state which avoids the need to express footprints using inductively defined predicates that traverse data structures.

Contributions

- We augment module interface specifications by including a *dynamic encapsulation boundary* which must be respected by clients. The dynamic boundary is described via read effects that approximate, in a way suitable to appear in the interface, the footprint of the hidden invariant.
- We extend *region logic* (RL) [3] to include procedures and correctness judgements with hypotheses (Sec. 3 and 4). We improve the treatment of hiding sketched briefly in that paper. On this basis we give our SOF rule (Sec. 5). We prove its soundness in a straightforward operational semantics, which validates standard proof rules, such as Hoare’s rule of Conjunction, and can use a deterministic or nondeterministic memory allocator.
- By contrast with SL, (a) our SOF rule allows the hidden invariant to depend on state that is also read by client programs (e.g., global entities in an application framework); and (b) it allows invariants that are not *precise* [27], i.e., they do not have a unique minimal footprint. Precise predicates typically traverse a data structure (e.g., “every reachable node is balanced”) whereas imprecise ones involve existential quantification (e.g., “some queue is nonempty”). In Sec. 7 we examine why our setting is not susceptible to Reynolds’ conundrum [27] which shows that in SL the SOF rule needs to be restricted to precise invariants or precise specifications to be consistent with the rule of Conjunction.
- We show by examples (Sec. 2 and 6) that our SOF rule can hide invariants that pertain to several objects with a single owner, as well as design patterns in which several peers cooperate (which are incompatible with ownership and remain as challenge problems in the current literature [20, 6, 19]). A program may link together multiple modules, each with its own hidden invariant and dynamic boundary. Our approach encompasses alias confinement disciplines that are enforceable by static analysis [11] as well as less restrictive disciplines that impose proof obligations on clients as in ownership transfers that are “in the eye of the asserter” [27].

Our explicit use of ghost state for frame specifications and separation reasoning was directly inspired by the *dynamic frames* of Kassios [18], whence our term “dynamic boundary”. Similar uses of ghost state, including pure first order encodings of reachability properties, have been found effective in verification tools based on automated theorem provers [16, 10, 35]. One of our aims is to justify the axiomatic semantics used in such tools, which often embody Hoare’s mismatch. We pay a price in verbosity compared with SL, in that ghost fields and variables need to be declared, assigned, and used in specifications. For example, the code of class *Set* in Sec. 2 assigns newly allocated list nodes to a “*rep*” field, and the code in our memory manager example (Sec. 6) updates a region variable holding the objects currently “owned” by the manager.

For readability, the formalization (Sec. 3 and 4) is for barebones programs with rudimentary procedures. Related work is discussed further in Sec. 7.

Our aim is not to advocate a particular proof system (and certainly not a module system) but rather to move beyond competing attempts to provide the “right” specification notation and programming discipline for hiding invariants on mutable state. We provide a logical bridge between internal and external views of a module interface, in which the encapsulation boundary is fully described through dynamic framing (to augment the conventional syntactic signature of the interface). On this basis we hope to shift attention to creating a corpus of specification patterns, with a range of granularity, generality, and balance of proof obligations between modules

```
class Node { val : int; next : Node; own : Object; }
globalvar pool : rgn;
class Set {
  lst : Node; rep : rgn;
  Set()
  { self.lst := null; self.rep := ∅; pool := pool ∪ {self}; }
  add(i : int)
  { if not(contains(i)) then
    var n : Node := new Node();
    n.val := i; n.own := self; n.next := self.lst;
    self.lst := n; self.rep := self.rep ∪ {n}; }
  contains(i : int) : boolean { “linear search for i” }
  remove(i : int) { “remove first node containing i, if any” }
```

Figure 1. Library code for *Set* example. Class *Set* and variable *pool* comprise a module.

Method	Post-condition
<i>Set</i> ()	$elements = empty$ $\wedge pool = old(pool) \cup \{self\}$
<i>add</i> (<i>i</i> : int)	$elements = old(elements) \cup \{i\}$
<i>contains</i> (<i>i</i> : int)	$result = (i \in elements)$
<i>remove</i> (<i>i</i> : int)	$elements = old(elements) - \{i\}$
Method	Effects
<i>Set</i> ()	wr <i>pool</i>
<i>add</i> (<i>i</i> : int)	wr {self}‘any, wr alloc
<i>contains</i> (<i>i</i> : int)	(none)
<i>remove</i> (<i>i</i> : int)	wr {self}‘any, wr <i>pool</i> ‘ <i>rep</i> ‘any

Figure 2. Public specifications for class *Set*. Preconditions: *true*.

```
var s : Set := new Set(); var n : Node := new Node();
s.add(1); s.add(2); n.val := 1; s.remove(1);
b := s.contains(1);
```

Figure 3. Example client, in context of variable *b* : boolean.

and their clients. Modular verification tools will eventually be able to support integrated use of complementary disciplines for a wide range of program design patterns and application frameworks.

2. Synopsis

For a first example of second order framing, we consider the simple program in Fig. 1.¹ This section begins by introducing region logic as used by the interface specifications in Fig 2. Then we consider reasoning about the implementation (Fig. 1) using an invariant, which serves to illustrate local reasoning via a first order “frame rule” akin to Hoare’s rule of Invariance. Finally, we consider reasoning about a client program, using the interface specifications. To justify hiding the invariant, the interface includes a dynamic boundary for the invariant and the client code is obliged not to write within that boundary. We conclude by sketching how soundness of this reasoning is captured by the SOF rule.

Interface specifications —the client’s view. The interface specifications for methods of collection class *Set* do not expose the internal representation but rather refer to the abstraction of interest.

¹The programming notation is similar to Java, in particular a value of a class type like *Node* is either null or a reference to an allocated object with the fields declared in the class. Methods have an implicit parameter, *self*.

The specifications (Fig. 2) are expressed in terms of an integer set, *elements*, that could be formalized as a “model field”.² Abstraction of this sort is commonplace and important, but not the focus of this paper so we don’t formalize *elements*.

What is important is that we shall add to the interface a dynamic boundary which abstracts from the state to be encapsulated for a hidden invariant. To this end, the implementation in Fig. 1 is instrumented using ghost state of type *rgn*; a *region* is a set of allocated references of any type, possibly also containing null. Inclusion of null is one of several small deviations from our paper [3], which remains useful for a more detailed introduction to the basic logic.

The effects clause in Fig. 2 for the constructor, *Set()*, is a conventional modifies specification that says variable *pool* is allowed to be assigned. Our chosen notation and terminology reflects a feature that deals with framing issues beyond the scope of the paper, namely, that both read and write effects may be ascribed to commands and procedures.

For *add*, the write effect $wr \{self\}^{\text{any}}$ is also a conventional one that says fields of *self* may be written (cf. *self.state* in a JML or Spec# modifies clause). The expression $\{self\}$ denotes a *singleton region* containing the value of *self* (an object reference). For any field name *f* and region expression *G*, the effect $wr G^f$ allows update of the *f* fields of objects in *G* (and is read “write *G* image *f*”). The special name “any” abstracts from specific field names to allow any field to be written; this can be refined to “data groups” [24, 8] that abstract from specific field names that should not be exposed to clients.

The primitive region expression “alloc” denotes the current domain of the heap, i.e., the set of allocated references. The effect of *add* includes $wr \text{alloc}$ which means new objects may be allocated.

For *remove*, the most interesting effect is $wr \text{pool}^{\text{rep}}^{\text{any}}$. Here pool^{rep} is a region expression; it denotes the union of the *rep* fields of all objects in *pool*. The effect says fields of objects in pool^{rep} (in the initial state) may be written. For any *s* of type *Set*, field *rep* is intended to hold the (references to) the nodes reached from *s.lst*, and pool^{rep} is the union of those regions, i.e., all the nodes used by *Sets* in *pool*. (To allow for alternate implementations it might be wise for *add* to also have the effect $wr \text{pool}^{\text{rep}}^{\text{any}}$.)

The effect $wr \text{pool}^{\text{rep}}^{\text{any}}$ is a “dynamic frame” because *rep* is a mutable field and *pool* a mutable variable: the meaning of the effect is state dependent.³

A module invariant. For efficiency, our implementation of *remove* relies on the invariant that no integer value is duplicated in the list rooted at *lst*. This invariant must be established by the constructor and preserved by all methods of class *Set*. However, as per what we call *Hoare’s mismatch* [17], the invariant does not appear in the interface specifications as viewed by clients.

The invariant “no duplicates” pertains to a single instance *s* of *Set*, together with its list. The ghost field *s.rep* is intended to refer to the set of nodes reachable from field *s.lst*. To avoid the need for reachability, which is costly for automated verification and not available in all assertion languages, we approximate it

²The term is from JML [8]. Various specification languages offer different ways to express abstractions represented by concrete data. For example, *elements* might be defined as a zeroary function that traverses the list and accumulates the set of its elements.

³Because regions are untyped, *pool* could potentially contain references of many types; nonetheless, the region pool^{rep} involves only those objects in *pool* that have field *rep*, which by uniqueness of field names means objects of type *Set*. Also, we have extended the first version of the logic [3]: if *G* is a region and *f* is a field of some reference type, G^f is a region of the *f*-images, but if *f* is of type *rgn* then G^f is a region, formed as the union of the *f*-images.

using the invariant that *s.rep* contains only nodes “owned” by *s*. Consider the following condition. The first conjunct says there are no duplicates. The next two say that *s.rep* is *nxt*-closed and contains *s.lst*. The last says all nodes in *s.rep* are owned by *s*.

$$\begin{aligned} \text{SetI}(s : \text{Set}) : & (\forall n, m : \text{Node} \in s.\text{rep} \mid n = m \vee n.\text{val} \neq m.\text{val}) \\ & \wedge s.\text{lst} \in s.\text{rep} \wedge s.\text{rep}^{\text{nxt}} \subseteq s.\text{rep} \\ & \wedge s.\text{rep}^{\text{own}} \subseteq \{s\} \end{aligned}$$

Note that “ $n \in s.\text{rep}$ ” is considered false in case $s = \text{null}$.

Were we to extend the example by adding iterators, we might find that the natural granularity for an invariant would be a *Set* together with its iterators as well as its list. To avoid commitment to invariants of a specific granularity, we consider *module invariants*, associated with the program syntax and unit of scope: a “module”, consisting of one or more class declarations. For example, the module invariant could say *SetI* holds for every instance of *Set*. To illustrate flexibility in choosing encapsulation boundaries, we choose instead to say *SetI* holds for only those instances that are in *pool* (which, e.g., could be those returned by a factory method):

$$I : \text{null} \notin \text{pool} \wedge \forall s : \text{Set} \in \text{pool} \mid \text{SetI}(s)$$

We shall find a “frame” or footprint for *I*, that can serve as a dynamic boundary expressing the state-dependent aspect of the encapsulation that will allow *I* to be hidden from clients. But first we seek a frame for the object invariant *SetI(s)*, which will be used for “local reasoning” [28] at the granularity of a single instance of *Set*. The frame is given by read effects that describe the part of the state on which the value of *SetI(s)* may depend:

$$\bar{\delta}_0 : rd\ s, rd\ \{s\}^{\text{rep, lst}}, rd\ s.\text{rep}^{\text{nxt, val, own}}$$

Region logic provides rules for a judgement that says certain read effects bound the footprint of a formula, in this case:

$$\text{true} \vdash \bar{\delta}_0 \text{ frm } \text{SetI}(s) \quad (1)$$

The effect $rd\ s$ says *SetI(s)* may depend on variable *s*, and $rd\ \{s\}^{\text{rep, lst}}$ says it may depend on *s.rep* and *s.lst*. The dynamic part is $rd\ s.\text{rep}^{\text{nxt, val, own}}$ which says it depends on the *nxt*, *val*, and *own* fields of some objects in *s.rep*. The judgement (1) involves a formula, here *true*, because effects can involve region expressions like *s.rep* that depend on mutable state, so framing relationships may hold only under some conditions on that state. For example, we can derive

$$s \in \text{pool} \vdash rd\ \text{pool}^{\text{rep, lst}} \text{ frm } s.\text{lst} \in s.\text{rep} \quad (2)$$

Using judgements including (2) and (1) we can derive a frame judgement $\text{true} \vdash \bar{\delta}_I \text{ frm } I$ for the module invariant, where

$$\bar{\delta}_I : rd\ \text{pool}, \text{pool}^{\text{rep, lst}}, \text{pool}^{\text{rep}}^{\text{nxt, val, own}}$$

An abstraction of $\bar{\delta}_I$ will be used later, when we turn to verification of the client.

Framing for local reasoning using the invariant. For the implementations in Fig. 1, we would like to reason “locally” in terms of a single *Set*. The ownership conditions in *SetI(s)* yield an “island confinement” property (where # denotes disjointness of sets):

$$I \Rightarrow (\forall s, t : \text{Set} \in \text{pool} \mid s = t \vee \{s\}^{\text{rep}} \# \{t\}^{\text{rep}}) \quad (3)$$

because if $s \neq t$, $n \neq \text{null}$, and *n* is in $s.\text{rep} \cap t.\text{rep}$ then $n.\text{own} = s$ and $n.\text{own} = t$, a contradiction. To verify that method *add* preserves *I*, we exploit island confinement in order to focus on the object invariant. Now, *I* is equivalent to $\text{SetI}(\text{self}) \wedge I_{\text{except}}$, where

$$I_{\text{except}} : \text{null} \notin \text{pool} \wedge \forall s \in \text{pool} - \{\text{self}\} \mid \text{SetI}(s)$$

We can frame *Iexcept* by the effects

$$\bar{\delta}_x: \text{ rd self, pool, (pool - \{self\})}^*(\text{rep, lst}), \\ (\text{pool - \{self\}})^* \text{rep}^*(\text{next, val, own})$$

The footprint of I_{except} is disjoint from the footprint of $SetI(\text{self})$. More to the point, let B_{add} be the body of method add . By ordinary means we can verify the following Hoare triple:

$$\{SetI(\text{self})\} B_{add} \{SetI(\text{self}) \wedge \text{elements} = \text{old}(\text{elements}) \cup \{i\}\}$$

Next, we exploit separation to conjoin I_{except} to the pre and post conditions because the write effects of add (Fig. 2) are separate from the read effect of I_{except} . To make this precise, we define an operator: If $\bar{\delta}$ is a set of read effects and $\bar{\varepsilon}$ is a set of write effects then $\bar{\delta} \star \bar{\varepsilon}$ is a conjunction of disjointness formulas, validity of which ensures that writes allowed by $\bar{\varepsilon}$ cannot affect the value of a formula with footprint $\bar{\delta}$. (The formula $\bar{\delta} \star \bar{\varepsilon}$ is defined by induction on the syntax of the effects.) Here is our first order frame rule (from [3]):

$$\text{FRAME} \frac{\vdash \{P\} C \{P'\} [\bar{\varepsilon}] \quad P \vdash \bar{\delta} \text{ frm } Q \quad P \Rightarrow \bar{\delta} \star \bar{\varepsilon}}{\vdash \{P \wedge Q\} C \{P' \wedge Q\} [\bar{\varepsilon}]}$$

It happens that $\bar{\delta}_x \star (\text{wr } \{\text{self}\}^* \text{any, wr alloc})$ is *true*, so we can take Q to be I_{except} in rule FRAME to complete the proof of $\{I\} B_{add} \{I \wedge \text{elements} = \text{old}(\text{elements}) \cup \{i\}\}$.

Reasoning about a client while hiding the invariant. Besides verification of the module’s method bodies with respect to specifications in which the invariant I is explicit, there is another obligation on the module. It must declare a dynamic boundary that frames I . It would not be appropriate to use $\bar{\delta}_I$, which mentions field lst that would be declared private. We choose to use

$$\bar{\theta}_I: \text{ rd pool, pool}^* \text{any, pool}^* \text{rep}^* \text{any}$$

The obligation is $I \vdash \bar{\theta}_I \text{ frm } I$, which is derivable from *true* $\vdash \bar{\theta}_I \text{ frm } I$ by a subsumption rule.

Something is needed to ensure that I is initially true. Typical formalizations include an initializer command, so the client program takes the form *let* m be B in $(\text{init}; C)$. With dynamic allocation, it is constructors that do much of the work to establish invariants. In the present example, let us define $Init$ to be the condition $\text{pool} = \emptyset$ which is suitable to be declared in the module interface. Note that $Init \Rightarrow I$ is valid.

Finally, we can turn to verifying the simple client command in Fig. 3. We will verify that under precondition $Init$ the client establishes postcondition $b = \text{false}$. Here is a proof outline.

```

{Init}
s := new Set();
{s.elements = \emptyset \wedge pool = \{s\}} // by spec of Set
n := new Node();
s.add(1); s.add(2);
{s.elements = \{1, 2\} \wedge pool = \{s\}} // by spec of add
n.val := 1;
s.remove(1);
{s.elements = \{2\} \wedge pool = \{s\}} // by spec of remove
b := s.contains(1);
{b = false} // by spec of contains

```

As it should, this reasoning uses the interface specifications (Fig. 2). However, there is an additional proof obligation. For it to be sound to hide I from the client, we need that the client’s write effects are separate from the dynamic bound $\bar{\theta}_I$ —not just pre/post effects but also effects at intermediate steps, at which module methods like add are called. We include the assignment $n.val := 1$ as a simple example of how encapsulation might be violated (but is not). If this assignment was replaced by $s.lst.val := 1$ then indeed it would break the invariant and render the program incorrect.

The effect of $n.val := 1$ is $\text{wr } \{n\}^* \text{val}$ and it must be shown to be outside the boundary $\bar{\theta}_I$. After all, I reads field val as seen in

the more precise footprint, $\bar{\delta}_I$, subsumed by $\bar{\theta}_I$. By definition of \star , we have that $\bar{\theta}_I \star \text{wr } \{n\}^* \text{val}$ is $\{n\} \# \text{pool} \wedge \{n\} \# \text{pool}^* \text{rep}$. We must show it holds just before the assignment $n.val := 1$.

The condition $\{n\} \# \text{pool}$ is equivalent to $n \notin \text{pool}$ which clearly holds—none of the client code writes pool and the value of n is fresh.⁴ There are several ways to show $\{n\} \# \text{pool}^* \text{rep}$. One way is to notice that a form of “package confinement” [15] applies here: references to the instances of $Node$ used by the Set implementation are never made available to client code. (That is stronger than the property required here, which is merely that such references are not misused by clients.)

To be very explicit about the reasoning that could be embodied by a static analysis for confinement, we shall exploit field own of class $Node$. The analysis might ensure the all-states invariant that any $Node$ pointer p available to the client code has $p.\text{own} = \text{null}$ (the default from $Node$ ’s constructor). Moreover, nodes are not “leaked” by the module code, which can be shown using

$$R: \text{ pool}^* \text{rep}^* \text{own} \subseteq \text{pool} \wedge \text{null} \notin \text{pool}$$

Unlike I , formula R is suitable to appear in the interface. A general fact about region images is that $G \cdot f \subseteq H$ and $x.f \notin H$ imply $x \notin G$ (where G, H are any region expressions and f any field). So using R we get $n \notin \text{pool}^* \text{rep}$ as required by the obligation to respect the dynamic encapsulation boundary.

A dynamic boundary is expressed in terms of state potentially mutated by the module implementation, e.g., the effect of add in Fig. 1 allows writing state on which θ_I depends.⁵ So interface specifications need to provide clients with sufficient information to reason about the boundary. In the example, R could be explicitly conjoined with the interface’s method specifications, or declared as a public invariant [21]. In our memory manager example (Sec 6), it is not a fixed invariant but rather the individual method specifications that allow clients to reason about the boundary.⁶

In summary, the verification of our client and its module are justified by the following rule which embodies Hoare’s mismatch:

$$\frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}] \quad \Delta(\bar{\delta}); (\Theta \circledast I) \langle \rangle \vdash \{Q \wedge I\} B \{Q' \wedge I\} [\bar{\varepsilon}] \quad I \vdash \bar{\theta} \text{ frm } I \quad \text{Init} \Rightarrow I}{\Delta(\bar{\delta}) \vdash \{P \wedge \text{Init}\} \text{let } m \text{ be } B \text{ in } C \{P'\} [\bar{\varepsilon}]}$$

Here for simplicity Θ consists of a single method specification, $\{Q\} m \{Q'\} [\bar{\eta}]$. The first antecedent expresses the proof obligation on the client program C : in addition to the usual pre/post/modifies specification, it must respect the dynamic boundary $\bar{\theta}$ which (please note!) is written together with the specification of m . In general, a boundary is associated with each group of methods taken to comprise the interface of a module, e.g., here we include a second module interface $\Delta(\bar{\delta})$. To respect the boundary means that every primitive action within C must stay outside $\bar{\theta}$ and $\bar{\delta}$. It is not enough that the end-to-end effect $\bar{\varepsilon}$ of C is outside $\bar{\theta}, \bar{\delta}$, as $\bar{\varepsilon}$ only describes effects on objects that exist in the initial state. Moreover, $\bar{\varepsilon}$ includes

⁴ Had we chosen to use typed regions, or to maintain an invariant that every element of pool is of type Set , then $n \notin \text{pool}$ would hold because the type $Node$ of n is not a subtype of Set .

⁵ Dynamic framing raises the issue of interference with effects, which is handled in the rule for sequential composition in [3].

⁶ Nonetheless, we believe the present situation, where $I \Rightarrow R$, is typical. It suggests that the module interface specification may include, in addition to the dynamic bound, a public invariant that is framed by the bound and can be seen as part of the encapsulation discipline. In this case clients are not responsible for maintaining R because (from $I \Rightarrow R$) it is protected by the bound they are already required to respect. Nor is there an additional proof obligation on the module implementation.

$$\begin{aligned}
x, y, r &\in \text{VarName} & f, g &\in \text{FieldName} & K &\in \text{DeclaredClassNames} \\
T &::= \text{int} \mid K \mid \text{rgn} \\
E &::= x \mid c \mid \text{null} \mid E \oplus E & c &\text{ is in } \mathbb{Z} \text{ and } \oplus \text{ is in } \{=, +, -, *, >, \dots\} \\
G &::= x \mid \{E\} \mid \text{emp} \mid \text{alloc} \mid G^{\prime}f \mid G \otimes G & \otimes &\text{ is in } \{\cup, \cap, -\} \\
F &::= E \mid G \\
C &::= x := F \mid x := \text{new } K \mid x := x.f \mid x.f := F \mid \dots
\end{aligned}$$

Figure 5. Programming language (excerpts). Category E is for program expressions and G for region expressions.

effects from calls to methods of Θ and those effects are not required to be outside the boundary $\bar{\theta}$.

The second antecedent is the proof obligation on the implementation, B , of m , in which the invariant I is explicit; the dynamic boundary for Θ is empty, so $\bar{\theta}$ is not imposed on B , but the bound $\bar{\delta}$ on the ambient library must be respected by B . Any recursive calls to m in B are verified using the specifications $\Theta \otimes I$ which means I is conjoined, here $\{Q \wedge I\}m\{Q' \wedge I\}[\bar{\eta}]$. The side conditions $I \vdash \bar{\theta}$ frm I and $\text{Init} \Rightarrow I$ are obligations on the module implementation.

The rule above is a derived rule. Fig. 4 gives a generic derivation, using an ordinary rule that links a procedure to its implementation with no mismatch in specifications, together with the SOF rule and a rule to forget the dynamic boundary once it has served its purpose. The beauty of the SOF, the form of which is due to O’Hearn et al [28], is that it distills the essence of Hoare’s mismatch—which, with proper encapsulation, is an intricate match.

3. Background: Region logic without procedures

Programming language. Fig. 3 gives the programming language syntax, except for procedures. The semantics is standard and deferred to Sec. 4. A program consists of a command C in the context of some class declarations. Commands are standard: assignment, object allocation, field access, field update, conditionals, loops, and local variable blocks.

A class declaration class $K \{ \overline{T} \overline{f} \}$ introduces a type name K . The values of this type are null and all allocated references to mutable objects of type K with fields $\overline{f} : \overline{T}$. Here and throughout, identifiers with an overline range over lists. In addition to int and reference types, there is type rgn with values ranging over finite sets of references of any type.

Program expressions, E , do not depend on the heap: $y.f$ is not an expression but rather part of the primitive field access command, $x := y.f$, for reading a field, as in SL. Region expressions, G , cannot influence control flow or the value of non-region fields/variables (as no integer expressions have subexpressions of type rgn); their purpose is to serve as ghosts for reasoning.

Typing. There is an ambient class table comprising a well formed collection of class declarations. We write $\text{fields}(K)$ for the field declarations $\overline{f} : \overline{T}$ of class K . Judgement $\Gamma \vdash F : T$ says that in context Γ that assigns types to variable names, region or program expression F has type T , and $\Gamma \vdash C$ says C is a well formed command in Γ . Type int is separated from reference types: there is no pointer arithmetic. The typing rule for singleton region $\{E\}$ enforces that E is of reference type. The rule for region dereference, $G^{\prime}f$, checks that f is declared in some class K and either f is a field of class type, K' , or f is of type rgn. The primitive command for field access, $x := y.f$, is restricted to non-region fields. For fields of type rgn, note that, e.g., $x := \{y\}^{\prime}f$ is permitted and is an instance of ordinary assignment, $x := F$. An implicit side condition on all typing and also proof rules is that both the consequent and the antecedents are well formed.

$$\begin{aligned}
P &::= E = E \mid x.f = E \mid G \subseteq G \mid G \# G \mid \text{type}(K, G) \\
&\quad \mid (\forall x : \text{int} \mid P) \mid (\forall x : K \in G \mid P) \mid P \wedge P \mid \neg P \\
\sigma &\models x.f = E && \text{iff } \sigma(x) \neq \text{null} \text{ and } \sigma(x.f) = \llbracket E \rrbracket \sigma \\
\sigma &\models G_1 \# G_2 && \text{iff } \llbracket G_1 \rrbracket \sigma \cap \llbracket G_2 \rrbracket \sigma \subseteq \{\text{null}\} \\
\sigma &\models^{\Gamma} \forall x : K \in G \mid P && \text{iff } \text{extend}(\sigma, x, o) \models^{\Gamma, x:K} P \text{ for all } \\
&&& o \in \llbracket G \rrbracket \sigma \text{ s.t. } o \neq \text{null} \text{ and } \text{type}(o, \sigma) = K
\end{aligned}$$

Figure 6. Formulas: grammar and semantics (excerpts).

Semantics. We assume given a set Ref of reference values including a distinguished value, null . A Γ -state assigns values to the variables in Γ and also has a heap. We abstract from the concrete representation of states and merely assume the following operations are available for a state σ : $\sigma(x)$ is the value of x , $\text{alloc}(\sigma)$ is the set of all allocated references, $\text{type}(o, \sigma)$ is the type of an allocated reference o , $\text{update}(\sigma, o.f, v)$ overrides σ to map field f of o to v (for $o \in \text{alloc}(\sigma)$), $\text{extend}(\sigma, x, v)$ extends σ to map x to value v (for $x \notin \text{dom}(\sigma)$), $\text{new}(\sigma, o, K, \bar{v})$ extends σ to map fresh o (i.e., assuming $o \notin \text{alloc}(\sigma)$) to a K -record with field values \bar{v} and type K . States are assumed to be well typed and to have no dangling references.

The semantics of program expressions E , written $\llbracket E \rrbracket \sigma$, is straightforward and omitted. Note that $\llbracket E \rrbracket \sigma$ is always a value (of appropriate type), never fault ; moreover it only depends on the store, not the heap. For region expressions, the meaning of singleton region $\{E\}$ is $\{\llbracket E \rrbracket \sigma\}$. The meaning of emp is the empty set and that of alloc is $\text{alloc}(\sigma)$, i.e., the set of all allocated references. The meaning of $G^{\prime}f$, if f is a reference type, is the set containing the f -images of all non-null references in G ’s denotation; but if $f : \text{rgn}$ then $G^{\prime}f$ denotes the union of the f -images.

Assertions and effects. The assertion language appears in Fig. 6. Quantification is over int and reference types only (not over rgn). For reference types, a bounding region is required (e.g., region expression G in $\forall x : K \in G \mid P$) and the bound variable must not appear in the bound of the quantification. This facilitates framing. Fig. 6 also gives the semantics of a well formed formula, $\Gamma \vdash P$, as a satisfaction relation $\sigma \models^{\Gamma} P$ that is defined for all Γ -states σ . A formula is *valid* iff it is true in all states.

Effects are given by the grammar

$$\varepsilon ::= \text{rd } x \mid \text{rd } G^{\prime}f \mid \text{wr } x \mid \text{wr } G^{\prime}f \mid \text{rd alloc} \mid \text{wr alloc} \mid \text{fr } G$$

Effect $\text{wr } G^{\prime}f$ says f fields of any pre-existing object in G may be written. Note that, e.g., in $\text{wr } \{x\}^{\prime}f^{\prime}g$ the first apostrophe is forming a region expression $\{x\}^{\prime}f$ and then the second is indicating what field (of objects in the region) is allowed to be written. Effect wr alloc allows allocation and $\text{fr } G$ says that all elements of G in the final state are freshly allocated. Read effect $\text{rd } G^{\prime}f$ says f fields of any pre-existing object in G may be read and rd alloc allows reading the set alloc. An *effect set* is a comma-separated list of effects, ranged over by $\bar{\varepsilon}$ etc. Effects are well formed in Γ if their constituent expressions are.

Freshness effects are needed to annihilate write effects of freshly allocated objects in sequences. Consider, e.g., $x := \text{new } K; x.f := 0$; While the effect of the field update is $\text{wr } \{x\}^{\prime}f$, the pre-state of the entire sequence does not contain the freshly allocated object. Indeed, in this case no pre-existing objects are updated. See the sequence rule in [3].

Definition 1 (allows transition) Let effect set $\bar{\varepsilon}$ be well formed in Γ and let σ, σ' be Γ -states. We say $\bar{\varepsilon}$ *allows transition* from σ to σ' , written $\sigma \rightsquigarrow \sigma' \models \bar{\varepsilon}$, iff σ' extends⁷ σ and the following all hold:

⁷ σ' extends σ provided $\text{alloc}(\sigma) \subseteq \text{alloc}(\sigma')$ and $\text{type}(o, \sigma) = \text{type}(o, \sigma')$ for all $o \in \text{alloc}(\sigma)$.

$$\begin{array}{c}
\frac{\Delta\langle\bar{\delta}\rangle; \Theta\langle\bar{\theta}\rangle \vdash \{P\} C \{P'\} [\bar{\varepsilon}]}{\Delta\langle\bar{\delta}\rangle; (\Theta \otimes I)\langle\bar{\theta}\rangle \vdash \{P \wedge I\} C \{P' \wedge I\} [\bar{\varepsilon}]} \text{SOF} \\
\frac{\Delta\langle\bar{\delta}\rangle; (\Theta \otimes I)\langle\bar{\theta}\rangle \vdash \{Q \wedge I\} B \{Q' \wedge I\} [\bar{\eta}]}{\Delta\langle\bar{\delta}\rangle; (\Theta \otimes I)\langle\bar{\theta}\rangle \vdash \{P \wedge I\} C \{P' \wedge I\} [\bar{\varepsilon}]} \text{DYNBNDELIM} \\
\frac{\Delta\langle\bar{\delta}\rangle \vdash \{P \wedge I\} \text{ let } m \text{ be } B \text{ in } C \{P' \wedge I\} [\bar{\varepsilon}]}{\Delta\langle\bar{\delta}\rangle \vdash \{P \wedge \text{Init}\} \text{ let } m \text{ be } B \text{ in } C \{P'\} [\bar{\varepsilon}]} \text{LINK} \\
\text{CONSEQ}
\end{array}$$

Figure 4. Schematic derivation for second order framing, assuming that Θ is a single specification $\{Q\}m(x:T)\{Q'\}[\bar{\eta}]$. For SOF the side condition is $I \vdash (\bar{\theta}, \text{rd alloc}) \text{ frm } I$ and for CONSEQ it is validity of $\text{Init} \Rightarrow I$.

- (a) for every y in $\text{dom}(\Gamma)$, either $\sigma(y) = \sigma'(y)$ or $\text{wr } y$ is in $\bar{\varepsilon}$
- (b) for every $o \in \text{alloc}(\sigma)$ and every $f \in \text{fields}(o, \sigma)$, either $\sigma(o.f) = \sigma'(o.f)$ or there is $\text{wr } G^*f$ in $\bar{\varepsilon}$ such that $o \in \llbracket G \rrbracket \sigma$
- (c) if $\text{alloc}(\sigma') \neq \text{alloc}(\sigma)$ then wr alloc is in $\bar{\varepsilon}$.
- (d) for each $\text{fr } G$ in $\bar{\varepsilon}$, we have $\llbracket G \rrbracket \sigma' \subseteq \text{alloc}(\sigma') - \text{alloc}(\sigma)$.

Read effects in $\bar{\varepsilon}$ are not significant in Def. 1. As per part (b), expressions in write effects are interpreted in the initial state whereas (see (d)) the region in a freshness effect is interpreted in the final state.

We formalize the separation of read effects from write effects (as in rules FRAME and SOF) via conditions under which two states σ, σ' look the same when viewed through some read effects: define $\text{Agree}(\sigma, \sigma', \bar{\varepsilon})$ to hold just if σ' extends σ and moreover (a) $\sigma(x) = \sigma'(x)$ for all $\text{rd } x$ in $\bar{\varepsilon}$; (b) $\text{alloc}(\sigma) = \text{alloc}(\sigma')$ if rd alloc in $\bar{\varepsilon}$; (c) $\sigma(o.f) = \sigma'(o.f)$ for all $\text{rd } G^*f$ in $\bar{\varepsilon}$ and all $o \in \llbracket G \rrbracket \sigma$ with $f \in \text{fields}(o, \sigma)$.

Subeffects. A natural situation for effect subsumption is when the effect is very detailed or refers to local variable(s) and is thus unsuitable to be exposed for use in interfaces. Subeffecting allows abstraction of such an effect. For example, effect $\text{wr } \text{pool}^* \text{rep}^*$ any of the *remove* method in Sec. 2 is obtained by subeffecting and it abstracts from whatever particular object fields are written.

The judgement for subeffecting takes the form $P \vdash \bar{\varepsilon}_1 \leq \bar{\varepsilon}_2$ and express that the writes/reads in a bigger effect are more permissive in P -states. There are axioms like $G_0 \subseteq G_1 \vdash \text{wr } G_0^*f \leq \text{wr } G_1^*f$ and rules like strengthening the antecedent.

Framing. Recall from Sec. 2 that rule FRAME has a side condition that the write effect of a command be separated from the read footprint of the framed formula. The judgement, $P \vdash \bar{\varepsilon} \text{ frm } Q$ is intended to say the truth or falsity of predicate Q depends only on the state read according to $\bar{\varepsilon}$, i.e., $\bar{\varepsilon}$ covers the footprint of Q in P -states. We first provide a simple syntactic analysis for calculating read effects. For any expression F , define the read footprint of F , written $\text{ftpt}(F)$, as follows: If F is a program expression, E , define $\text{ftpt}(E) = \{\text{rd } x \mid x \in \text{Vars}(E)\}$. For a region expression G , define $\text{ftpt}(G)$ by: $\text{ftpt}(G^*f) = \{\text{rd } G^*f\} \cup \text{ftpt}(G)$; $\text{ftpt}(\text{alloc}) = \{\text{rd alloc}\}$; $\text{ftpt}(\text{emp}) = \emptyset$. Other expressions and primitive formulas are straightforward, e.g., $\text{ftpt}(x.f = E) = \{\text{rd } x, \text{rd } \{x\}^*f\} \cup \text{ftpt}(E)$. This is lifted to judgements by the axiom $\text{true} \vdash \text{ftpt}(P) \text{ frm } P$ for primitive formula P . Proof rules of non-primitive formulas are inductively defined relative to validity of formulas (see [3]), e.g., $\frac{\text{ftpt}(G) \subseteq \bar{\varepsilon} \quad P \wedge x \in G \vdash \bar{\varepsilon}, \text{rd } x, \text{rd } \{x\}^*f \text{ frm } P'}{P \vdash \bar{\varepsilon}, \text{rd } G^*f \text{ frm } \forall x: K \in G \mid P'}$

Lemma 2 (agreement) For any states, σ, σ' , any expression F , any predicates P, P' , and any set of effects $\bar{\varepsilon}$: (a) Suppose $\text{Agree}(\sigma, \sigma', \text{ftpt}(F))$. Then $\llbracket F \rrbracket \sigma = \llbracket F \rrbracket \sigma'$. (b) Suppose $P \vdash \bar{\varepsilon} \text{ frm } P'$ and $\sigma \models P$ and $\text{Agree}(\sigma, \sigma', \bar{\varepsilon})$. Then $\sigma \models P'$ iff $\sigma' \models P'$.

$$\begin{array}{c}
\vdash \{x \neq \text{null} \wedge y = F\} x.f := F \{x.f = y\} [\text{wr } \{x\}^*f] \\
\frac{\text{fields}(K) = \bar{f}: \bar{T} \quad Q' \equiv (\text{type}(\{x\}, K) \wedge x.\bar{f} = \text{default}(\bar{T}))}{\vdash \{\text{true}\} x := \text{new } K \{Q'\} [\text{wr } x, \text{wr alloc}, \text{fr } \{x\}]} \\
\text{SUBEFF} \frac{\vdash \{P\} C \{P'\} [\bar{\varepsilon}] \quad P \vdash \bar{\varepsilon} \leq \bar{\varepsilon}'}{\vdash \{P\} C \{P'\} [\bar{\varepsilon}']} \\
\text{CONJ} \frac{\vdash \{P_1\} C \{P'_1\} [\bar{\varepsilon}] \quad \vdash \{P_2\} C \{P'_2\} [\bar{\varepsilon}]}{\vdash \{P_1 \wedge P_2\} C \{P'_1 \wedge P'_2\} [\bar{\varepsilon}]} \\
\text{SUBST} \frac{\vdash \{P\} C \{P'\} [\bar{\varepsilon}] \quad x \text{ is specification-only} \quad (P/x \rightarrow F) \Rightarrow \text{ftpt}(F) \star (\bar{\varepsilon}/x \rightarrow F)}{\vdash \{P/x \rightarrow F\} C \{P'/x \rightarrow F\} [\bar{\varepsilon}/x \rightarrow F]}
\end{array}$$

Figure 7. Selected correctness rules and axioms for commands.

Separators. Given effect sets $\bar{\delta}$ and $\bar{\varepsilon}$, the separator formula $\bar{\delta} \star \bar{\varepsilon}$ is defined to be the conjunction of certain disjointness formulas. For example, $\text{rd } G_1^*f \star \text{wr } G_2^*g$ is defined to be $G_1 \# G_2$ provided $f \equiv g$; otherwise, it is just *true*.⁸ In a state where $\bar{\delta} \star \bar{\varepsilon}$ holds, nothing that the read effects in $\bar{\delta}$ allow to be read can be written according to the write effects $\bar{\varepsilon}$. The following result, together with Lemma 2, proves soundness of FRAME.

Lemma 3 (separator agreement) Consider any effect sets $\bar{\delta}$ and $\bar{\varepsilon}$. Suppose $\sigma \rightsquigarrow \sigma' \models \bar{\varepsilon}$ and $\sigma \models \bar{\delta} \star \bar{\varepsilon}$. Then $\text{Agree}(\sigma, \sigma', \bar{\delta})$.

Proof rules and correctness. A correctness judgement takes the form $\vdash^\Gamma \{P\} C \{P'\} [\bar{\varepsilon}]$ and is *well formed* in Γ just if $P, P', \bar{\varepsilon}$ are well formed in Γ and also $\Gamma \vdash C$. *Validity* of the judgement consists of a standard part—from any initial state that satisfies P , C does not fault (terminate with error), and if it terminates then the final state satisfies P' —and a non-standard one: any allocation and update effects are allowed by $\bar{\varepsilon}$ (Def. 1).

Selected proof rules appear in Fig. 7. For field update we choose a “small axiom”, inspired by [28], which snapshots F with an auxiliary variable y in the precondition. The effect in the rule is a write to the f field of the single object in region $\{x\}$ that exists in the pre-state. The other rules shown are structural rules. The rule of substitution refers to *specification-only* variables that are not allowed to appear in code.⁹

4. Region logic with procedures

In this section, procedures are added to the programming language. Correctness judgements are extended with hypotheses, i.e., proce-

⁸ We do not need to recurse further with G_1 and G_2 . The ftpt function at the core of the frames judgement generates convex sets; if the set contains $\text{rd } G^*f$ then it also contains $\text{ftpt}(G)$.

⁹ Note that [3] conjectures mistakenly that read effects of commands can be used for the substitution rule; here we give a sound rule.

cedure specifications. The program semantics is given and used to define the semantics of correctness judgements with hypotheses.

Program syntax and semantics. Command syntax has these additional forms

$$C ::= m(z) \mid \text{let } m(x : T) \text{ be } C \text{ in } C \mid \text{skip} \mid \text{end}(x) \mid \text{end}(m)$$

where m ranges over procedure names. The form “let $m(x : T)$ be B in C ” binds m to B in C . Typing rules enforce that procedure body B is typed in a context with its parameter x and also m in scope — recursion is allowed. To streamline notation we avoid generalizing to multiple parameters and multiple, mutually recursive procedures, but there is no difficulty in that. The command end is a technical device, not allowed in source programs and used only in the semantics to mark the end of the scope of a variable or method. The *primitive commands* are method calls, skip, and assignments including field update and new.

Let Γ range over contexts like $y : T, m : (x : U)$ that, in addition to variable declarations, declare procedure parameters. Then the typing rule for linking to procedures is standard.

$$\frac{\Gamma, x : T, m : (x : T) \vdash B \quad \Gamma, m : (x : T) \vdash C}{\Gamma \vdash \text{let } m(x : T) \text{ be } B \text{ in } C}$$

We refrain from giving other typing rules as they are quite standard.

We use a small-step operational semantics in which configurations carry a *procedure environment* that binds procedure names to their bodies. Configurations take the form $\langle C, \sigma, \mu \rangle$ where C is a command, σ is a state (as defined in Sec. 3), and procedure environment μ is a partial function from procedure names to parameterized commands of the form $(\lambda x : T. C)$. We consider only well formed programs and initial configurations, so reachable configurations enjoy standard well-formedness properties that will not be discussed in detail. Every reachable configuration has a successor unless the command is skip. A terminating computation ends in a configuration of the form $\langle \text{skip}, \sigma, \mu \rangle$ or else *fault*.

The only unusual feature of the semantics is that it depends on assumed specifications of procedures that are not bound in the procedure environment. The identifier Δ will range over such assumptions. The semantics is given as a transition relation \mapsto_{Δ} for commands that may have free occurrences of procedures specified in Δ . The procedures in Δ are to be distinct from those in the procedure environment. Semantics of a call $m(z)$, for m in Δ , depends on the specification; its definition is deferred (to Fig. 9) and discussed later.

The transition semantics \mapsto_{Δ} for all other cases is independent from Δ . It is defined in Fig. 8. For allocation, we assume *fresh* is an arbitrary function such that $\text{fresh}(\sigma)$ is a non-empty set of non-null references that are not allocated in σ . Thus our results encompass deterministic allocators as well as the maximally non-deterministic one used in separation logic. We write \mapsto_{Δ}^* for the reflexive, transitive closure of \mapsto_{Δ} .

In addition to the well-formedness properties of reachable configurations that were already mentioned, we note that the *context procedures* specified in Δ are always distinct from the *linked procedures* bound in the procedure environment. Furthermore, if C is not an end then $\langle C, \sigma, \mu \rangle \mapsto_{\Delta}^* \langle \text{skip}, \sigma', \mu' \rangle$ implies $\mu' = \mu$. There is no deallocation so the domain of the heap only grows, and once allocated the type of a reference never changes.

Procedure specifications and hypothetical judgements. A hypothesis Δ is a comma-separated list of procedure specifications, each of the form

$$\{Q\}m(x : T)\{Q'\}[\bar{\varepsilon}] \quad (4)$$

$$\begin{array}{c} \frac{\sigma(x) = \text{null}}{\langle x.f := F, \sigma, \mu \rangle \mapsto \text{fault}} \\ \frac{\sigma(x) = o \quad o \neq \text{null}}{\langle x.f := F, \sigma, \mu \rangle \mapsto \langle \text{skip}, \text{update}(\sigma, o.f, \llbracket F \rrbracket \sigma), \mu \rangle} \\ \frac{o \in \text{fresh}(\sigma) \quad \text{fields}(K) = \bar{f} : \bar{T} \quad \text{default}(\bar{T}) = \bar{v}}{\langle x := \text{new } K, \sigma, \mu \rangle \mapsto \langle \text{skip}, \text{new}(\sigma, o, K, \bar{v}), \mu \rangle} \\ \langle (\text{skip}; C), \sigma, \mu \rangle \mapsto \langle C, \sigma, \mu \rangle \\ \frac{\langle C_0, \sigma, \mu \rangle \mapsto \langle C'_0, \sigma', \mu' \rangle}{\langle (C_0; C_1), \sigma, \mu \rangle \mapsto \langle (C'_0; C_1), \sigma', \mu' \rangle} \\ \frac{\langle \text{let } m(x : T) \text{ be } B \text{ in } C, \sigma, \mu \rangle \mapsto \langle (C; \text{end}(m)), \sigma, \text{extend}(\mu, m, (\lambda x : T. B)) \rangle}{\mu(m) = \lambda x : T. B \quad x' \notin \text{dom}(\sigma) \quad B' = B/x \rightarrow x'} \\ \langle m(z), \sigma, \mu \rangle \mapsto \langle (B'; \text{end}(x')), \text{extend}(\sigma, x', \sigma(z)), \mu \rangle \\ \langle \text{end}(x), \sigma, \mu \rangle \mapsto \langle \text{skip}, \text{retract}(\sigma, x), \mu \rangle \\ \langle \text{end}(m), \sigma, \mu \rangle \mapsto \langle \text{skip}, \sigma, \text{retract}(\mu, m) \rangle \end{array}$$

Figure 8. Definition of transition relation \mapsto_{Δ} . For calls of procedures in Δ , for which see Fig. 9. Here Δ is the same throughout and omitted. Also $\text{retract}(\sigma, x)$ removes x from the state. Rules for ordinary assignment, if, and while are standard.

where m is a procedure name and x is the parameter name. For (4) to be well formed, $Q, Q', \bar{\varepsilon}$ should be wf in $\Gamma, x : T$. Less obviously, $\bar{\varepsilon}$ must not contain $\text{wr } x$; this enforces the usual constraint on value-parameters: so that use of x in postconditions or in effects like $\text{wr } x.f$ refers to its initial value).

Correctness judgements now take the form

$$\Delta \vdash^{\Gamma} \{P\} C \{P'\} [\bar{\varepsilon}] \quad (5)$$

Judgement (5) is well formed just if $P, P', \bar{\varepsilon}$, and all specifications in Δ are well formed in Γ . Moreover C is well formed with respect to the procedure signatures in Δ , i.e., $\Gamma; \text{signs}(\Delta) \vdash C$ where *signs* extracts the procedure signatures from a set of procedure specifications.

There are two plausible interpretations of a hypothetical correctness judgement (5), both expressing that C is “modularly correct” with respect to the specifications of procedures of its context Δ . The first interpretation can be described as follows: Consider the semantics of C started from a procedure environment that links each procedure in Δ to an arbitrary command that satisfies the procedure’s specification. That is, C is correct with respect to all correct implementations of Δ . The second interpretation avoids this universal quantification by interpreting C with respect to a single transition relation (not necessarily denotable by a command) that models the least refined or “worst” meaning that satisfies the specification. This embodies all possible behaviors of correct implementations. The second interpretation is popular in work on program refinement. It is used as well by O’Hearn et al [27, 29]. Although typically the second interpretation is associated with a big step semantics, we use it in mixed step form (as in, e.g., [31]).

Fig. 9 completes the definition of transition relation \mapsto_{Δ} , by giving the step from configuration $\langle m(z), \sigma, \mu \rangle$ where m is a context procedure rather than being in the environment μ . To explain the definition, let us first consider a specification $\{P\}m(x : T)\{P'\}[\bar{\varepsilon}]$. There are two cases. If σ satisfies the pre-

$$\frac{\Delta \text{ contains } \{P\}m(x:T)\{P'\}[\bar{\varepsilon}] \quad \sigma \rightsquigarrow \sigma' \models \bar{\varepsilon} \quad \text{extend}(\sigma, x, \sigma(z)) \models P \quad \text{extend}(\sigma', x, \sigma(z)) \models P'}{\langle m(z), \sigma, \mu \rangle \xrightarrow{\Delta} \langle \text{skip}, \sigma', \mu \rangle}$$

$$\frac{\Delta \text{ contains } \{P\}m(x:T)\{P'\}[\bar{\varepsilon}] \quad \text{extend}(\sigma, x, \sigma(z)) \not\models P}{\langle m(z), \sigma, \mu \rangle \xrightarrow{\Delta} \langle \text{skip}, \sigma', \mu \rangle \quad \text{and also} \quad \langle m(z), \sigma, \mu \rangle \xrightarrow{\Delta} \text{fault}}$$

Figure 9. Definition of $\xrightarrow{\Delta}$ for calls of *context procedures*, i.e., those in Δ .

$$\text{CALL} \frac{\Delta \text{ contains } \{P\}m(x:T)\{P'\}[\bar{\varepsilon}] \quad Q \equiv P/x \rightarrow z \quad Q' \equiv P'/x \rightarrow z \quad \bar{\varepsilon}' \equiv \bar{\varepsilon}/x \rightarrow z}{\Delta \vdash \{Q\}m(z)\{Q'\}[\bar{\varepsilon}']}$$

LINK

$$\frac{\Theta \text{ is } \{Q\}m(x:T)\{Q'\}[\bar{\delta}] \quad \Delta, \Theta \vdash^{\Gamma} \{P\}C\{P'\}[\bar{\varepsilon}] \quad \Delta, \Theta \vdash^{\Gamma, x:T} \{Q\}B\{Q'\}[\bar{\delta}]}{\Delta \vdash^{\Gamma} \{P\} \text{let } m(x:T) \text{ be } B \text{ in } C\{P'\}[\bar{\varepsilon}]}$$

Figure 10. Proof Rules for hypothetical judgements.

condition, then the configuration can step to $\langle \text{skip}, \sigma', \mu \rangle$ for any σ' such that σ' satisfies P' and moreover $\bar{\varepsilon}$ allows the transition from σ to σ' . If σ does not satisfy the precondition, then the configuration can step to any σ' whatsoever, and may also step to *fault*.

A correctness judgement is *valid*, written $\Delta \models^{\Gamma} \{P\}C\{P'\}[\bar{\varepsilon}]$, iff the following holds. Let μ_0 be an arbitrary procedure environment disjoint from the procedures bound in C or present in Δ . Let $C_0 = C$. Then for all Γ -states σ_0 such that $\sigma_0 \models P$ we require (i) It is not the case that $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta}^* \text{fault}$; and (ii) For all computations $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta}^* \langle \text{skip}, \sigma_n, \mu_n \rangle$ of n steps (noting that $n \geq 0$ and $\mu_n = \mu_0$) we have both $\sigma_n \models P'$ and $\sigma_0 \rightsquigarrow \sigma_n \models \bar{\varepsilon}$. The new proof rules are given in Fig. 10. The proof rules of Sect. 3 are all adapted by adding hypotheses Δ to all correctness judgements.

Extension to a richer language. In examples we use an extended programming language, where methods have an implicit parameter named *self* and may have a return value, and there are constructor methods. Formalization of the additional features complicates the rules ALLOC, LINK, and CALL. Calls without return value take the form $y.m(z)$. The specification may mention both *self* and parameter x , so y is substituted for *self* (and z for x as before). With a return value, there is an additional substitution.

The LINK rule no longer has an exact match between the specification used by the client and the specification with respect to which a method implementation is verified. That is because a method call involves first testing nullity of the receiver and only then dispatching the body of the method; so precondition $\text{self} \neq \text{null}$ is added. For the constructor of some class K , the proof obligation also reflects in its precondition that the fields are initialized to defaults before the constructor runs; the effect specification reflects that *self* is freshly allocated and thus writes to fields of *self* must not be reported.

5. Module invariants and second order framing

This section augments hypotheses with dynamic boundaries. A dynamic boundary $\bar{\delta}$ is associated with a group (list) Δ of procedure specifications using notation $\Delta(\bar{\delta})$. The general form for correctness judgement would have a sequence $\Delta_0(\bar{\delta}_0); \dots; \Delta_n(\bar{\delta}_n)$ of

hypotheses, and rules for permuting the groups. But for readability it suffices to spell out the case of two groups. So a correctness judgement has the form

$$\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash^{\Gamma} \{P\}C\{P'\}[\bar{\varepsilon}] \quad (6)$$

The judgement is wf if the conditions for (5) hold and the read effects $\bar{\delta}$ and $\bar{\theta}$ are wf in Γ .¹⁰

The current command in a configuration can always be written as a sequence of one or more commands that are not themselves sequences; the first is the *active command*, the one that is rewritten in the next step. Formally we define $\text{active}(C_1; C_2) = \text{active}(C_1)$ and $\text{active}(C) = C$ if there are no C_1, C_2 such that C is $C_1; C_2$.

Definition 4 A correctness judgement is *valid*, written

$$\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models^{\Gamma} \{P\}C\{P'\}[\bar{\varepsilon}]$$

iff the following holds. Let Δ' be the union (Δ, Θ) , let C_0 be C , and let μ_0 be an arbitrary procedure environment disjoint from the procedures linked within C or present in Δ, Θ . Then for all Γ -states σ_0 such that $\sigma_0 \models P$

- It is not the case that $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta'}^* \text{fault}$.
- For all computations $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta'}^* \langle \text{skip}, \sigma_n, \mu_n \rangle$
 - (a) $\sigma_n \models P'$ and $\sigma_0 \rightsquigarrow \sigma_n \models \bar{\varepsilon}$
 - (b) for all $0 < i \leq n$, either $\text{active}(C_{i-1})$ is a call to some procedure m in Δ or else $\text{Agree}(\sigma_{i-1}, \sigma_i, \bar{\delta})$
 - (c) like (b) but *mutatis mutandis* for Θ and $\bar{\theta}$.

Note that in the case $\bar{\delta}$ and $\bar{\theta}$ are empty, this coincides with the definition of validity in Sec. 4.

The axioms for assignment forms are now given hypotheses $\Delta(\bar{\delta}); \Theta(\bar{\theta})$ with empty dynamic bounds. Proper rules are revised to have $\Delta(\bar{\delta}); \Theta(\bar{\theta})$ in both antecedents and conclusions, except that CALL and LINK are replaced by the ones in Fig. 11. Note that LINK' requires the procedure body to respect the other dynamic bounds and linked procedures must have empty dynamic bound. In addition, the Figure adds rules used to introduce and eliminate the dynamic boundary.

Rule SOF uses an operation $\otimes I$ that conjoins a formula I to pre- and post-conditions of specifications. It is defined by

$$(\{Q\}m(x:T)\{Q'\}[\bar{\varepsilon}]) \otimes I = \{Q \wedge I\}m(x:T)\{Q' \wedge I\}[\bar{\varepsilon}]$$

Define $\Delta \otimes I$ by distributing $\otimes I$ over the specifications in Δ . Rule SOF also imposes a mild admissibility condition on I . The problem is that some useful invariants include *alloc* in their effect, e.g., in the example of Sec. 2 we drop variable *pool* and instead let I quantify over all allocated *Sets*. Typical clients do allocation, and thus write *alloc*, which would conflict with a dynamic boundary containing *rd alloc*. The solutions are based on the limited way that *alloc* gets written: only by the new command and only by adding the newly allocated object.

Our first solution to the above problem is similar to one used in some specific invariant disciplines [26, 32].

Definition 5 Formula Q is *admissible* with respect to the procedure specifications Δ , written $\text{admiss}(Q, \Delta)$, iff it is not falsifiable by allocation. That is, for any σ with $\sigma \models Q$, we also have $\sigma' \models Q$ where σ' is $\text{new}(\sigma, o, K, \text{default}(\bar{T}))$ and \bar{T} is the field types for K and o is any fresh reference (not in $\text{alloc}(\sigma)$).

¹⁰ In a proof, the dynamic bounds will appear in many judgements, including those for procedure implementations and in the scope of local variables. Thus to be useful they must be well formed in a number of contexts — typically, they involve only some global variables and public field names.

$$\begin{array}{c}
\text{SOF} \frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}]}{I \vdash (\bar{\theta}, \text{rd alloc}) \text{ frm } I \quad \text{admiss}(I, \Theta)} \\
\Delta(\bar{\delta}); (\Theta \circledast I)(\bar{\theta}) \vdash \{P \wedge I\} C \{P' \wedge I\} [\bar{\varepsilon}] \\
\\
\text{DYNBND\text{ELIM}} \\
\frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}]}{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}]} \\
\\
\text{DYNBND\text{INTRO}} \\
\frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}] \quad C \text{ is primitive} \quad P \Rightarrow \bar{\theta} \star \bar{\varepsilon}}{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} C \{P'\} [\bar{\varepsilon}]} \\
\\
\text{DYNBND\text{INTROM}} \\
\frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} m(z) \{P'\} [\bar{\varepsilon}] \quad m \text{ is in } \Theta}{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{P\} m(z) \{P'\} [\bar{\varepsilon}]} \\
\\
\text{CALL}' \\
\frac{\Delta \text{ or } \Theta \text{ contains } \{P\} m(x:T) \{P'\} [\bar{\varepsilon}]}{Q \equiv P/x \rightarrow z \quad Q' \equiv P'/x \rightarrow z \quad \bar{\varepsilon}' \equiv \bar{\varepsilon}/x \rightarrow z} \\
\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \{Q\} m(z) \{Q'\} [\bar{\varepsilon}'] \\
\\
\text{LINK}' \\
\frac{\Theta \text{ is } \{Q\} m(x:T) \{Q'\} [\bar{\eta}] \quad \Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \Gamma \{P\} C \{P'\} [\bar{\varepsilon}]}{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \Gamma, x:T \{Q\} B \{Q'\} [\bar{\eta}]} \\
\Delta(\bar{\delta}) \vdash \Gamma \{P\} \text{ let } m(x:T) \text{ be } B \text{ in } C \{P'\} [\bar{\varepsilon}]
\end{array}$$

Figure 11. Rules involving dynamic boundaries. Omitted is a rule that swaps $\Delta(\bar{\delta}); \Theta(\bar{\theta})$ with $\Theta(\bar{\theta}); \Delta(\bar{\delta})$, unconditionally.

This is formulated as a semantic property of Q , unlike all the other side conditions in the proof rules. We include Δ in the notation, even though it is not used in the definition, in order to highlight the way admissibility should be treated in a richer language that includes constructor methods. Before discussing that further, we note the following which follows immediately using Lemma 2.

Lemma 6 Suppose $P \vdash (\bar{\delta}, \text{rd alloc}) \text{ frm } Q$ and $\text{admiss}(Q, \Delta)$. Suppose $\sigma \models P \wedge Q$ and let $\sigma' = \text{new}(\sigma, o, K, \bar{v})$. If $\text{Agree}(\sigma, \sigma', \bar{\delta})$ then $\sigma' \models Q$.

The point is that by condition $P \vdash (\bar{\delta}, \text{rd alloc}) \text{ frm } Q$, it appears that Q depends on alloc, but by semantic condition $\text{admiss}(Q, \Delta)$ it does not.

Suppose Q is of the form $\forall x: K \in \text{alloc} \mid x.\text{init} \Rightarrow P(x)$ with init a boolean field. Although the footprint of Q includes alloc, it is not falsifiable by allocation (because the default value for $x.\text{init}$ is false). Such a formula would be suitable as an invariant in a program where $x.\text{init}$ only gets truthified by procedures that also establish $P(x)$. In a richer language with constructor methods, this pattern is in some sense part of the semantics and we can define $\text{admiss}(Q, \Delta)$ by the following syntactic condition: every occurrence of alloc in Q occurs in a subformula of the form $\forall x: K \in \text{alloc} \mid \dots$ where the constructor for class K is among the procedures of Δ .

A verification condition generator could impose the conditions of rule DYNBNDINTRO on each primitive in a program. However, for general purpose encapsulation disciplines based on type systems or other static analyses, a dynamic bound $\bar{\eta}$ could be introduced by a rule of the form

$$\frac{\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \Gamma \{P\} C \{P'\} [\bar{\varepsilon}] \quad \text{Ok}(\bar{\eta}, C, \Gamma, \Delta(\bar{\delta}), \Theta(\bar{\theta}), P)}{\Delta(\bar{\delta}); \Theta(\bar{\theta}, \bar{\eta}) \vdash \Gamma \{P\} C \{P'\} [\bar{\varepsilon}]}$$

It posits a condition $\text{Ok}(\bar{\eta}, C, \Gamma, \Delta(\bar{\delta}), \Theta(\bar{\theta}), P)$ which performs some static analysis of C , taking into account that methods of Θ

are exempt from the effect bound $\bar{\theta}$. What is required of Ok is that it makes this rule sound, i.e., it ensures that every step of C , other than calls to methods in Θ , respects the bound $\bar{\eta}$.

We leave this topic and turn to the main result of the paper.

Theorem 7 Any judgement $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \vdash \Gamma \{P\} C \{P'\} [\bar{\varepsilon}]$ that is derivable is valid, i.e., $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models \Gamma \{P\} C \{P'\} [\bar{\varepsilon}]$.

Proof: By induction on the derivation and by cases on the last rule used. Each case simply appeals to Lemma 9. \square

Proposition 8 For all $C, C', \sigma, \sigma', \mu, \mu', \Delta, \Delta'$, such that Δ and Δ' declare the same methods and $\text{active}(C)$ is not a call to a method in Δ , we have: $\langle C, \sigma, \mu \rangle \xrightarrow{\Delta} \langle C', \sigma', \mu' \rangle$ if and only if $\langle C, \sigma, \mu \rangle \xrightarrow{\Delta'} \langle C', \sigma', \mu' \rangle$.

Proof of the proposition is by induction on the structure of C ; under the given conditions, only the transition rules in Fig. 8 are at play.

Lemma 9 (rule soundness) Every axiom is valid. For every proper rule, derives valid conclusions from valid antecedents.

Proof: By cases on the axioms and rules. We give the main case, SOF. Assume the antecedent is valid, $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models \Gamma \{P\} C \{P'\} [\bar{\varepsilon}]$. To prove validity of the conclusion, i.e.,

$$\Delta(\bar{\delta}); (\Theta \circledast I)(\bar{\theta}) \models \{P \wedge I\} C \{P' \wedge I\} [\bar{\varepsilon}] \quad (7)$$

let Δ_c be the hypotheses $(\Delta, \Theta \circledast I)$, for the conclusion of the rule, and let Δ_a be the hypotheses (Δ, Θ) for the antecedent. Consider any σ_0 with $\sigma_0 \models P \wedge I$, let μ_0 be any method environment disjoint from Δ, Θ , and let C_0 be C .

Claim A: Consider any computation from $\langle C_0, \sigma_0, \mu_0 \rangle$ under transition relation $\xrightarrow{\Delta_c}$. Then (a) that sequence of configurations is also a computation of $\xrightarrow{\Delta_a}$ and (b) if $\sigma_0 \models I$ then I holds in every configuration. Claim A implies (7) as follows.

It is not the case that $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta_c} \text{*fault}$, because that would imply $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta_a} \text{*fault}$ which contradicts validity of the antecedent. Furthermore, by validity of the antecedent we get $\sigma_n \models P'$ and $\sigma_0 \rightsquigarrow \sigma_n \models \bar{\varepsilon}$. We get $\sigma_n \models I$ from the claim. What remains is to show that for all $0 < i \leq n$, either $\text{active}(C_{i-1})$ is a call to some method m in $(\Theta \circledast I)$ or else $\text{Agree}(\sigma_{i-1}, \sigma_i, \bar{\theta})$, and *mutatis mutandis* for Δ and $\bar{\delta}$. This follows immediately from the corresponding condition in validity of the antecedent.

It remains to prove Claim A, which is by induction on the length of the computation sequence for $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta_c} \langle C_n, \sigma_n, \mu_n \rangle$. The induction step uses **Claim B:** For any $D, D', \sigma, \sigma', \mu, \mu'$, suppose that $\langle D, \sigma, \mu \rangle$ is reachable under $\xrightarrow{\Delta_a}$ from an initial configuration $\langle C_0, \mu_0, \sigma_0 \rangle$. Suppose $\sigma \models I$ and $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_c} \langle D', \sigma', \mu' \rangle$. Then $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_a} \langle D', \sigma', \mu' \rangle$ and $\sigma' \models I$. To prove Claim B there are 3 cases: **(a)** If $\text{active}(D)$ is not a call to a context method (noting that Δ_a and Δ_c declare the same methods) then by Proposition 8, $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_a} \langle D', \sigma', \mu' \rangle$. If $\sigma' \neq \text{new}(\sigma, o, K, \bar{v})$ then $\text{alloc}(\sigma) = \text{alloc}(\sigma')$ and from $\text{Agree}(\sigma, \sigma', \bar{\theta})$ (by valid antecedent) and $\sigma \models I$, we get $\sigma' \models I$ by Lemma 2. In case $\sigma' = \text{new}(\sigma, o, K, \bar{v})$, from $\sigma \models I$ and side conditions $I \vdash \text{rd alloc}, \bar{\theta} \text{ frm } I$ and $\text{admiss}(I, \Theta)$ of SOF, we get $\sigma' \models I$ by Lemma 6.

For the second and third cases, suppose that $\text{active}(D)$ is a call to a method m with specification $\{V\}m(x:T)\{V'\}[\bar{\eta}]$.

Case **(b)** $m \in \Delta$: It must be that $\sigma \in \llbracket V \rrbracket$. For, if not, then by context call semantics (Fig. 9) we would have $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_a} \text{fault}$

which, because $\langle D, \sigma, \mu \rangle$ is reachable, contradicts the antecedent $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models^{\Gamma} \{P\} C \{P'\} [\bar{\varepsilon}]$. So by the first rule in Fig. 9 for both Δ_a and Δ_c we get $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_a} \langle D', \sigma', \mu' \rangle$. We show that I is maintained, by cases on whether $\sigma' = \text{new}(\sigma, o, K, \bar{v})$. If so, from $\sigma \models I$ and side conditions $I \vdash \text{rd alloc}, \bar{\theta} \text{ frm } I$ and $\text{admiss}(I, \Theta)$ we get $\sigma' \models I$ by Lemma 6. Otherwise, from $\text{Agree}(\sigma, \sigma', \bar{\theta})$ and $\sigma \models I$, we have $\sigma' \models I$ by Lemma 2.

Case (c) $m \in \Theta \otimes I$. Then V is $Q \otimes I$ and V' is $Q' \otimes I$ for some Q, Q' such that $\{Q\}m(x: T)\{Q'\}[\bar{\eta}]$ is in Θ . It must be that $\sigma \models Q$: for, if not, then by method call semantics (Fig. 9) we would have $\langle D, \sigma, \mu \rangle \xrightarrow{\Delta_a} \text{fault}$, contradicting $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models^{\Gamma} \{P\} C \{P'\} [\bar{\varepsilon}]$. Because $\sigma \models Q$, by specification of m in $\Theta \otimes I$ we get $\sigma' \models Q' \wedge I$.¹¹ \square

Case DynBndIntro. For DYNBNDINTRO, suppose C is a primitive command and the side condition $P \Rightarrow \bar{\theta} \star \bar{\varepsilon}$ holds. Suppose the antecedent is valid: $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models \{P\} C \{P'\} [\bar{\varepsilon}]$. To show validity of the conclusion, $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models \{P\} C \{P'\} [\bar{\varepsilon}]$, consider any μ and any σ with $\sigma \models P$. If C is skip then there is no transition and the only thing to prove is $\sigma \models P'$ — which we have by validity of the antecedent. If C is an assignment, field update, or call of a method, the possible transitions have the form $\langle C, \sigma, \mu \rangle \mapsto \langle \text{skip}, \sigma', \mu' \rangle$ (by the antecedent we do not have to consider faults) and by the antecedent we have $\sigma' \models P'$. To show $\text{Agree}(\sigma, \sigma', \bar{\theta})$, we can use Lemma 3 owing to side condition $P \Rightarrow \bar{\theta} \star \bar{\varepsilon}$. The fact that $\bar{\delta}$ is respected follows from the antecedent.

Case Link. Suppose Θ is $\{Q\}m(x: T)\{Q'\}[\bar{\eta}]$. Suppose both antecedents of the rule are valid, i.e., $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models^{\Gamma} \{P\} C \{P'\} [\bar{\varepsilon}]$ and $\Delta(\bar{\delta}); \Theta(\bar{\theta}) \models^{\Gamma, x: T} \{Q\} B \{Q'\} [\bar{\eta}]$. To show

$$\Delta(\bar{\delta}) \models^{\Gamma} \{P\} \text{ let } m(x: T) \text{ be } B \text{ in } C \{P'\} [\bar{\varepsilon}]$$

suppose $\sigma_0 \models P$ and μ_0 be any environment. Let C_0 be let $m(x: T)$ be B in C . We must show that (a) it is not the case that $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta} \text{fault}$; and (b) if $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta} \langle \text{skip}, \sigma_n, \mu_n \rangle$ then $\sigma_n \models P'$ and $\sigma_0 \rightsquigarrow \sigma_n \models \bar{\varepsilon}$ and $\text{Agree}(\sigma_{i-1}, \sigma_i, \bar{\delta})$ for all steps $\sigma_i \mapsto \sigma_{i+1}$. Validity of the antecedents tells us about $\xrightarrow{\Delta'}$, where Δ' is Δ, Θ . We must therefore relate $\xrightarrow{\Delta'}$ to $\xrightarrow{\Delta}$. By semantics, the first step is $\langle C_0, \sigma_0, \mu_0 \rangle \xrightarrow{\Delta} \langle C_1, \sigma_1, \mu_1 \rangle$ where C_1 is C , $\sigma_1 = \sigma_0$, and μ_1 extends μ_0 to map m to $\lambda x: T. B$. It respects $\bar{\delta}$ because the state is unchanged. Subsequent steps are matched exactly by steps via $\xrightarrow{\Delta'}$ —which respect $\bar{\delta}$ by validity of the antecedent for C — until we reach a configuration $\langle C_i, \sigma_i, \mu_i \rangle$ where the active command in C_i is a call of m , i.e., C_i is a sequence $m(z); D$, for some D . In this case, $\xrightarrow{\Delta'}$ does not fault (as otherwise the antecedent for C would not be valid).

Thus by definition of $\xrightarrow{\Delta'}$ it must be that $\sigma \models Q$ (as otherwise it could fault). Instead, it goes in one step to $\langle \text{skip}; D, \sigma', \mu_i \rangle$ for all σ' such that $\sigma' \models Q'$ and $\sigma_i \rightsquigarrow \sigma' \models \bar{\eta}$. The next step goes to $\langle D, \sigma', \mu_i \rangle$, which we will show is matched.

As for $\xrightarrow{\Delta}$, it steps first to $\langle (B'; \text{end}(x'); D), \sigma'', \mu_i \rangle$ where $\mu_i(m) = \lambda x: T. B$ and B' has x renamed to fresh identifier x' that is initialized in σ'' to $\sigma_i(z)$. From there, by validity of the antecedent for B and using $\sigma \models P$, there is no fault and if the overall computation terminates then it first reaches the end of the

B' -computation, i.e., a configuration $\langle (\text{skip}; \text{end}(x'); D), \sigma', \mu' \rangle$, and moreover here $\sigma' \models P'$ and $\sigma_i \rightsquigarrow \sigma' \models \bar{\eta}$. Also, by hypothesis for B , $\bar{\delta}$ is respected at each step. After a couple more steps we reach $\langle D, \sigma'', \mu' \rangle$ where $\sigma'' = \text{retract}(x', \sigma')$. This is one of the possible configurations $\langle D, \sigma', \mu_i \rangle$ reached by $\xrightarrow{\Delta'}$ so we are at a matching point in the computations and can proceed by induction on the rest of the computation of $\xrightarrow{\Delta}$.

6. Examples

Toy memory manager. This example is adapted from the one of O’Hearn et al [27] who use it to exemplify ownership transfer. It also exemplifies situations where clients retain access to locations on which which they are not currently allowed to act. Those references must be known when reasoning about the client. This might be done by tracking a “typestate” that distinguishes between, e.g., updatable and non-updatable phases of an object’s life [5]. We will do it by tracking the set of objects that are currently in the “freed” state. This is a “static” module, i.e., unlike our other examples it is not individual instances of a class that represent an abstraction, but rather some global variables.

```
global-vars result : Node; freed : rgn;
module-vars flist : Node; count : int;
alloc()
{ if count = 0 then result := new Node();
  else freed := freed - {flist};
  result := flist; flist := flist.nxt; count--=1; }
free(n : Node)
{ n.nxt := flist; flist := n; count++=1; freed := freed ∪ {n}; }
```

The interface specifications are in Fig. 12. The module invariant is I_M defined as $FC(\text{flist}, \text{freed}, \text{count})$. Here FC is defined by induction on the size of its region parameter:

```
FC(f : Node, r : rgn, c : int) :
  (f = null ⇒ r = ∅ ∧ c = 0)
  ∧ (f ≠ null ⇒ f ∈ r ∧ c > 0 ∧ FC(f.nxt, r - {f}, c - 1))
```

So I_M says freed is the nodes reached from flist and count is the size. As dynamic boundary we choose

$$\bar{\delta}_M : \text{rd freed}, \text{freed}^* \text{nxt}$$

To be precise, the read footprint of I_M includes also flist and count , but we shall assume they are hidden from the client by ordinary scoping.

The implementation of alloc relies on accuracy of count . In particular, it relies on $\text{count} \neq 0 \Rightarrow \text{flist} \neq \text{null}$ (as otherwise the dereference $f.\text{nxt}$ could fault), but for this to hold on subsequent calls the stronger condition I_M needs to be maintained as invariant.

Consider this strange client that both reads and writes data in the free list—but not in a way that interferes with the module.

```
var x, y : Node; x := new Node();
{x ∈ freed}
alloc(); y := result;
{x ∈ freed ∧ y ∈ freed} // by spec of alloc
free(x); free(y);
while y ≠ null { y.val := 7; y := y.nxt; }
{x ∈ freed ∧ y ∈ freed}
```

This is provable using the specifications (Fig. 12) and rules DYNBND-INTRO and DYNBNDINTROM. The point is that although the loop updates val , that effect is separate from the dynamic boundary, $\bar{\delta}_M$.

Combining modules. Consider a client program that uses both the memory manager and the Set module of Sec. 2:

```
let alloc, free be ... in let add, contains, remove be ... in Cli
```

¹¹ In part (c) of the proof of the Claim B, the argument would also apply to a σ' arising as a result of a call to a constructor method. The postcondition of such a method must be of the form $Q' \otimes I$ and thus I must be established by σ' (replacing the appeal to Lemma 6).

Method	Pre-condition	Post-condition	Effects
$alloc()$	true	$result \neq null \wedge freed = old(freed) - result$	$wr\ result, freed, alloc, freed^*nxt$
$free(n : Node)$	$n \neq null \wedge n \notin freed$	$freed = old(freed) \cup \{n\}$	$wr\ freed, freed^*nxt$

Figure 12. Interface specifications for memory manager. Dynamic boundary: rd $freed, freed^*nxt$

(eliding the parameter lists, implementations, and constructor to save space). Let us write Δ for the two specifications in Fig. 12 and Θ for the four in Fig. 2. The main program Cli is verified in context $\Delta\langle\bar{\delta}_M\rangle; \Theta\langle\bar{\theta}_I\rangle$, i.e., it must respect the dynamic boundary $\bar{\delta}_M$ for the memory manager and $\bar{\theta}_I$ for the *Set* module. The implementations of *add*, *contains*, etc. are verified in context $\Delta\langle\bar{\delta}_M\rangle; \Theta\langle\rangle$, i.e., they may use the memory manager but must respect its dynamic boundary. Finally, *alloc* and *free* are verified in context $\Delta\langle\rangle$.

Observer pattern. This example illustrates the verification of a client of the Observer pattern (Figs. 13, 14). A subject has a list of observers cognizant of its internal state (here represented as an integer value held in field *val* of a subject). A new observer registers itself with a subject (the observer’s *sub* field tracks its subject) and is notified of its subject’s state which it caches in field *cacheVal*. When a subject’s state is updated it notifies each of its observers. Each observer then calls back into the subject interface and gets the new value.

A subject and its observers together form a cooperating cluster of objects that are not in an ownership relation. Accordingly we can consider a module containing both classes *Subject* and *Observer* whose interface exports the methods *update*, *get* and the constructor for *Observer*¹². In Fig. 13, $s.O$ is a region containing the observers of subject s . The subject employs a list data structure to manage its observers — this is hidden from clients as is the list header, $s.obs$. Let $SubH(s)$ be the predicate $List(s.obs, s.O)$ that says “for subject s , list beginning at $s.obs$ lies in region $s.O$ ”. The exact definition of $List$ is immaterial here (but see [3]); it suffices to know that the read effects of $SubH(s)$ are $rd\ s, \{s\}^*(obs, O), \{s\}^*O^*nxt$.

Hidden invariant, I , is defined as $I(\emptyset, \emptyset)$ where $I(m, n)$ is

$$\begin{aligned} & (\forall s : Subject \in alloc - m \mid SubH(s)) \wedge \\ & (\forall o : Observer \in alloc - n \mid o.sub \neq null \Rightarrow o \in o.sub^*O) \end{aligned}$$

The second conjunct says that any observer, o , with non-null subject $o.sub$, is in the subject’s O region. The specifications use the predicate $Obs(b, s, v) : b.sub = s \wedge b.cacheVal = v$. Its read effect is $rd\ b, s, v, \{b\}^*(sub, cacheVal)$.

Suppose there are two subjects s, t with $s.val = 0$ and $t.val = 5$. Here is a client, C :

```
o := new Observer(s); p := new Observer(t); s.update(2);
```

Then we can show, using only the specifications in Fig. 14 that after the call to *update*, $p.cacheVal = 5$ while $o.cacheVal = 2$.

Let $P : \forall b : Observer \in \{s\}^*O \mid Obs(b, s, 0)$
 $P' : \forall b : Observer \in \{s\}^*O \mid Obs(b, s, 2)$
 $Q : \forall b : Observer \in \{t\}^*O \mid Obs(b, t, 5)$
 $W : P \wedge Q \wedge s.val = 0 \wedge t.val = 5$

in the proof outline

¹² It also exports the constructor for *Subject* but we elide it. The other procedures are module scoped.

```
class Subject { obs : Observer; val : int; O : rgn;
  Subject() { obs := null; val := 0; O := emp; }
  register(b : Observer) { add(b); b.notify(); }
  update(n : int) { val := n; b : Observer := obs;
    while b != null { b.notify(); b := b.nxt; } }
  get() : int { return self.val; }
  add(b : Observer) { O := O ∪ {b}; b.nxt := obs; obs := b; } }

class Observer { sub : Subject; cacheVal : int; nxt : Observer;
  Observer(u : Subject) { sub := u; u.register(self); }
  notify() { cacheVal := sub.get(); } }
```

Figure 13. Subject/Observer implementation.

```
{ W }
{ P ∧ Q ∧ s ≠ t ... } // by CONSEQ
o := new Observer(s);
{ P ∧ Q ∧ o ∈ {s}^*O } // by s ≠ t, spec of Observer, FRAME
p := new Observer(t);
{ P ∧ Q ∧ o ∈ {s}^*O ∧ p ∈ {t}^*O ... } // ditto
s.update(2);
{ P' ∧ Q ∧ o ∈ {s}^*O ∧ p ∈ {t}^*O ... } // s ≠ t, update, FRAME
{o.cacheVal = 2 ∧ p.cacheVal = 5 ... } // by CONSEQ
```

Verification of Cli : let $get()$ be ...
 $update(n : int)$ be ...
 $Observer(u : Subject)$ be ... in C
requires showing (by LINK') $\vdash \{ W \wedge I \} Cli \{ W' \wedge I \} [\bar{\varepsilon}]$
where W' is $s.val = 2 \wedge t.val = 5$ and $\bar{\varepsilon}$ is $wr\ o, p, alloc,$
 $wr\ \{s\}^*(O, val, dg), s.O^*(nxt, cacheVal), \{t\}^*(O, dg), t.O^*nxt$.
Dynamic boundary $\bar{\theta}$ is $rd\ alloc^*(O, dg), alloc^*O^*nxt$. It is easy
to see $I \vdash (\bar{\theta}, rd\ alloc)$ frm I . Read effect $alloc^*sub^*O$ of I is
subsumed by $alloc^*O$ because $alloc^*sub \subseteq alloc$. We first need

$$\Theta(\bar{\theta}) \vdash \{ W \} C \{ W' \} [\bar{\varepsilon}] \quad (8)$$

where Θ is a list containing the specifications in Fig. 14. But using the proof outline above established $\Theta\langle\rangle \vdash \{ W \} C \{ W' \} [\bar{\varepsilon}]$ from which (8) follows by a few applications of DYNBNDINTROM and sequencing. Therefore, first by SOF, we can conjoin I and then by DYNBNDDELIM drop $\bar{\theta}$ to obtain

$$(\Theta \otimes I)\langle\rangle \vdash \{ W \wedge I \} C \{ W' \wedge I \} [\bar{\varepsilon}]$$

Now for each $\{Q\}m\{Q'\}[\bar{\eta}]$ in Θ we must show

$$(\Theta \otimes I)\langle\rangle \vdash \{ Q \wedge I \} B_m \{ Q' \wedge I \} [\bar{\eta}]$$

where Q and Q' are the pre- and post-conditions and $\bar{\eta}$ are the effects according to Fig. 14 and B_m is the implementation of m . For example, to verify *Observer*’s body, B_o , we factor I into $J \wedge I(\{u\}, \{self\})$ and establish $\vdash \{ Q \wedge J \} B_o \{ Q' \wedge J \} [\bar{\varepsilon}_o]$, where J is $SubH(u) \wedge sub \neq null \Rightarrow self \in sub^*O$. Now FRAME yields $\vdash \{ Q \wedge I \} B_o \{ Q' \wedge I \} [\bar{\varepsilon}_o]$.

7. Related work

The influence of SL on our work is clear. We follow the direction of several recent works [10, 16] that have exploited ghost state

Method	Pre-condition	Post-condition	Effects
update(n) get	$\forall b: Observer \in \{\text{self}\}^* O \mid Obs(b, \text{self}, \text{self}.val)$ <i>true</i>	$\forall b: Observer \in \{\text{self}\}^* O \mid Obs(b, \text{self}, n)$ <i>res = val</i>	$wr \{\text{self}\}^* val, \text{self}.O^* cacheVal$
Observer(u)	$\forall b: Observer \in \{u\}^* O \mid Obs(b, u, u.val)$	$\forall b: Observer \in \{u\}^* O \mid Obs(b, u, u.val)$ $\wedge \text{self} \in \{u\}^* O$	$wr \{u\}^* O^* next, \{u\}^*(O, dg)$

Figure 14. Specifications for Subject/Observer example. *Observer*'s effects publishes data group *dg* that abstracts from *obs*.

to abstract away low-level reasoning (e.g. maintain reachability information as needed while not saying so explicitly in specs.) The price is that programs need to be instrumented with ghost updates and our specifications are less compact in contrast to [27].

Both our Frame rule and our SOF rule use ordinary conjunction to introduce an invariant, together with side conditions that designate a footprint with respect to which the invariant is separated from the write effect of a command. By contrast, in SL these rules use the separating conjunction $*$ which expresses the existence of such a footprint. Reynolds gave a derivation using the rule of conjunction that shows the SOF is not sound in SL (the “conundrum”) without restriction to predicates that are “precise” in the sense of determining a unique footprint [27]. A predicate I is precise iff $(I * _)$ distributes over \wedge , i.e., $I * (Q \wedge R) \Leftrightarrow (I * Q) \wedge (I * R)$. In an unpublished proof of admissibility of SOF in region logic, Naumann proved the case of the conjunction rule using that $(I \wedge _)$ distributes over \wedge . The admissibility proof essentially says that a use of SOF can be replaced by explicitly conjoining the invariant throughout the proof, justified by many uses of ordinary FRAME to introduce not only the invariant but also a chosen footprint.

In earlier work [1] we automatically verified the Observer pattern in Boogie, with the invariant hidden — by postulate. Our current work validates the postulate by way of the SOF. Through Peter Müller (personal communication) we have learnt that our Boogie encodings of regions have been adapted to successfully verify the priority inheritance protocol in real time OS code.

Drossopoulou et al. [13] introduce a general framework to describe verification techniques for invariants based on visible state semantics which requires all invariants to hold on all method call/return boundaries. The framework handles subtyping as well. A general result shows that certain constraints on the framework parameters are sufficient for soundness. A number of ownership type disciplines from the literature are studied as instances of the framework, which promises great improvement in comparing and assessing disciplines. In its present form, the framework does not encompass disciplines like Boogie [4] that rely on state-based encapsulation rather than types or disciplines that deal with design patterns. However, the framework does handle callbacks.

One reaction to the difficulty of hiding invariants is to abandon hiding entirely, in favor of abstraction as advocated by Bierman and Parkinson [6]. They use second order separation logic: method specifications refer to “abstract predicates” that are existentially quantified over the specification. This has been implemented in the jStar prototype [12] based on symbolic execution of SL assertions. Parkinson has reacted to the challenge of invariants over object clusters by advocating explicit but abstract invariants in method specifications [30]. We think more experience with realistic clients is needed to evaluate the practicality of explicitly carrying around a conjunct for the invariant of each abstraction in use.

Smans et al. [34] build a prototype verifier to explore ways to do dynamic framing in the setting of ordinary assertion languages (and also SL), gaining precision through use of location sets and abstraction via abstraction (model fields and pure methods) also in their recent work [33]. Leino’s Dafny language features dynamic frames of the

form G^* any as well as pure functions that can compute recursive predicates like our *FC* and *List* examples [22].

8. Conclusion

We presented an imperative notion of module interface, which complements static scoping constructs with state dependent expression of the dynamic part of the encapsulation boundary. Each primitive action of the client must respect this boundary — it is not enough that the end to end effect of the client respects it. This entails a small step interpretation of specifications. (It is achieved in SL using a big step interpretation but a stronger separation property.) What is achieved is sound and flexible modular reasoning that encompasses various design patterns. We conclude that Hoare’s “mis”match is rather an intricate match between modules and their clients and we show how to get it right.

Apart from the specifications for interface methods, a specifier needs to make explicit the dynamic encapsulation boundary. Then the verifier can fulfill its usual obligations, namely, generate verification conditions (a) for the client by using the interface method specifications and (b) for the bodies of interface methods taking the module invariant into account. But the verifier also needs to verify that each primitive action of the client respects the boundary.

In future work we plan to extend our results to subclassing and inheritance as well as state based representation independence [2]. A particularly exciting challenge is callbacks. Our rendition of the Observer pattern (following Parkinson [30]) has callbacks within the module. It appears that the small-step correctness property enables a SOF rule that supports callbacks back and forth across dynamic encapsulation boundaries.

References

- [1] A. Banerjee, M. Barnett, and D. A. Naumann. Boogie meets regions: A verification experience report. In *VSTTE*, pages 177–191, 2008.
- [2] A. Banerjee and D. A. Naumann. State based ownership, reentrance, and encapsulation. In *ECOOP*, pages 387–411, 2005.
- [3] A. Banerjee, D. A. Naumann, and S. Rosenberg. Regional logic for local reasoning about global invariants. In *ECOOP*, 2008.
- [4] M. Barnett, K. R. M. Leino, and W. Schulte. The Spec# programming system: An overview. In *CASSIS*, volume 3362 of *LNCS*, pages 49–69, 2005.
- [5] N. E. Beckman, K. Bierhoff, and J. Aldrich. Verifying correct usage of atomic blocks and tpestate. In *OOPSLA*, pages 227–244, 2008.
- [6] G. Bierman and M. Parkinson. Separation logic and abstraction. In *POPL*, pages 247–258, 2005.
- [7] C. Boyapati, B. Liskov, and L. Shrira. Ownership types for object encapsulation. In *POPL*, pages 213–223, 2003.
- [8] L. Burdy, Y. Cheon, D. R. Cok, M. D. Ernst, J. R. Kiniry, G. T. Leavens, K. Rustan M. Leino, and E. Poll. An overview of JML tools and applications. *STTT*, 7(3):212–232, 2005.
- [9] D. Clarke and S. Drossopoulou. Ownership, encapsulation and the disjointness of type and effect. In *OOPSLA*, pages 292–310, 2002.
- [10] M. Dahlweid, M. Moskal, T. Santen, S. Tobies, and W. Schulte. VCC: Contract-based modular verification of concurrent C. In *ICSE-Companion*, pages 429–430, 2009.

- [11] W. Dietl and P. Müller. Universes: Lightweight ownership for JML. *Journal of Object Technology*, 4:5–32, 2005.
- [12] D. Distefano and M. J. Parkinson. jstar: Towards practical verification for java. In *OOPSLA*, pages 213–226, 2008.
- [13] S. Drossopoulou, A. Francalanza, P. Müller, and A. J. Summers. A unified framework for verification techniques for object invariants. In *ECOOP*, pages 412–437, 2008.
- [14] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *PLDI*, 2002.
- [15] C. Grothoff, J. Palsberg, and J. Vitek. Encapsulating objects with confined types. *ACM TOPLAS*, 29(6), 2006.
- [16] C. Hawblitzel and E. Petrank. Automated verification of practical garbage collectors. In *POPL*, pages 441–453, 2009.
- [17] C. A. R. Hoare. Proofs of correctness of data representations. *Acta Inf.*, 1:271–281, 1972.
- [18] I. T. Kassios. Dynamic framing: Support for framing, dependencies and sharing without restriction. In *Formal Methods*, 2006.
- [19] N. R. Krishnaswami, J. Aldrich, L. Birkedal, K. Svendsen, and A. Buisse. Design patterns in separation logic. In *TLDI*, 2009.
- [20] G. T. Leavens, K. R. M. Leino, and P. Müller. Specification and verification challenges for sequential object-oriented programs. *Formal Aspects of Computing*, 19(2):159–189, 2007.
- [21] G. T. Leavens and P. Müller. Information hiding and visibility in interface specifications. In *ICSE*, pages 385–395, 2007.
- [22] K. R. M. Leino. Specification and verification in object-oriented software. Marktobendorf lecture notes, 2008.
- [23] K. R. M. Leino and P. Müller. Object invariants in dynamic contexts. In *ECOOP*, pages 491–516, 2004.
- [24] K. R. M. Leino, A. Poetzsch-Heffter, and Y. Zhou. Using data groups to specify and check side effects. In *PLDI*, pages 246–257, 2002.
- [25] P. Müller, A. Poetzsch-Heffter, and G. T. Leavens. Modular invariants for layered object structures. *Sci. Comput. Programming*, 62(3):253–286, 2006.
- [26] D. A. Naumann and M. Barnett. Towards imperative modules: Reasoning about invariants and sharing of mutable state (extended abstract). In *LICS*, pages 313–323, 2004.
- [27] P. O’Hearn, H. Yang, and J. Reynolds. Separation and information hiding. In *POPL*, pages 268–280, 2004.
- [28] P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL*, pages 1–19, 2001.
- [29] P. W. O’Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. *ACM TOPLAS*, 31(3):1–50, 2009.
- [30] M. Parkinson. Class invariants: The end of the road? In *IWACO*, 2007.
- [31] C. Pierik. Validation techniques for object-oriented proof outlines. Dissertation, Universiteit Utrecht, 2006.
- [32] C. Pierik, D. Clarke, and F. S. de Boer. Controlling object allocation using creation guards. In *Formal Methods*, pages 59–74, 2005.
- [33] J. Smans, B. Jacobs, and F. Piessens. Implicit dynamic frames: Combining dynamic frames and separation logic. In *ECOOP*, 2009.
- [34] J. Smans, B. Jacobs, F. Piessens, and W. Schulte. An automatic verifier for Java-like programs based on dynamic frames. In *FASE*, pages 261–275, 2008.
- [35] K. Zee, V. Kuncak, and M. C. Rinard. Full functional verification of linked data structures. In *PLDI*, pages 349–361, 2008.