

# A Logical Analysis of Framing for Specifications with Pure Method Calls

Anindya Banerjee <sup>\*1</sup> and David A. Naumann <sup>\*\*2</sup>

<sup>1</sup> IMDEA Software Institute, Madrid, Spain

<sup>2</sup> Stevens Institute of Technology, Hoboken, USA

**Abstract.** For specifying and reasoning about object-based programs it is often attractive for contracts to be expressed using calls to pure methods. It is useful for pure methods to have contracts, including read effects to support local reasoning based on frame conditions. This leads to puzzles such as the use of a pure method in its own contract. These ideas have been explored in connection with verification tools based on axiomatic semantics, guided by the need to avoid logical inconsistency, and focusing on encodings that cater for first order automated provers. This paper adds pure methods and read effects to region logic, a first-order program logic that features frame-based local reasoning and a proof rule for linking of clients with modules to achieve end-to-end correctness by modular reasoning. Soundness is proved with respect to a conventional operational semantics and using the extensional (i.e., relational) interpretation of read effects.

## 1 Introduction

In reasoning about programs, a frame condition is the part of a method’s contract that says what part of the state may be changed by an invocation of the method. Frame conditions make it possible to retain a global picture while reasoning locally: If predicate  $Q$  can be asserted at some point in a program where method  $m$  is called,  $Q$  still holds after the call provided that the locations on which  $Q$  depends are disjoint from the locations that may be written according to  $m$ ’s frame condition. This obvious and familiar idea is remarkably hard to formalize in a way that is useful for sound reasoning about programs acting on dynamically allocated mutable objects (even sequential programs, to which we confine attention here). One challenge is to precisely describe the writable state in case it involves heap allocated objects. Another challenge is to determine what part of such state may be read by  $Q$  (its ‘footprint’). For reasons of abstraction,  $Q$  may be expressed in terms of named functions. To hide information about data representation, the function definitions may not be visible in the client program where  $m$  is called. This paper provides a foundational theory addressing these and related challenges.

Consider a class `Cell` with instances holding an integer value, used in the following client code.

```
method get(): int  
method set(v: int) ensures self.get() = v  
var c, d: Cell; c:= new Cell; d:= new Cell; c.set(5); d.set(6); assert c.get() = 5;
```

---

\* Currently on leave at the US National Science Foundation.

\*\* Partially supported by US NSF award CNS-1228930.

```

class Cell { private val: int; ghost footprint: rgn;
  pure method get(): int ensures self.get() = result
              reads self.footp {any
  method set(v: int) ensures self.get() = v
              writes self.footp {any //(we elide read effects for impure methods)
              ensures  $\forall o: \text{Cell} . o \neq \text{self} \Rightarrow o.\text{footp} \# \text{self.footp}$  }
}

```

**Fig. 1.** Example: Cell. Type **rgn** (for ‘region’) means sets of object references and **footp {any** denotes a set of locations, namely all fields of objects in **footp**.

The goal is to prove the assertion by reasoning that the state read by `c.get()` is disjoint from the state written by `d.set(6)`. Suppose the internal representation of Cell objects consists of an integer field `val`. The specifications could say `set` writes `self.val` and `get` reads `self.val`. Then the frame condition of `d.set(6)` would allow the postcondition of the call `c.set(5)`, i.e., the predicate `c.get() = 5`, to be framed over the call `d.set(6)`, yielding the assertion. But such specifications expose the internal representation. It would preclude, for example, an alternative implementation that uses, instead of integer field `val`, a pointer to a character string that represents the number using 0s and 1s.

Better specifications appear in Fig. 1, using ghost state to describe the ‘footprint’ of each cell, and postconditions from which the client can deduce disjointness of the representations of `c` and `d`. Use of ghost state for footprints is a key part of the ‘dynamic frames’ approach [8], and in addition to explicit disjointness conditions it supports separation reasoning based on freshness.

The example illustrates another challenging issue: one method (`get`) is used in the specification of others (`get,set`). Here is an example of calling a pure method in a frame condition: Instead of the ghost field `footp` one might choose to define a region-valued method `footpm`. If `footpm` is only used in specifications, one may argue it should be defined as part of the mathematical theory in which reasoning is carried out (though its read effect would still be useful). But there are practical benefits to using programmed methods in specifications, which can be justified provided that they are *pure* in the sense of having no effects other than reading.

Use of pure methods, especially ones in the program rather than part of the ambient mathematical theory, poses challenges. One is how to model such specifications without inconsistency. For example, care must be taken for sound treatment of specifications like **method** `f(x: int): int ensures result = f(x)+1`. Recursion aside, one may also wonder about soundness of using a pure method in its own postcondition, e.g., `get` in Fig. 1, or in its own frame condition: e.g., the read effect of `footpm` might be `footpm() {any` (making it ‘self-framing’). Another issue is that the specifications of `get` and `set` are abstract, in the sense that they are consistent with many interpretations of the function `get` (e.g., `get` could return `self.val+7`). Client code should respect the abstraction, i.e., be correct with respect to any interpretation, whereas the expected implementations of `get` and `set` are only correct with respect to the interpretation that returns `self.val`.

The issues discussed so far have been addressed in prior work, especially in the context of verification-condition generation (VC-gen); see Sec. 7. However, most of the VC-gen work takes axiomatic semantics for granted rather than defining and proving

soundness with respect to operationally grounded program semantics; the focus is on methodological considerations and on encodings that work effectively with SMT-based theorem provers. In these works, hypotheses are encoded as axioms, and linking of separately verified methods is implicit in the implementation of the VC-gen. The intricacies of dealing with heap structure, framing, purity, and self-framing frame conditions have led to soundness bugs in implemented verification systems (see [6]).

This paper provides a foundational account, by way of a conventional logic of programs that caters for SMT-provers by reasoning about framing using ghost state and FOL, and that is proved sound with respect to a standard operational semantics. Our account focuses on a *proof rule for linking* the implementation of an interface (i.e., collection of method specifications) with a client that relies on that interface.

The approach we take is motivated by two additional challenges. The first is information hiding, in the sense that implementations rely on invariants on module-internal data structures, but these invariants do not appear in the interface specification [7]. As a contrived example, representing the integer cell using a string might have the invariant that only 0 and 1 characters appear, without leading zeros. The invariant might be exploited by a method `getAsString`, but it has no place in the interface specification of method `get` which returns an integer. An alternative to hiding is to rely on abstraction: a predicate whose definition is opaque in the interface can be defined internally to be the invariant [15,11,9].

The second additional challenge arises from the practical need to use programmed methods that are only *observationally pure* in the sense that they do have side effects but these effects are benevolent [7] and not observable to clients. There are many examples, including memoization, lazy initialization, and path compression in Union-Find structures. These may involve allocation of fresh objects and mutation of existing ones.

Strong encapsulation is critical both for hiding of invariants [7,14,1] and for observational purity [13]. Both involve linking a client with the implementation of a module, where that implementation is verified against specifications different from those used by clients of the module —hiding invariants and hiding effects. In prior work we developed *region logic* (RL), a Hoare logic for sequential object-based programs, using standard FOL for assertions. By contrast with separation logic and permission-based systems, in RL separation is expressed as disjointness of explicit footprints, following the approach of dynamic frames. A benefit is that the verifier does not need to support separating conjunction; it comes at the cost of more verbose specifications. The language features expressions that denote *regions*, i.e., sets of object references. The logic provides a frame rule for local reasoning, based on frame conditions of methods and a subsidiary judgment for framing of formulas (Fig. 10). In addition to ordinary frame conditions, the logic formalizes encapsulation boundaries for modules, again in the manner of dynamic frames. This supports a second-order frame rule for linking method implementations to clients, hiding invariants [7,14].

In ongoing work, we have extended RL to a relational version, akin to [3,20,12] but featuring a proof rule for representation independence. We plan to use this as basis for a proof system that allows use of observationally pure methods in specifications, which relies on relational consequences of encapsulation [7,13]. The problem is that general relational reasoning depends on read effects in frame conditions, a non-trivial if not

earthshaking extension of RL. It deserves to be studied and presented in isolation from the complications needed for encapsulation and information hiding.

This paper builds on RLI [2] and RLII [1], extending RL with pure method calls in specifications and read effects in frame conditions. This involves adding read effects to frame conditions for commands and for pure and impure methods.

*Outline and contributions.* Sec. 2 introduces the programming language and specifications, as well as the judgment of correctness under hypotheses. The latter is written  $\Delta \vdash C : P \rightsquigarrow Q [\varepsilon]$ . It says that under precondition  $P$  command  $C$  does not fault; if it terminates its final state satisfies  $Q$  and the computation's effects are allowed by  $\varepsilon$ . Moreover this conclusion is under hypothesis  $\Delta$ , a list of method specifications. What's new in this paper is read effects in  $\varepsilon$  and  $\Delta$ , and pure methods used in  $\Delta, P, C, Q, \varepsilon$ , specified in  $\Delta$ . Sec. 3 takes the first step towards defining semantics, sketching two ways to interpret the hypotheses and pointing out a potential circularity. To dodge this circularity, semantics of expressions and formulas is parameterized on the interpretation of pure methods. Sec. 4 formalizes an extensional semantics of read effects; this is used to define correct interpretations of pure method specifications and to define the denotation of impure method specifications. The latter is like a specification statement, and is used in the operational semantics of programs; its first order semantics is justified by a closure property that is our first technical result. Sec. 5 completes the semantics of the correctness judgment, suitably instantiating the interpretation of pure methods as motivated by the rule for linking. For  $C$  verified under hypothesis  $\Delta$  that specifies pure method  $p$  called in  $C$  and/or used in the specification of  $C$ , linking discharges the hypothesis. If the specification of  $p$  is unsatisfiable, it is not possible to instantiate the interpretation as required by the linking rule. This shows that a separate satisfiability check is not needed in a tool that correctly verifies the linked program, though the check maybe be helpful to flag problems early.

Sec. 6 gives selected proof rules and states the main result, soundness of the rules. Proofs and technical details that we gloss over can be found in the full version of the paper. The proofs are intricate because we work directly with small-step operational semantics, yet this is essential for the use of dynamic frames to provide flexible encapsulation of modules in RLII. But the proofs are elementary and do not involve fixpoints.

Sec. 7 briefly discusses related work. For future work, the next steps towards observational purity are (a) to extend the logic with second order framing, as in RLII but with hiding of effects, and (b) to add weak purity which allows allocation though not other effects (this is not hard but does add a few complications). Another step is to add read effects and pure methods to our prototype SMT-based verifier for RL [16,17], which already provides limited support for pure function definitions with framing, based on a version of Leino's Dafny. As a first step, we have successfully checked versions of the Cell example by manual encoding in Why3, using SMT-provers only.

## 2 Programs, specifications, and correctness judgments

Fig. 2 illustrates features of our programming and specification notations, by way of the Composite pattern, a well-known verification challenge problem [16,4]. A Comp is the root of a tree, nodes of which are accessible to clients. Here is an example client:

```

class Comp {
  private children: listOf(Comp);
  specpublic parent: Comp; // (private but visible in specifications)
  specpublic size: int := 1; // number of descendants
  ghost desc: rgn; // set of descendants

  method addChild(x: Comp)
    requires x ≠ null ∧ x.parent = null ∧ ...
    ensures // x is added as a child self
    writes children, x.parent, ancestors(self)‘size, ancestors(self)‘desc
  pure method getSize(): int
    reads size ensures result = size
  pure method ancestors(p: Comp): rgn
    reads ancestors(p)‘desc, ancestors(p)‘parent
    ensures result = { o ∈ alloc | type(o,Comp) ∧ p ∈ o.desc } }

```

**Fig. 2.** Composite example, adapted from RLI. Ghost code maintains the invariant that desc is the set of descendants.

```

var b, c, d: Comp; var i: int; ... i := d.getSize(); b.add(c); assert i = d.getSize();

```

To prove the assertion we want to frame the formula  $i = d.\text{getSize}()$  over the call  $b.\text{add}(c)$ . The frame condition of `addChild` says it is allowed to write `self.children`, `x.parent`, and the `size` and `desc` fields of the ancestors of `self`. In method set (Fig. 1) we use ‘`any`’ to abstract from field names, but here both `size` and `desc` are appropriate to make visible in the interface. (See RLI for more discussion of this facet of information hiding.) The frame condition would be less precise using `ancestors(self)‘any`.

In order to reason using the frame rule (Fig. 10), we establish a subsidiary judgment written  $\vdash \text{rd } i, d, d.\text{size} \text{ frm } i = d.\text{getSize}()$  which says the formula  $i = d.\text{getSize}()$  depends only on the values of  $i$ ,  $d$ , and  $d.\text{size}$ . The rules let us establish this judgment based on the specification of `getSize`. The frame rule also requires us to establish validity of a so-called *separator formula*. This formula is determined from the frame of the formula and from the write effect of `addChild`. The function  $\cdot/\cdot$  generates the separator formula and is defined by recursion on syntax.<sup>3</sup> In the example, we compute  $\varepsilon \cdot/\cdot (\text{rd } i, d, d.\text{size})$ , where  $\varepsilon$  is the write effect of `addChild`. The formula is the disjointness  $\{d\} \# \text{ancestors}(b)$ , which says the singleton region  $\{d\}$  is disjoint from the set of ancestors. It needs to hold following the elided part of the example client. In general,  $\eta \cdot/\cdot \varepsilon$  is a formula which implies that the locations writable according to  $\varepsilon$  are disjoint from the locations readable according to  $\eta$ .

Fig. 3 gives the grammar of programs, revised from RLII to allow method calls in expressions. We assume given a fixed collection of classes. A class has a name and some typed fields. We do not formalize dynamic dispatch or even associate methods with classes; so the term ‘method’ is just short for procedure. For expository clarity methods have exactly one parameter (plus `res` for pure methods).

<sup>3</sup> Please note that  $\cdot/\cdot$  is not syntax in the logic; it’s a function in the metalanguage that is used to obtain formulas from effects; see Sec. 6.

$m, p \in \text{MethName}$	$x, y \in \text{VarName}$	$f, g \in \text{FieldName}$	$K \in \text{DeclaredClassNames}$
(Types)	$T ::= \text{int} \mid \text{rgn} \mid \text{Obj} \mid K$		
(Program Expressions)	$E ::= x \mid c \mid \text{null} \mid E \oplus E \mid m(E)$ where $c$ is in $\mathbb{Z}$ , $\oplus$ is in $\{=, +, \dots\}$		
(Region Expressions)	$G ::= \emptyset \mid x \mid \{E\} \mid G^{\cdot} f \mid G \otimes G \mid m(F)$ where $\otimes$ is in $\{\cup, \cap, \setminus\}$		
(Expressions)	$F ::= E \mid G$		
(Commands)	$C ::= \text{skip} \mid x := F \mid x := \text{new } K \mid x := x.f \mid x.f := F$ $\quad \mid \text{if } E \text{ then } C \text{ else } C \mid \text{while } E \text{ do } C \mid C; C \mid \text{var } x:T \text{ in } C$ $\quad \mid m(x) \mid \text{let } m(x:T) = C \text{ in } C \mid \text{let } m(x:T, \text{res}:U) = C \text{ in } C$		

**Fig. 3.** Programming language, highlighting additions to RLII [1].

The linking construct  $\text{let } m(x:T, \text{res}:U) = C \text{ in } C'$  designates that  $m$  is pure, with return type  $U$ , as indicated by the distinguished variable name  $\text{res}$ . It binds  $x$ ,  $\text{res}$ , and  $m$  in  $C$ , and  $m$  in  $C'$ . Calls of  $m$  are expressions and pass a single argument. The body  $C$  is executed in a state with both  $x$  and  $\text{res}$ , the latter initialized to the default value for type  $U$ . The final value of  $\text{res}$  is the value of the call expression. The linking construct  $\text{let } m(x:T) = C \text{ in } C'$  designates that  $m$  is impure; command  $m(y)$  depicts its call.

Typing contexts, ranged over by  $\Gamma$ , are finite maps, written in conventional form. The judgment  $\Gamma \vdash E : T$  means that  $E$  is well-formed and has type  $T$ . The typing rules straightforward. A command  $C$  is **well-formed in context**  $\Gamma$  provided that it is typable, i.e.,  $\Gamma \vdash C$ , and in addition method call expressions  $m(F)$  occur only in assignments  $x := m(F)$  to a simple variable and with  $F$  free of method calls.

Values of type  $K$  are references to objects of class  $K$  (including the improper reference **null**). Value of type  $\text{rgn}$  are sets of references of any type. If  $\Gamma \vdash G : \text{rgn}$  then  $\Gamma \vdash G^{\cdot} f : \text{rgn}$  for any field name  $f$  of region or reference type. In case  $f : K$ , the value of  $G^{\cdot} f$  is the set of  $f$ -values of objects in  $G$ . In case  $f : \text{rgn}$ , the value of  $G^{\cdot} f$  is the union of the  $f$ -values. Aside from allocation and dereference (in the command forms  $x := y.f$  and  $y.f := F$ ), the only operation on references is equality test.

The syntax of formulas is standard and unchanged from RLI (Sec. 4.2), except that now the expressions include method calls, as in the points-to predicate  $x.f = E$  and region containment  $G \subseteq G'$ .

$$P ::= E = E \mid x.f = E \mid G \subseteq G \mid (\forall x : K \in G. P) \mid P \wedge P \mid \neg P$$

The formula  $\forall x : K \in G. P$  quantifies over all non-null references of type  $K$  in  $G$ . For disjointness of regions it is convenient to write  $G \# H$  for  $G \cap H \subseteq \{\text{null}\}$ .

*Specifications.* **Effects** are given by  $\varepsilon ::= \text{rd } x \mid \text{rd } G^{\cdot} f \mid \text{wr } x \mid \text{wr } G^{\cdot} f \mid \text{fr } G \mid \varepsilon, \varepsilon \mid$  (*empty*). Effects must be **syntactically well-formed (swf)** for the context  $\Gamma$  in which they occur:  $\text{rd } x$  and  $\text{wr } x$  are swf if  $x \in \text{dom}(\Gamma)$ ;  $\text{rd } G^{\cdot} f$ ,  $\text{wr } G^{\cdot} f$ , and  $\text{fr } G$  are swf if  $G$  is swf in  $\Gamma$ . In particular, if  $G$  is a call  $m(F)$  to a pure method, then it must be that  $\Gamma \vdash m(F) : \text{rgn}$ . The freshness effect  $\text{fr } G$  says the value of  $G$  in the final state contains only (but not necessarily all) references that were not allocated in the initial state. Later we use the term 'well-formed', without qualification, to mean in addition that the expressions do not depend on pure methods invoked outside their preconditions.

**Specifications** for impure methods take the form  $(x:T)R \rightsquigarrow S[\eta]$  and for pure methods the form  $(x:T, \text{res}:U)R \rightsquigarrow S[\eta]$ :  $x$  is the parameter (passed by value),  $R$  the

$$\begin{aligned}
df(m(F), \Delta) &= df(P_F^x, \Delta) \wedge df(F, \Delta) \wedge P_F^x \text{ where } \Delta(m) = (x : T, \text{res} : U)P \rightsquigarrow Q[\varepsilon] \\
df(x.f = E, \Delta) &= x \neq \text{null} \Rightarrow df(E, \Delta) \\
df(G_1 \subseteq G_2, \Delta) &= df(G_1, \Delta) \wedge df(G_2, \Delta) \\
df(\forall x : K \in G. P, \Delta) &= df(G, \Delta) \wedge \forall x : K \in G. df(P, \Delta) \\
df(P_1 \wedge P_2, \Delta) &= df(P_1, \Delta) \wedge (P_1 \Rightarrow df(P_2, \Delta))
\end{aligned}$$

**Fig. 4.** Definedness formulas for expressions and formulas (selected), for swf method context  $\Delta$ .

precondition,  $S$  the postcondition, and  $\eta$  the effects. For these specifications to be swf in context  $\Gamma$ ,  $\eta$  must not include  $wr x$ . (This is standard in Hoare logic; postconditions refer to initial parameter values.) Moreover,  $R$  must be typable in  $\Gamma, x : T$ . Both  $S$  and  $\eta$  must be typable in  $\Gamma, x : T$ , for the impure form or  $\Gamma, x : T, \text{res} : U$ , for the pure form. Finally, for the pure method there can be no write effects in  $\eta$ . Although the body of a pure method will write  $\text{res}$ , the semantics is a return value, not an observable mutation of state. In this paper, there’s no need for impure methods to have read effects, but they will be needed for reasoning about data abstraction and observational purity. In any context  $\Gamma$ , there is a read effect that imposes no restriction:  $\text{rd vars}(\Gamma), \text{rd alloc}$  any.

A **method context**  $\Delta$  is a finite map from method names to specifications. We are interested in specifications that may refer to global variables declared in some typing context  $\Gamma$ . Moreover, specifications in  $\Delta$  are allowed to refer to any of the pure methods in  $\Delta$ ; the specification of  $p$  may have calls to  $p$  in its post-condition and effect, or  $p$  and  $m$  may refer mutually to each other—subject to the restriction that calls in preconditions must exhibit acyclic dependency. To make this restriction precise, we define a relation  $\prec_\Delta$  on method names:  $m \prec_\Delta m'$  iff  $m'$  occurs in the precondition of  $\Delta(m)$ . Now we can define what it means for a context  $\Delta$  to be swf in  $\Gamma$ . First, the transitive closure,  $\prec_\Delta^+$ , is irreflexive. Second, the domains of  $\Gamma$  and  $\Delta$  are disjoint and each specification is swf in the context  $\text{vars}(\Gamma), \text{sigs}(\Delta)$ . Here  $\text{sigs}$  extracts the types of methods. For example, let  $\Delta_0$  be  $m : (x : T)R \rightsquigarrow S[\eta]$ ,  $p : (y : V, \text{res} : U)P \rightsquigarrow Q[\varepsilon]$ . Then  $\text{sigs}(\Delta_0)$  is  $m : (x : T), p : (y : V, \text{res} : U)$ . Also  $\text{vars}$  discards method declarations.

A **correctness judgment** takes the form  $\Delta \vdash^\Gamma C : P \rightsquigarrow Q[\varepsilon]$ . It is swf iff  $\Delta$  is swf in  $\Gamma$  and  $C, P, Q, \varepsilon$  are all swf in  $\text{vars}(\Gamma), \text{sigs}(\Delta)$ . We often elide  $\Gamma$ .

Sound proof rules for correctness judgments prevent a pure method from being applied outside its precondition, to avoid the need to reason about undefined or faulty values. As is common in VC-generation, we use **definedness formulas**, see Fig. 4. The idea is that in states where  $df(P, \Delta)$  holds, evaluation of  $P$  does not depend on values of pure methods outside their preconditions. Although the clause for  $df(m(F), \Delta)$  refers to a method specification that may refer to another pure method in its precondition,  $df$  is well-defined, owing to the requirement that  $\prec_\Delta^+$  is irreflexive (and  $\text{dom } \Delta$  is finite).

An expression or formula will be considered well-formed if its definedness formula is valid, in addition to it being swf. To define validity, we proceed to semantics.

### 3 Semantics of expressions and formulas

There are two approaches to semantics of a judgment  $\Delta \vdash C : P \rightsquigarrow Q[\varepsilon]$ . The first goes by quantifying over all correct implementations of the procedures specified by  $\Delta$ . The second goes by using nondeterminacy to represent a ‘worst implementation’ of each

procedure, akin to the ‘specification statement’ used in axiomatic semantics. The second avoids a quantification and has been found to be quite effective [14,1]; we use it for impure methods (and in so doing show how the specification statement can include read effects). However, for pure method calls in formulas conventional semantics requires determinate values, so we use the first approach for pure methods.<sup>4</sup>

The transition semantics uses an environment for let-bound methods. A call to such  $m$  results in execution of the body found in the method environment. By contrast, if  $m$  is declared in  $\Delta$  then its call is a single step in accord with its specification. If  $m$  is impure, the step goes to any state allowed by the specification  $\Delta(m)$ ; we describe this by a relation  $\llbracket \Delta(m) \rrbracket$  (Def. 2). If  $m$  is pure, we need a determinate result value but no change of state. So we use a function  $\theta(m)$  to provide this value. The semantics of a correctness judgment (Def. 5) quantifies over all  $\theta$  such that  $\theta(m)$  conforms to the specification  $\Delta(m)$  for each  $m$  in  $\text{dom } \Delta$  (Def. 3). This is similar to an axiomatic semantics where  $\theta(m)$  is an uninterpreted function constrained by  $\Delta(m)$ .

To define what it means for  $\theta(m)$  to conform, and to define  $\llbracket \Delta(m) \rrbracket$ , we need semantics of expressions, formulas, and effects —and these depend on the meaning of pure method calls. To break this circularity, we define in this section a notion of candidate interpretation, and define the semantics of formulas and expressions with respect to any candidate interpretation  $\theta$ .

We assume given an infinite set  $\text{Ref}$  of reference values including a distinguished ‘improper reference’  $\text{null}$ . A  $\Gamma$ -*state* is comprised of a global heap and a store. The store is a type-respecting assignment of values to the variables in  $\Gamma$  and to the variable  $\text{alloc} : \text{rgn}$  which is special. Updates of  $\text{alloc}$  are built into the program semantics so that  $\text{alloc}$  holds the set of all allocated references. We write  $\sigma(x)$  for the value of variable  $x$  in state  $\sigma$ ,  $\sigma(o.f)$  to look up field  $f$  of object  $o$  in the heap,  $\text{Dom}(\sigma)$  for the variables of  $\sigma$ , and  $\llbracket \Gamma \rrbracket$  for the set of  $\Gamma$ -states. We write  $\llbracket T \rrbracket \sigma$  for the set of values of type  $T$  in state  $\sigma$ . Thus  $\llbracket \text{int} \rrbracket \sigma = \mathbb{Z}$  and  $\llbracket K \rrbracket \sigma = \{\text{null}\} \cup \{o \mid o \in \sigma(\text{alloc}) \wedge \text{Type}(o, \sigma) = K\}$ . Besides states, we use the faulting outcome  $\frac{1}{2}$  for runtime errors (null-dereference), and also to signal precondition violations as described later.

For a typing context  $\Gamma$ , a *candidate  $\Gamma$ -interpretation*  $\theta$  is a mapping from the pure method names in  $\Gamma$  such that if  $\Gamma(m) = (x : T, \text{res} : U)$  then  $\theta(m)$  is a function such that for any  $T$ -value  $t$  and state  $\sigma$ ,  $\theta(\sigma, t)$  is a  $U$ -value or  $\frac{1}{2}$ . To be precise,  $\theta(m)$  has the dependent type  $(\sigma \in \llbracket \Gamma \rrbracket) \times \llbracket T \rrbracket \sigma \rightarrow (\llbracket U \rrbracket \sigma \cup \{\frac{1}{2}\})$ . A *candidate  $\Delta$ -interpretation* is just a candidate  $\text{sigs}(\Delta)$ -interpretation.

The denotation of an expression  $F$  in candidate  $\Gamma$ -interpretation  $\theta$  and state  $\sigma$  is written  $\llbracket F \rrbracket_{\theta} \sigma$  and defined straightforwardly, see Fig. 5. The second line in the figure is for application  $m(F)$  of a pure method: evaluate  $F$  to get a value  $v$ , then apply the function  $\theta(m)$  to the pair  $(\sigma, v)$ . Using the semantics for expressions we define the 3-valued semantics of formulas in Fig 6. We also define  $\sigma \models_{\theta}^{\Gamma} P$  iff  $\llbracket P \rrbracket_{\theta} \sigma = \text{true}$ . Fig. 7 shows that when the definedness formulas hold, the usual 2-valued clauses hold.

<sup>4</sup> This does not preclude nondeterminacy modulo an equivalence relation, which is especially important for ‘weakly pure’ methods that return freshly allocated references [13]. For VCs this is explored in [10].

$$\begin{aligned}
\llbracket E_1 + E_2 \rrbracket_{\theta} \sigma &= \text{let } v_1 = \llbracket E_1 \rrbracket_{\theta} \sigma \text{ in let } v_2 = \llbracket E_2 \rrbracket_{\theta} \sigma \text{ in } v_1 + v_2 \\
\llbracket m(F) \rrbracket_{\theta} \sigma &= \text{let } v = \llbracket F \rrbracket_{\theta} \sigma \text{ in } \theta(m)(\sigma, v) \\
\llbracket \{E\} \rrbracket_{\theta} \sigma &= \text{let } v = \llbracket E \rrbracket_{\theta} \sigma \text{ in } \{v\} \\
\llbracket G^{\cdot} f \rrbracket_{\theta} \sigma &= \text{let } X = \llbracket G \rrbracket_{\theta} \sigma \text{ in } \{ \sigma(o.f) \mid o \in X \wedge o \neq \text{null} \wedge \text{Type}(o, \sigma) = \text{DeclClass}(f) \} \\
&\quad \text{if } f : K \text{ for some } K \\
&= \text{let } X = \llbracket G \rrbracket_{\theta} \sigma \text{ in } \bigcup \{ \sigma(o.f) \mid o \in X \wedge o \neq \text{null} \wedge \text{Type}(o, \sigma) = \text{DeclClass}(f) \} \\
&\quad \text{if } f : \text{rgn}
\end{aligned}$$

**Fig. 5.** Semantics of selected program and region expressions, for state  $\sigma$  and candidate interpretation  $\theta$ . We use the  $\zeta$ -strict let-binder, i.e., ‘let  $v = X$  in  $Y$ ’ denotes  $\zeta$  if  $X$  denotes  $\zeta$ .

$$\begin{aligned}
\llbracket E_1 = E_2 \rrbracket_{\theta} \sigma &= \text{let } v_1 = \llbracket E_1 \rrbracket_{\theta} \sigma \text{ in let } v_2 = \llbracket E_2 \rrbracket_{\theta} \sigma \text{ in if } v_1 = v_2 \text{ then true else false} \\
\llbracket x.f = E \rrbracket_{\theta} \sigma &= \text{if } \sigma(x) = \text{null} \text{ then false else let } v = \llbracket E \rrbracket_{\theta} \sigma \text{ in} \\
&\quad \text{if } \sigma(x.f) = v \text{ then true else false} \\
\llbracket x.f = E \rrbracket_{\theta} \sigma &= \text{let } v = \llbracket E \rrbracket_{\theta} \sigma \text{ in if } \sigma(x) \neq \text{null} \text{ and } \sigma(x.f) = v \text{ then true else false} \\
\llbracket G_1 \subseteq G_2 \rrbracket_{\theta} \sigma &= \text{let } X_1 = \llbracket G_1 \rrbracket_{\theta} \sigma \text{ in let } X_2 = \llbracket G_2 \rrbracket_{\theta} \sigma \text{ in if } X_1 \subseteq X_2 \text{ then true else false} \\
\llbracket \Gamma \vdash \forall x : K \in G. P \rrbracket_{\theta} \sigma &= \zeta \text{ if } \llbracket G \rrbracket_{\theta} \sigma = \zeta \text{ or } \llbracket \Gamma, x : K \vdash P \rrbracket_{\theta} \text{Extend}(\sigma, x, o) = \zeta \\
&\quad \text{for some } o \text{ in } (\llbracket G \rrbracket_{\theta} \sigma) \setminus \{\text{null}\} \text{ with } \text{Type}(o, \sigma) = K \\
&= \text{true if } \llbracket \Gamma, x : K \vdash P \rrbracket_{\theta} \text{Extend}(\sigma, x, o) = \text{true} \\
&\quad \text{for all } o \text{ in } (\llbracket G \rrbracket_{\theta} \sigma) \setminus \{\text{null}\} \text{ with } \text{Type}(o, \sigma) = K \\
&= \text{false otherwise}
\end{aligned}$$

**Fig. 6.** Formulas: three-valued semantics,  $\llbracket \Gamma \vdash P \rrbracket_{\theta} \sigma \in \{\text{true}, \text{false}, \zeta\}$ . Typing context is elided in most cases.

## 4 Semantics of effects and programs

*Effects.* A **location** is either a variable name  $x$  or a heap location comprised of a reference  $o$  and field name  $f$ . We write  $o.f$  for such pairs. Define  $\text{rlocs}(\sigma, \theta, \varepsilon)$ , the locations designated by read effects of  $\varepsilon$ , in  $\sigma$ , by  $\text{rlocs}(\sigma, \theta, \varepsilon) = \{x \mid \varepsilon \text{ contains rd } x\} \cup \{o.f \mid \varepsilon \text{ contains rd } G^{\cdot} f \text{ with } o \in \llbracket G \rrbracket_{\theta} \sigma\}$ . Define  $\text{wlocs}$  similarly, for write effects.

Write effects constrain what locations are allowed to change between one state and another. We say  $\varepsilon$  **allows change from  $\sigma$  to  $\tau$  under  $\theta$** , written  $\sigma \rightarrow \tau \models_{\theta} \varepsilon$ , provided (a) if  $y$  changed value (i.e.,  $\tau(y) \neq \sigma(y)$ ) then  $\text{wr } y$  is in  $\varepsilon$ ; (b) if  $o.f$  changed value then there is  $\text{wr } G^{\cdot} f$  in  $\varepsilon$  such that  $o \in \llbracket G \rrbracket_{\theta} \sigma$ ; and (c) if  $\text{fr } G$  is in  $\varepsilon$  then elements of  $G$  in  $\tau$  are fresh. Reads are ignored, so  $\sigma \rightarrow \tau \models_{\theta} \varepsilon$  iff  $\sigma \rightarrow \tau \models_{\theta} \text{writes}(\varepsilon)$ . In (b), region expressions  $G$  are interpreted in the initial state because frame conditions need only report writes to fields of pre-existing objects and not freshly allocated objects.

Read effects constrain what locations an outcome can depend on. Dependency is expressed by considering two initial states that agree on the locations deemed readable. Agreement needs to take into account variation in allocation, as two states may have isomorphic pointer structure but differently chosen references.

Let  $\pi$  range over **partial bijections** on  $\text{Ref}$ . We write  $\pi(p) = p'$  to express that  $\pi$  is defined on  $p$  and has value  $p'$ . A **refperm** from  $\sigma$  to  $\sigma'$  is partial bijection  $\pi$  such that

- $\text{dom}(\pi) \subseteq \sigma(\text{alloc}) \cup \{\text{null}\}$  and  $\text{rng}(\pi) \subseteq \sigma'(\text{alloc}) \cup \{\text{null}\}$
- $\pi(\text{null}) = \text{null}$
- $\pi(p) = p'$  implies  $\text{Type}(p, \sigma) = \text{Type}(p', \sigma')$  for all proper references  $p, p'$

$$\begin{aligned}
\sigma \models_{\theta} E_1 = E_2 & \quad \text{iff } \llbracket E_1 \rrbracket_{\theta} \sigma = \llbracket E_2 \rrbracket_{\theta} \sigma \\
\sigma \models_{\theta} x.f = E & \quad \text{iff } \sigma(x) \neq \text{null} \text{ and } \sigma(x.f) = \llbracket E \rrbracket_{\theta} \sigma \\
\sigma \models_{\theta} G_1 \subseteq G_2 & \quad \text{iff } \llbracket G_1 \rrbracket_{\theta} \sigma \subseteq \llbracket G_2 \rrbracket_{\theta} \sigma \\
\sigma \models_{\theta}^{\Gamma} \forall x : K \in G. P & \quad \text{iff } \text{Extend}(\sigma, x, o) \models_{\theta}^{\Gamma, x:K} P \\
& \quad \text{for all } o \text{ in } (\llbracket G \rrbracket_{\theta} \sigma) \setminus \{\text{null}\} \text{ with } \text{Type}(o, \sigma) = K \\
\sigma \models_{\theta} P_1 \wedge P_2 & \quad \text{iff } \sigma \models_{\theta} P_1 \text{ and } \sigma \models_{\theta} P_2 \\
\sigma \models_{\theta} \neg P & \quad \text{iff } \sigma \not\models_{\theta} P
\end{aligned}$$

**Fig. 7.** Two-valued semantics. These clauses hold when  $\sigma \models_{\theta} df(P, \Delta)$  (Lemma 6).

Define  $p \overset{\pi}{\sim} p'$  to mean  $\pi(p) = p'$ . We extend  $\overset{\pi}{\sim}$  to a relation on integers by  $i \overset{\pi}{\sim} j$  iff  $i = j$ . For reference sets  $X, Y$  we define  $X \overset{\pi}{\sim} Y$  iff  $\pi(X) \supseteq Y$  and  $X \subseteq \pi^{-1}(Y)$  (where  $\pi(X)$  is the direct image of  $X$ ). That is,  $\pi$  forms a bijection between  $X$  and  $Y$ .

Define  $\text{freshLocs}(\sigma, \tau) = \{p.f \mid p \in \text{freshRefs}(\sigma, \tau) \wedge f \in \text{Fields}(\text{Type}(p, \tau))\}$  where  $\text{freshRefs}(\sigma, \tau) = \tau(\text{alloc}) \setminus \sigma(\text{alloc})$ . For a set  $W$  of variables and heap locations, define  $\text{Lagree}(\sigma, \sigma', W, \pi)$  iff  $\forall x \in W. \sigma(x) \overset{\pi}{\sim} \sigma'(x)$  and  $\forall (o.f) \in W. o \in \text{dom}(\pi) \wedge \sigma(o.f) \overset{\pi}{\sim} \sigma'(\pi(o).f)$ .

**Definition 1 (agreement on read effects)** Let  $\varepsilon$  be an effect that is swf in  $\Gamma$ . Consider states  $\sigma, \sigma'$ . Let  $\pi$  be a partial bijection. Let  $\theta$  be a candidate interpretation (for some  $\Delta$  that is swf in  $\Gamma$ ). Say that  $\sigma$  and  $\sigma'$  **agree on  $\varepsilon$  modulo  $\pi$** , written  $\text{Agree}(\sigma, \sigma', \varepsilon, \pi, \theta)$ , iff  $\text{Lagree}(\sigma, \sigma', \text{rlocs}(\sigma, \theta, \varepsilon), \pi)$ . Define  $\text{Agree}(\sigma, \sigma', \varepsilon, \theta) = \text{Agree}(\sigma, \sigma', \varepsilon, \pi, \theta)$  where  $\pi$  is the identity on  $\sigma(\text{alloc}) \cap \sigma'(\text{alloc})$ .

*Programs.* In the following we consider a method context  $\Delta$  that is well-formed in some typing context  $\Gamma$  (often elided). For substitution we use the notation  $P_e^x$ . For clarity we use substitution notation in satisfaction statements, even though strictly speaking the syntax does not (and should not) include reference literals. If  $\Gamma, x : T \vdash P$  and  $\sigma \in \llbracket \Gamma \rrbracket$  and  $v \in \llbracket T \rrbracket \sigma$ , we may write  $\sigma \models_{\theta}^{\Gamma} P_v^x$  to abbreviate  $\text{Extend}(\sigma, x, v) \models_{\theta}^{\Gamma, x:T} P$ .

The transition relation depends on a method context  $\Delta$ . Configurations take the form  $\langle C, \sigma, \mu \rangle$  where  $\mu$  is a method environment. The call of a let-bound method  $m$  executes the body  $\mu(m)$  with variables renamed to avoid clashes with the calling context. In case of a pure method the call takes the form  $y := m(F)$  and there is some extra bookkeeping to assign the final value of  $\text{res}$  (or rather, a fresh instance thereof) to  $y$ .

The transition semantics for pure method call  $y := m(F)$  takes a step that assigns to  $y$  the value  $\llbracket m(F) \rrbracket_{\theta} \sigma$  (defined in Fig. 5). The transition semantics of a call  $m(z)$ , for impure  $m$  in  $\Delta$ , takes a single step to a final state (or  $\downarrow$ ) that satisfies the specification  $\Delta(m)$ . Such states are described by the denotation  $\llbracket \Delta(m) \rrbracket$  of the specification.

**Definition 2 (Denotation of impure method spec)** Let  $\Delta$  be swf and let  $(x:T)R \rightsquigarrow S[\eta]$  be in  $\Delta$ . Let  $\theta$  be a candidate interpretation of  $\Delta$  and  $z$  a variable name. Then  $\llbracket (x:T)R \rightsquigarrow S[\eta] \rrbracket(\theta, z)$  is defined as follows, for any  $\Gamma_1 \supseteq \Gamma$  and  $\Gamma_1$ -states  $\sigma, \tau$ :

- (i)  $\llbracket (x:T)R \rightsquigarrow S[\eta] \rrbracket(\theta, z) \sigma \not\downarrow$  iff  $\sigma \not\models_{\theta} R_z^x$
- (ii)  $\llbracket (x:T)R \rightsquigarrow S[\eta] \rrbracket(\theta, z) \sigma \tau$  iff
  - (a)  $\sigma \models_{\theta}^{\Gamma_1} R_z^x$  and  $\tau \models_{\theta}^{\Gamma_1} S_z^x$  and  $\sigma \rightarrow \tau \models_{\theta} \eta_z^x$  and

$$\begin{array}{c}
\frac{\llbracket \Delta(m) \rrbracket(\theta, z) \sigma \tau}{\langle m(z), \sigma, \mu \rangle \xrightarrow{\Delta, \theta} \langle \text{skip}, \tau, \mu \rangle} \qquad \frac{\llbracket \Delta(m) \rrbracket(\theta, z) \sigma \not\downarrow}{\langle m(z), \sigma, \mu \rangle \xrightarrow{\Delta, \theta} \not\downarrow} \\
\frac{\llbracket m(F) \rrbracket_{\theta} \sigma \neq \not\downarrow \quad \tau = [\sigma \mid x : \llbracket m(F) \rrbracket_{\theta} \sigma]}{\langle x := m(F), \sigma, \mu \rangle \xrightarrow{\Delta, \theta} \langle \text{skip}, \tau, \mu \rangle} \qquad \frac{\llbracket m(F) \rrbracket_{\theta} \sigma = \not\downarrow}{\langle x := m(F), \sigma, \mu \rangle \xrightarrow{\Delta, \theta} \not\downarrow}
\end{array}$$

**Fig. 8.** Transition rules for calls of impure and pure procedures in context  $\Delta$ .

- (b) for all  $\sigma', \pi$ , if  $\text{Agree}(\sigma, \sigma', \eta_z^x, \pi, \theta)$  and  $\sigma' \models_{\theta}^{\Gamma_1} R_z^x$  then there are  $\tau', \rho$  with
- $\tau' \models_{\theta}^{\Gamma_1} S_z^x$  and  $\sigma' \rightarrow \tau' \models_{\theta} \eta_z^x$
  - $\rho \supseteq \pi$  and  $\text{freshRefs}(\sigma', \tau') \subseteq \rho(\text{freshRefs}(\sigma, \tau))$
  - $\text{Lagree}(\tau, \tau', X, \rho)$  where  $X = \text{freshLocs}(\sigma, \tau) \cup \text{wlocs}(\sigma, \theta, \eta_z^x)$

It is item (ii)(b) that is new in this paper; the rest is from RLII. Note that  $X$  is defined by interpreting  $\eta$  in the initial state.

A state  $\sigma$  may have no successor because the specification is unsatisfiable at  $\sigma$ . Unsatisfiability may be due to the postcondition, but it can also happen that  $\tau$  satisfies the postcondition but not the read effect. The specification  $(x:\text{Cell}) \text{true} \rightsquigarrow y = x.\text{val} [\text{wr } y]$  is unsatisfiable:  $y$  cannot be set to  $x.\text{val}$  without reading  $\{x\}^{\text{val}}$  or having a stronger precondition like  $y = x.\text{val}$ .

Although specifications include read effects—a relational property—the denotation of a specification need not be defined as an extreme solution to constraints including that relational property. The elementary definition above has the property that any  $\tau'$  that satisfies the conditions in (ii) is a possible successor of  $\sigma'$ , i.e., the denotation is closed in the sense that it includes the pair  $\sigma', \tau'$ . This is made precise in Thm. 9. The condition  $\text{freshRefs}(\sigma', \tau') \subseteq \rho(\text{freshRefs}(\sigma, \tau))$  in (ii)(b) was not immediately obvious but is crucial for Thm. 9.

With all the ingredients in hand, the transition semantics can be defined; see Fig. 8.

## 5 Semantics of correctness judgments

To link a client  $C$  with implementation  $B$  of a method  $m$  used by  $C$  we want  $C$  to be correct for all interpretations of the method context. But reasoning about  $B$  can use a particular interpretation for  $m$ . Such an interpretation might be provided directly, as a mathematical definition provided by the programmer, or it might be derived from the code as it is in work on VC generation for pure methods [5]. Here we treat such interpretations semantically. To that end, we generalize the correctness judgment form to  $\Delta; \theta \vdash^{\Gamma} C : P \rightsquigarrow Q [\varepsilon]$ . For this to be swf,  $\theta$  should be a candidate interpretation of some subset of  $\Delta$ , and  $\Delta \vdash^{\Gamma} C : P \rightsquigarrow Q [\varepsilon]$  should be swf as defined in Sec. 2. The original form is essentially the special case where  $\theta$  is the empty function. The generalized correctness judgment is important for the linking rule, which we introduce here in abridged form. We consider a single method specification  $\Theta \equiv m : (x:T, \text{res}:U) Q \rightsquigarrow Q'$ , we elide effects, and the partial interpretation of the

ambient library  $\Delta$  is empty.

$$\frac{\Delta, \Theta; \emptyset \vdash C : P \rightsquigarrow P' \quad \Delta, \Theta; \theta \vdash B : Q \rightsquigarrow Q' \quad \text{dom } \theta = \{m\} \quad \theta \models \Delta, \Theta}{\Delta; \emptyset \vdash \text{let } m(x:T, \text{res}:U) = B \text{ in } C : P \rightsquigarrow P'} \quad (1)$$

A client  $C$  is linked with the implementation  $B$  of a pure procedure  $m$ . The verification condition for  $C$  is under the hypothesis of some specifications  $\Delta, \Theta$  which include the specification  $\Theta$  of  $m$ . The rule may only be instantiated with swf judgments, so  $\Delta$  is swf (as it appears in the conclusion) and the larger method context  $\Delta, \Theta$  is also swf.

According to the semantics to follow, the judgment for  $C$  means that it is correct with respect to any interpretation  $\varphi$  of all the pure procedures in  $\Delta, \Theta$ . The verification condition for  $B$  also has hypothesis  $\Delta, \Theta$  for procedures that may be called in  $B$  or used in its specification, and  $B$  must be correct with respect to any interpretation of the pure procedures in  $\Delta$ , but fixed interpretation  $\theta$  of  $m$ . The rule requires that in fact  $\theta$  is an interpretation of  $\Theta$ , meaning that  $\theta(m)$  satisfies the specification of  $m$ . Because this specification may refer to pure methods in the ambient context  $\Delta$ , satisfaction is expressed as  $\theta \models \Delta, \Theta$ . This is defined in terms of the following.

**Definition 3 (context interpretation)** Let  $\Delta$  be swf in  $\Gamma$  and let  $\theta$  be a candidate  $\Delta$ -interpretation. (Note that  $\text{dom } \theta = \text{dom } \Delta$ .) Say  $\theta$  is a  $\Delta$ -*interpretation* iff the following holds for each  $m : (x:T, \text{res}:U)P \rightsquigarrow Q [\varepsilon]$  in  $\Delta$ . For any  $\sigma \in \llbracket \Gamma \rrbracket$  and  $v \in \llbracket T \rrbracket \sigma$ ,

$$(a) \quad \theta(m)(\sigma, v) = \perp \text{ iff } \sigma \not\models_{\theta} P_v^x$$

Furthermore, if  $\sigma \models_{\theta} P_v^x$  then letting  $w = \theta(m)(\sigma, v)$  we have

$$(b) \quad \sigma \models_{\theta} Q_{v,w}^{x,\text{res}}$$

$$(c) \quad \text{for any } \sigma' \in \llbracket \Gamma \rrbracket, v' \in \llbracket T \rrbracket \sigma' \text{ with } \sigma' \models_{\theta} P_{v'}^x, \text{ and any refperm } \pi \text{ from } \sigma \text{ to } \sigma', \\ \text{if } v \stackrel{\pi}{\sim} v' \text{ and } \text{Agree}(\sigma, \sigma', \varepsilon, \pi, \theta) \text{ then } w \stackrel{\pi}{\sim} w' \text{ where } w' = \theta(m)(\sigma', v')$$

**Definition 4 (partial context interpretation)** Let  $\Delta$  be swf and  $\Delta, \Theta$  be swf. Let  $\theta$  be a candidate interpretation of  $\Theta$ . We say  $\theta$  is a *partial interpretation* of  $\Delta, \Theta$ , written  $\theta \models \Delta, \Theta$ , provided that for any  $\Delta$ -interpretation  $\delta$ , the candidate  $\delta \cup \theta$  is a  $(\Delta, \Theta)$ -interpretation.<sup>5</sup>

**Definition 5 (valid judgment)** A swf correctness judgment  $\Delta; \theta \vdash^{\Gamma} C : P \rightsquigarrow Q [\varepsilon]$  is *valid* iff the following conditions hold for all  $\Gamma$ -environments  $\mu$ , all  $\Delta$ -interpretations  $\delta$  such that  $\theta \subseteq \delta$ , and all states  $\sigma$  such that  $\sigma \models_{\delta}^{\Gamma, \text{sig}(\Delta)} P$ .

$$\text{(Safety)} \quad \text{It is not the case that } \langle C, \sigma, \mu \rangle \xrightarrow{\Delta, \delta}^* \perp.$$

$$\text{(Post)} \quad \tau \models_{\delta} Q \text{ for every } \tau \text{ with } \langle C, \sigma, \mu \rangle \xrightarrow{\Delta, \delta}^* \langle \text{skip}, \tau, \mu \rangle$$

$$\text{(Effect)} \quad \sigma \rightarrow \tau \models_{\delta} \varepsilon \text{ for every } \tau \text{ with } \langle C, \sigma, \mu \rangle \xrightarrow{\Delta, \delta}^* \langle \text{skip}, \tau, \mu \rangle$$

<sup>5</sup> Under these conditions, if the specifications in  $\Theta$  refer to methods in  $\Delta$ ,  $\theta$  is not swf on its own, and then it is not meaningful to call  $\theta$  a  $\Theta$ -interpretation.

(Read Effect) for any  $\tau$  such that  $\langle C, \sigma, \mu \rangle \xrightarrow{\Delta, \delta}^* \langle \text{skip}, \tau, \mu \rangle$ , and any  $\sigma', \pi, \tau$  such that  $\text{Agree}(\sigma, \sigma', \varepsilon, \pi, \delta)$  and  $\sigma' \models_{\delta}^{\Gamma} P$ , there are  $\tau', \rho$  such that  $\langle C, \sigma', \mu \rangle \xrightarrow{\Delta, \delta}^* \langle \text{skip}, \tau', \mu \rangle$  and  $\rho \supseteq \pi$  and  $\text{freshLocs}(\sigma', \tau') \subseteq \rho(\text{freshLocs}(\sigma, \tau))$  and  $\text{Lagree}(\tau, \tau', X, \rho)$  where  $X = \text{freshLocs}(\sigma, \tau) \cup \text{wlocs}(\sigma, \delta, \varepsilon)$ .

In case  $\Delta(m)$  is unsatisfiable (except possibly by divergence), no  $\Delta$ -interpretation exists. Then the judgment holds but the hypotheses cannot be discharged by linking because there is no way to instantiate  $\theta$  in rule (1).

The definitions up to this point apply even if pure methods are called outside their precondition. For understandable proof rules, and to stay within FOL for assertions, we will disallow such specifications and correctness judgments.

**Lemma 6 (two-valued semantics of formulas)** If  $\theta$  is a  $\Delta$ -interpretation and  $\sigma \models_{\theta} df(P, \Delta)$  then  $\llbracket P \rrbracket_{\theta} \sigma$  is not  $\perp$ . And for any  $\sigma$  and any  $\Delta$ -interpretation  $\theta$ , if  $\sigma \models_{\theta} df(P, \Delta)$  then  $\sigma \models_{\theta} P$  satisfies the usual defining clause, see Fig. 7.

**Definition 7** Let  $\Gamma$  be a typing context and let  $\Delta$  be a specification context that is swf in  $\Gamma$ . Let  $P$  be a formula that is swf in  $\text{vars}(\Gamma), \text{sigs}(\Delta)$ . Then  $P$  is  $\Delta$ -**valid**, written  $\Delta \models P$ , if and only if  $\sigma \models_{\theta} P$  for all  $\Delta$ -interpretations  $\theta$  and all states  $\sigma$ .

**Definition 8 (healthy, well-formed)** Let  $\Gamma$  and  $\Delta$  satisfy the conditions of Def. 7. A formula  $P$  that is swf is **healthy** iff  $df(P, \Delta)$  is valid. A swf specification  $P \rightsquigarrow Q[\eta]$  is **healthy** (with respect to  $\Gamma, \Delta$ ) iff the three formulas  $df(P, \Delta)$ ,  $P \Rightarrow df(Q, \Delta)$ , and  $P \Rightarrow df(\eta, \Delta)$  are  $\Delta$ -valid. A swf correctness judgment  $\Delta; \theta \vdash^{\Gamma} C : P \rightsquigarrow Q[\eta]$  is **healthy** iff the three formulas  $df(P, \Delta)$ ,  $P \Rightarrow df(Q, \Delta)$ , and  $P \Rightarrow df(\eta, \Delta)$  are  $\Delta$ -valid. The term **well-formed** means swf and healthy.

The definitions to this point are intricate but elementary; in particular, there are no fixpoints. But by contrast with axiomatic semantics, correctness is directly grounded in a conventional operational semantics. The one unconventional element is that transition semantics depends on method context. The ultimate confirmation that we *are* reasoning about program behavior is soundness of the linking rule, which can be used to discharge all hypotheses.

## 6 Proof system

The **framing judgment** has the form  $P; \Delta \vdash \eta \text{ frm } Q$  and is swf under evident conditions. It means that in  $P$ -states,  $Q$  reads within the read effect  $\eta$ . The judgment is **healthy** iff the formulas  $df(P, \Delta)$ ,  $P \Rightarrow df(\eta, \Delta)$ , and  $P \Rightarrow df(Q, \Delta)$  are all  $\Delta$ -valid. The judgment is **valid**, written  $P; \Delta \models^{\Gamma} \eta \text{ frm } Q$ , iff for all  $\Gamma$ -states  $\sigma, \sigma'$ , reperms  $\pi$ , and  $\Delta$ -interpretations  $\theta$ , if  $\text{Agree}(\sigma, \sigma', \eta, \pi, \theta)$ , and  $\sigma \models_{\theta}^{\Gamma} P \wedge Q$ , then  $\sigma' \models_{\theta}^{\Gamma} Q$ .

A verifier can check framing judgments in terms of the validity property, but our logic includes rules to derive framing judgments. A basic rule allows to infer, for atomic formula  $P$ , the judgment  $\text{true}; \Delta \vdash \text{ftpt}(P, \Delta) \text{ frm } P$  concerning a precise footprint computed by function  $\text{ftpt}$  which is defined in Fig. 9. For non-atomic formulas there are

$$\begin{array}{ll}
\text{ftpt}(x, \Delta) & = \text{rd } x & \text{ftpt}(E = E', \Delta) & = \text{ftpt}(E, \Delta), \text{ftpt}(E', \Delta) \\
\text{ftpt}(G'f, \Delta) & = \text{rd } G'f, \text{ftpt}(G, \Delta) & \text{ftpt}(G_1 \subseteq G_2, \Delta) & = \text{ftpt}(G_1, \Delta), \text{ftpt}(G_2, \Delta) \\
\text{ftpt}(\emptyset, \Delta) & = \emptyset & \text{ftpt}(x.f = F, \Delta) & = \text{rd } x, x.f, \text{ftpt}(F, \Delta) \\
\text{ftpt}(m(F), \Delta) & = \text{reads}(\varepsilon_F^x), \text{ftpt}(F, \Delta) & \text{for } \Delta(m) = (x : T, \text{res} : U)P \rightsquigarrow Q[\varepsilon]
\end{array}$$

**Fig. 9.** Footprints of region expressions and atomic assertions well-formed in  $\Delta$ .

syntax-directed rules, e.g., the rule for conjunction allows to infer  $P; \Delta \vdash \varepsilon$  from  $Q_1 \wedge Q_2$  from  $P; \Delta \vdash \varepsilon$  from  $Q_1$  and  $P \wedge Q_1; \Delta \vdash \varepsilon$  from  $Q_2$ . There are also subsidiary rules for subsumption of effects and for logical manipulation of  $P$ . These rules are adapted in a straightforward way from RLI (Sec. 6.1).

The point of establishing  $P; \Delta \vdash \eta$  from  $Q$  is that code that writes outside  $\eta$  cannot falsify  $Q$ . This is expressed in the frame rule by computing, from the frame  $\eta$  of  $Q$  and the frame condition  $\varepsilon$  of the code, a *separator formula* which is a conjunction of region disjointness formulas describing states in which writes allowed by  $\varepsilon$  cannot affect the value of a formula with read effect  $\eta$ . We define the separator formula as  $\eta \cdot \varepsilon$ , using function  $\cdot$  which recurses on syntax (see RLI Sec. 6.2). For example,  $\text{rd } G'f \cdot \text{wr } H'g$  is *true*, and  $\text{rd } G'f \cdot \text{wr } H'f$  is the disjointness formula  $G \# H$ . Also,  $\text{rd } x \cdot \text{wr } y$  is simply *false*, if  $x$  and  $y$  are the same variable, and *true* otherwise. Writes on the left and reads on the right are ignored, so  $\eta \cdot \varepsilon$  is the same as  $\text{reads}(\eta) \cdot \text{writes}(\varepsilon)$ .

The key property of a separator is to establish the agreement to which frame validity refers. To be precise, suppose  $\sigma \rightarrow \tau \models_{\theta} \varepsilon$  and  $\sigma \models_{\theta} \eta \cdot \varepsilon$ . Then  $\text{Agree}(\sigma, \tau, \eta, id, \theta)$ , where  $id$  is the identity on  $\sigma(\text{alloc})$ .

An effect  $\varepsilon$  is called *self-framing* in method context  $\Delta$  provided that for every  $\text{rd } G'f$  or  $\text{wr } G'f$  in  $\varepsilon$ ,  $\text{ftpt}(G, \Delta)$  is in  $\varepsilon$ . Such effects arise, for example, in case a method refers to itself in its frame condition. So are effects obtained using the *ftpt* function and most of the framing rules. For self-framing  $\varepsilon$ , if  $\text{Agree}(\sigma, \sigma', \varepsilon, \pi, \theta)$  then  $\llbracket G \rrbracket_{\theta} \sigma \stackrel{\pi}{\sim} \llbracket G \rrbracket_{\theta} \sigma'$  for any  $\text{rd } G'f$  or  $\text{wr } G'f$  in  $\varepsilon$ .

**Theorem 9 (denotation closure).** Suppose  $\eta$  is self-framing and  $\theta$  is a  $\Delta$ -interpretation. If  $\llbracket (x:T)R \rightsquigarrow S[\eta] \rrbracket(\theta, z)\sigma\tau$  then  $\llbracket (x:T)R \rightsquigarrow S[\eta] \rrbracket(\theta, z)\sigma'\tau'$ , provided that  $\sigma', \tau'$  satisfy the conditions in Def. 2(ii).

*Proof rules for correctness judgments.* Fig. 10 presents a few proof rules. They are to be instantiated only with well-formed premises and conclusions (Def. 8). The first, for field update, is a ‘local axiom’ that precisely describes the effect; it shows how read effects can easily be incorporated into the rules from RLI (Sec. 7.1) and RII (Sec. 7.1). Next is the frame rule, adapted from RLI/II by adding  $\Delta$  to the side conditions. Then come rules for impure and pure method calls. For reasons of parsimony we make self-framing an explicit premise where needed for soundness. It turns out that in non-trivial provable judgements the frame conditions in the method context will be self-framing.

We give an illustrative rule for linking a client with a method implementation. The general rule allows several pure and impure methods that may refer to each other in their specifications and code (of course, subject to the proviso concerning  $\prec_{\Delta}^+$  in the definition of swf method context in Sec. 2). Predicate  $\text{terminates}(Q, B)$  says that from any  $Q$ -state,  $B$  terminates (normally or abnormally). One premise is that partial  $(\Delta, \Theta)$ -

$$\begin{array}{c}
\Delta; \emptyset \vdash x.f := F : x \neq \text{null} \wedge y = F \rightsquigarrow x.f = y [\text{wr } x.f, \text{rd } x, \text{ftpt}(F, \Delta)] \\
\text{FRAME} \frac{\Delta; \theta \vdash C : P \rightsquigarrow Q [\varepsilon] \quad P; \Delta \vdash \eta \text{ frm } R \quad \Delta; \theta \models P \wedge R \Rightarrow \eta \cdot I. \varepsilon}{\Delta; \theta \vdash C : P \wedge R \rightsquigarrow Q \wedge R [\varepsilon]} \\
\\
\frac{\Delta; \theta \vdash C : P \rightsquigarrow Q [\varepsilon]}{\Delta; \theta \cup \theta' \vdash C : P \rightsquigarrow Q [\varepsilon]} \quad \frac{\varepsilon \text{ is self-framing}}{m : (x:T)P \rightsquigarrow Q [\varepsilon]; \emptyset \vdash m(z) : P_z^x \rightsquigarrow Q_z^x [\varepsilon_z^x]} \\
\\
\frac{x \notin \text{Vars}(H) \cup \text{FV}(Q)}{m : (y:T, \text{res}:U)P \rightsquigarrow Q [\varepsilon]; \emptyset \vdash x := m(H) : P_H^y \rightsquigarrow Q_{H,x}^{y, \text{res}} [\text{wr } x, \text{ftpt}(H, \Delta), \varepsilon_H^y]} \\
\\
\frac{\begin{array}{c} \Theta \text{ is } m : (x:T, \text{res}:U)Q \rightsquigarrow Q' [\eta] \quad \text{dom } \theta = \text{dom } \Theta \quad \theta \models \Delta, \Theta \\ \Delta, \Theta; \delta \vdash^F C : P \rightsquigarrow P' [\varepsilon] \quad \Delta, \Theta; \delta \cup \theta \vdash^F, x:T, \text{res}:U B : Q \rightsquigarrow Q' [\text{wr } \text{res}, \text{rd } x, \eta] \\ \eta \text{ is wr-free and self-framing} \quad \text{terminates}(Q, B) \end{array}}{\Delta; \delta \vdash^F \text{let } m(x:T, \text{res}:U) = B \text{ in } C : P \rightsquigarrow P' [\varepsilon]}
\end{array}$$

**Fig. 10.** Proof rules for field update, framing, interpreting, pure/impure calls, and linking.

interpretation  $\theta$  is provided; it gives the *chosen* interpretation for  $m$ , to be used in verifying the body  $B$ . By contrast, the premise for  $C$  requires correctness with respect to *all* interpretations of  $m$ .

**Theorem 10.** Any derivable correctness judgment is valid.

## 7 Related work

We take the Cell example from the most closely related work, [18], where read effects of pure methods are specified using dynamic frames and methods may be self-framing. They define (and implement) a VC-generator including VCs that encode the semantics of read effects, albeit only for a pair of states in succession. (That avoids the need for *reperms*, and suffices for framing but not relational reasoning for data abstraction and encapsulation.) They give a detailed proof of soundness with respect to transition semantics, by showing that the VCs ensure a small-step invariant that implies correctness and fault-avoidance. Axioms are included (and proved sound) to exploit read effects for framing. Different from our work, the body of a pure method is required to be a single ‘**return E**’ statement and  $E$  is visible to clients; and pure methods do not have postconditions. (Their implementation does include such postconditions.) Although VCs are generated modularly, we do not discern an explicit account of linking, or an easy adaptation to cater for hiding a pure method body or invariants from clients. As usual in practical systems, the syntax embeds specifications in programs, as opposed to judgments that ascribe properties to programs.

A number of earlier works point out the importance of read effects for pure methods and explore VC-generation, e.g. [5], explore weak purity which allows allocation, and shows consistency of a system of VCs (but not operational soundness). The analog of consistency, in our setting, is being able to discharge hypotheses in the linking rule.

Framing in separation logic encompasses read and write effects, implicitly in syntax but explicitly in the semantics (safety monotonicity, frame property [14]). Whereas self-framing is a property of effects, in our setting, it is a property of formulas in other settings. In separation logic, all assertions are effectively self-framing. The abstract predicates approach [15] to data abstraction has inspired several works that cater for SMT provers by using ghost instrumentation to encode intensional semantics of effects in terms of permissions. One provides a VC generator and sketches an argument for its operational soundness [6]. Another gives a detailed semantics and soundness proof for VCs that provide effective reasoning about recursively defined abstract predicates and abstraction functions [19]. The latter works have extensive pointers to related work.

## References

1. A. Banerjee and D. A. Naumann. Local reasoning for global invariants, part II: Dynamic boundaries. *Journal of the ACM*, 60(3):19:1–19:73, 2013.
2. A. Banerjee, D. A. Naumann, and S. Rosenberg. Local reasoning for global invariants, part I: Region logic. *Journal of the ACM*, 60(3):18:1–18:56, 2013.
3. N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *POPL*, 2004.
4. F. Bobot and J.-C. Filliâtre. Separation predicates: A taste of separation logic in first-order logic. In *ICFEM*, 2012.
5. A. Darvas and P. Müller. Reasoning about method calls in interface specifications. *Journal of Object Technology (JOT)*, June 2006.
6. S. Heule, I. T. Kassios, P. Müller, and A. J. Summers. Verification condition generation for permission logics with abstract predicates and abstraction functions. In *ECOOP*, 2013.
7. C. A. R. Hoare. Proofs of correctness of data representations. *Acta Inf.*, 1:271–281, 1972.
8. I. T. Kassios. The dynamic frames theory. *Formal Aspects of Comp.*, 23(3):267–288, 2011.
9. N. R. Krishnaswami, J. Aldrich, and L. Birkedal. Verifying event-driven programs using ramified frame properties. In *TLDI*, 2010.
10. K. R. M. Leino and P. Müller. Verification of equivalent-results methods. In *ESOP*, 2008.
11. A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *ESOP*, 2007.
12. A. Nanevski, A. Banerjee, and D. Garg. Dependent type theory for verification of information flow and access control policies. *ACM Trans. Program. Lang. Syst.*, 35(2):6, 2013.
13. D. A. Naumann. Observational purity and encapsulation. *Theoretical Comput. Sci.*, 376(3):205–224, 2007.
14. P. W. O’Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. *ACM Trans. Prog. Lang. Syst.*, 31(3):1–50, 2009.
15. M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *POPL*, 2005.
16. S. Rosenberg, A. Banerjee, and D. A. Naumann. Local reasoning and dynamic framing for the composite pattern and its clients. In *VSTTE*, 2010.
17. S. Rosenberg, A. Banerjee, and D. A. Naumann. Decision procedures for region logic. In *VMCAI*, 2012.
18. J. Smans, B. Jacobs, F. Piessens, and W. Schulte. Automatic verification of Java programs with dynamic frames. *Formal Aspects of Comp.*, 22(3-4):423–457, 2010.
19. A. J. Summers and S. Drossopoulou. A formal semantics for isorecursive and equirecursive state abstractions. In *ECOOP*, 2013.
20. H. Yang. Relational separation logic. *Theoretical Comput. Sci.*, 375(1-3):308–334, 2007.