

A Simple Model of Separation Logic for Higher-order Store

Lars Birkedal¹, Bernhard Reus², Jan Schwinghammer³, and Hongseok Yang⁴

¹ IT University of Copenhagen ² University of Sussex, Brighton

³ Saarland University, Saarbrücken ⁴ Queen Mary, University of London

Abstract. Separation logic is a Hoare-style logic for reasoning about pointer-manipulating programs. Its core ideas have recently been extended from low-level to richer, high-level languages. In this paper we develop a new semantics of the logic for a programming language where code can be stored (i.e., with higher-order store). The main improvement on previous work is the simplicity of the model. As a consequence, several restrictions imposed by the semantics are removed, leading to a considerably more natural assertion language with a powerful specification logic.

1 Introduction

Higher-order store is included in modern programming languages in the form of code pointers and storable objects. “Higher-order” here refers to the fact that one can keep not only data in the store but also procedures or commands that manipulate the store themselves. It is widely used in systems code, such as operating system kernels, device drivers and web servers. For instance, the Linux kernel keeps multiple linked lists whose nodes store code fragments, and calls those fragments in response to external events, such as a signal from a printer.

However, formal reasoning about higher-order store is still an open problem. Although several sound program logics for higher-order store have been proposed, they either are intended for machine code [4] or they fail to combine local reasoning with intuitive rules for stored code while maintaining the simplicity of Hoare logic for first-order store [6, 15]. The difficulty is that a logic for higher-order store should accommodate reasoning about “recursion through the store”, a tricky implicit recursion implemented by stored procedures.

The goal of our research is to solve the problem of reasoning about higher-order store using separation logic. Separation logic is a program logic for reasoning modularly about programs with pointers. It has been demonstrated that the logic substantially simplifies formal program verification in low-level C-like programming languages as well as richer, higher-level languages [17, 2, 1, 12, 9, 3, 8, 7, 13]. Our aim is to design program logics for higher-order store that keep all the benefits of separation logic, such as (higher-order) frame rules, while providing efficient, sound proof rules for recursion through the store.

In this paper, we investigate the semantic foundations for developing separation logic for higher-order store. We build on previous work of Reus and

Schwinghammer [15], which identified key semantic challenges for such a logic, and provided fairly sophisticated solutions based on functor categories. In this paper, we take different approaches to the various problems, and as a result obtain a more powerful logic and a substantially simpler semantic model.

We now give an overview of two key semantic challenges that are involved in developing separation logic for higher-order store. We outline how those challenges were addressed in earlier work [15], and compare with our new model.

The first challenge is to find a model that validates the frame rule known from separation logic [17]. In traditional models of separation logic [10], the soundness of the frame rule relies on programs satisfying a frame property, which says that the meaning of each program phrase only relies on its “footprint”. To ensure that all program phrases – in particular, memory allocation – satisfy the frame property, the models interpret commands as relations (i.e., functions from input states to *sets* of output states), and memory allocation denotes a function that nondeterministically picks new memory. Now, in a language with higher-order store, the semantics involves solving recursive domain equations. With nondeterministic memory allocation, one is naturally led to recursive domain equations using powerdomains. These are problematic not only because it is unclear whether they can be used to show the existence of recursive properties of the heap but also because programs would no longer denote ω -continuous functions, due to the *countable* nondeterminism arising from memory allocation. Instead, Reus and Schwinghammer considered a functor category, indexed over finite sets of locations, which made it possible to prove that programs obeyed a frame property without relying on a nondeterministic allocator. However, this involved two non-trivial aspects. First, recursive domain equations now had to be solved not in an ordinary category of domains, but in the functor category. Second, the frame property became a recursively defined property, whose existence required a separate non-trivial proof. While Reus and Schwinghammer succeeded in defining a model that validates the frame rule, the technical complications involved make it difficult to scale the ideas to richer languages and richer logics, e.g., with higher-order frame rules [3, 2, 11].

In this paper we validate the frame rule without relying on the frame property of programs. Instead, we “bake-in” the frame rule into the interpretation of Hoare triples, using an idea from [3]. (This is described in detail in Section 4.) In particular, this approach allows us to model memory allocation by a simple deterministic allocator, so that we can model the programming language using ordinary recursively defined domains, avoiding the complications in [15]. Furthermore, the approach also allows us to validate a whole range of higher-order frame rules and to include pointer arithmetic.

The second challenge is to validate proof rules for recursion through the store [16]. Such rules essentially amount to having recursively defined specifications, which denote recursive properties of the domain for commands. It is well-known that to establish the existence of such recursive properties of domains one needs additional conditions involving, in particular, admissibility and certain forms of downward closure [14]. In [15], these conditions were ensured

$e \in \text{EXP} ::= \dots \mid 'C'$	quote (command as expression)
$C \in \text{COM} ::= \text{skip} \mid C_1; C_2 \mid \text{if } (e_1=e_2) \text{ then } C_1 \text{ else } C_2$	no op, sequencing, conditional
$\mid \text{let } x=\text{new } (e_1, \dots, e_n) \text{ in } C \mid \text{free } e$	allocation, disposal
$\mid [e_1]:=e_2 \mid \text{let } y=[e] \text{ in } C \mid \text{eval } [e]$	assignment, lookup, unquote

Fig. 1. Syntax of expressions and commands

by restricting the assertion language of the logic. In the present paper, we avoid such restrictions by changing the interpretation of triples and slightly modifying the recursion rules. In particular, we use an admissible and downwards closure of the post-condition, similar to the use of $\perp\perp$ -closure in [3] (see Section 4).

2 Programs, assertions and specifications

Programs The abstract syntax of the programming language is presented in Fig. 1. It is essentially as in [15], with dynamic allocation (but here we assume a more realistic, deterministic memory allocator) and storable, parameter-less procedures. The language is deliberately kept simple so that we can study higher-order store without distraction. We point out two features of the language which proved problematic for the semantics given in *loc. cit.* First, the language assumes that addresses are natural numbers, so that it is possible to apply arithmetic operations on addresses. Next, the language includes an allocator that deterministically picks n -consecutive cells.

Assertions The assertions P, Q, \dots used in Hoare triples are built from the formulas of classical predicate logic and the additional separation logic assertions that describe the heap ($e \mapsto e'$, **emp**, $P * Q$ and $P \multimap Q$; cf. [17]). Note that expressions in formulas can point to quoted code, as in $x \mapsto 'C'$, so that they can be used to specify properties of stored procedures. We use two abbreviations:

$$e \mapsto _ \stackrel{\text{def}}{=} \exists x'. e \mapsto x', \quad e \mapsto e_1, \dots, e_n \stackrel{\text{def}}{=} e \mapsto e_1 * e+1 \mapsto e_2 * \dots * e+n-1 \mapsto e_n.$$

where $x' \notin \text{fv}(e)$. We write $\Gamma \vdash P(: \text{Assert})$ for some finite set of variables Γ , when the assertion P contains only free variables in Γ .

Specifications Specifications are formulas of first-order intuitionistic logic with equality. In addition, it includes Hoare triples as atomic formulas and invariant extensions $\varphi \otimes P$ (from [3]):

$$\varphi, \psi ::= e_1=e_2 \mid \{P\}C\{Q\} \mid \varphi \otimes P \mid \mathbf{T} \mid \mathbf{F} \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \Rightarrow \psi \mid \exists x. \varphi \mid \forall x. \varphi$$

While assertions express properties of states, specifications describe properties of programs (sometimes using assertions inside Hoare triples). For a finite set Γ of variables, $\Gamma \vdash \varphi(: \text{Spec})$ means that Γ includes all free variables of φ .

Proof rules Our specification logic includes all the usual proof rules of intuitionistic first-order logic with equality, and special rules for Hoare triples and

PROOF RULES FOR STORED CODE

$$\begin{aligned}
& ((\forall \vec{y}. \{P\} \text{eval } [e]\{Q\}) \Rightarrow \forall \vec{y}. \{P\}C\{Q\}) \Rightarrow \forall \vec{y}. \{P * e \mapsto 'C'\} \text{eval } [e]\{Q * e \mapsto 'C'\} \\
& \quad \text{(where } \vec{y} \notin \text{fv}(e, C)) \\
& (\forall x. (\forall \vec{y}. \{P * e \mapsto x\} \text{eval } [e]\{Q * e \mapsto x\}) \Rightarrow \forall \vec{y}. \{P * e \mapsto x\}C\{Q * e \mapsto x\}) \\
& \Rightarrow \forall \vec{y}. \{P * e \mapsto 'C'\} \text{eval } [e]\{Q * e \mapsto 'C'\} \quad \text{(where } x \notin \text{fv}(P, Q, \vec{y}, e, C), \vec{y} \notin \text{fv}(e, C)) \\
& (\forall x. (\forall \vec{y}. \{P * e \mapsto x\} \text{eval } [e]\{Q\}) \Rightarrow \forall \vec{y}. \{P * e \mapsto x\}C\{Q\}) \\
& \Rightarrow \forall \vec{y}. \{P * e \mapsto 'C'\} \text{eval } [e]\{Q\} \quad \text{(where } x \notin \text{fv}(P, Q, \vec{y}, e, C), \vec{y} \notin \text{fv}(e, C))
\end{aligned}$$

PROOF RULES FOR HOARE TRIPLES

$$\begin{aligned}
& (\forall x. \{P * x \mapsto e\}C\{Q\}) \Rightarrow \{P\} \text{let } x = \text{new } e \text{ in } C\{Q\} \quad \text{(where } x \notin \text{fv}(P, Q, e)) \\
& (\forall x. \{P * e \mapsto x\}C\{Q\}) \Rightarrow \{\exists x. P * e \mapsto x\} \text{let } x = [e] \text{ in } C\{Q\} \quad \text{(where } x \notin \text{fv}(Q, e)) \\
& \{e \mapsto _ \} \text{free}(e) \{ \text{emp} \} \quad \{e \mapsto _ \} [e] := e' \{e \mapsto e'\} \\
& \frac{[[P]]_\eta^A \subseteq [[P']]_\eta^A \text{ and } [[Q']]_\eta^A \subseteq [[Q]]_\eta^A \text{ for all } \eta \in [\Gamma]}{\Gamma \vdash \{P'\}C\{Q'\} \Rightarrow \{P\}C\{Q\}}
\end{aligned}$$

PROOF RULES FOR INVARIANT EXTENSION – $\otimes P$

$$\begin{aligned}
& \varphi \Rightarrow \varphi \otimes P & \{P\}C\{P'\} \otimes Q & \Leftrightarrow \{P * Q\}C\{P' * Q\} \\
& (e_0 = e_1) \otimes Q & \Leftrightarrow e_0 = e_1 & (\varphi \otimes P) \otimes Q & \Leftrightarrow \varphi \otimes (P * Q) \\
& (\varphi \oplus \psi) \otimes P & \Leftrightarrow (\varphi \otimes P) \oplus (\psi \otimes P) & (\kappa x. \varphi) \otimes P & \Leftrightarrow \kappa x. \varphi \otimes P \\
& & \text{(where } \oplus \in \{\Rightarrow, \wedge, \vee\}) & & \text{(where } \kappa \in \{\forall, \exists\}, x \notin \text{fv}(P))
\end{aligned}$$

Fig. 2. Some proof rules

invariant extension $\varphi \otimes P$. Fig. 2 lists some of those, where the context Γ for each specification is omitted. Note that the consequence rule uses semantically valid implications for assertions, some of which can be proved using the proof rules from classical logic and the logic of Bunched Implications. In this way, the consequence rule embeds reasoning about assertions into the specification logic without the need to commit to a specific proof system for assertions.

Most of the rules in the figure are standard and known from separation logic. The only exceptions are the three proof rules for stored procedures.¹ These rules are similar to the rule for calling a parameterless, recursive procedure p declared as $p \Leftarrow C$, where C is the body of p that may contain a recursive call to p :

$$\frac{(\forall \vec{y}. \{P\} \text{call } p\{Q\}) \vdash \forall \vec{y}. \{P\}C\{Q\}}{\forall \vec{y}. \{P\} \text{call } p\{Q\}} \quad (1)$$

This rule is usually proved sound via fixpoint induction (note that p in the premiss semantically refers to any procedure with the required properties, whereas in the conclusion p refers to the declared procedure). For the language of Fig. 1,

¹ For simplicity, we do not consider mutually recursive stored procedures here, but it is straightforward to generalize our rules to handle them. Also, the first rule for stored procedures can be derived from the second and the higher-order frame rules. We include it in order to point out the subtleties of reasoning about stored procedures.

the fact that stored procedures are in use means that the declaration of a procedure is now expressed by an assertion $e \mapsto 'C'$, stating that e is a reference to the procedure with body C .

In Fig. 2 there are three rules for stored code that might call itself recursively, establishing partial correctness (and hence do not feature in the logic for *total* correctness of [6]). The first rule prohibits any access to the storing location e except through `eval` [e], whereas the second and third are more permissive. Note also that only the first two rules establish that the stored procedure called has not been altered. This is important in cases where the procedure gets updated. Updating code after its first call is a general pattern of usage of stored code, found e.g. in device drivers [5]: the first call is used for initialisation that further calls rely on.

It would be preferable to have just one rule for recursive procedures, e.g.

$$\begin{aligned} (\forall x. (\forall \vec{y}. \{P * e \mapsto x\} \text{eval } [e] \{Q\}) &\Rightarrow \forall \vec{y}. \{P * e \mapsto x\} C \{Q\}) \\ \Rightarrow \forall \vec{y}. \{P * e \mapsto 'C'\} \text{eval } [e] \{Q['C'/x]\} &\text{ (where } x \notin \text{fv}(P, \vec{y}, e, C), \vec{y} \notin \text{fv}(e, C)) \end{aligned}$$

Alas we cannot easily prove such a rule sound. Our soundness proof for the recursion rules relies on pre- and post-condition satisfying properties (a) and (b), resp., as stated in the proof of Lemma 5 (see Section 4), and an arbitrary post-condition might violate property (b). We achieve soundness by restricting the shape of the post-condition to $Q * e \mapsto x$ and stipulating the side condition $x \notin \text{fv}(Q)$. As a consequence, we cannot instantiate the recursion rules with post-condition $e \mapsto x * e' \mapsto x$, as needed for a self-copying command `let $x=[e]$ in $[e']:=x$` stored at e . Yet, also $Q * e \mapsto x * e' \mapsto x$ satisfies property 2 mentioned above, and soundness of a corresponding recursion rule could be established analogously to Lemma 5. This means the logic is incomplete; our objective has been to find rules that are easy to apply on programs with “common” use of stored procedures.

Alternatively, we could have given a single rule for stored procedures, but with a *semantic* side-condition to rule out unsuitable post-conditions.

Example 1 (Factorial). Consider the following specification and implementation:

$$\begin{aligned} F_o &\stackrel{\text{def}}{=} \text{let } x=[o] \text{ in let } r=[o+1] \text{ in} \\ &\quad \text{if } (x=0) \text{ then skip else } ([o+1]:=r \cdot x; [o]:=x-1; \text{eval } [o+2]) \\ C &\stackrel{\text{def}}{=} [o+2]:= 'F_o'; \text{eval } [o+2] \quad o \vdash \{o \mapsto 5, 1, _ \} C \{o \mapsto 0, 5!, 'F_o'\} \end{aligned}$$

The command C implements the factorial function in an object-oriented style, using three consecutive cells $(o, o+1, o+2)$. The first two cells represent fields `arg` and `res`, and the third cell denotes a method that computes the factorial of `arg` (decrementing it as a side effect) and multiplies this onto `res`. Note that the procedure F_o stored in $o+2$ calls itself by recursion through the store; see the last instruction `eval` [$o+2$] of F_o .

The specification expresses that C computes $5!$ and stores it in cell $o+1$. The key step of the proof is the derivation

$$\frac{o \vdash (\forall i, j. \{o \mapsto i, j\} \text{eval } [o+2] \{o \mapsto 0, j \cdot i!\}) \Rightarrow (\forall i, j. \{o \mapsto i, j\} F_o \{o \mapsto 0, j \cdot i!\})}{o \vdash \forall i, j. \{o \mapsto i, j, 'F_o'\} \text{eval } [o+2] \{o \mapsto 0, j \cdot i!, 'F_o'\}}$$

which shows the correctness of the stored procedure F_o . This step applies the first rule for stored procedures, and it illustrates the benefit of the rule. Here, the rule lets us hide the cell $o+2$ for code F_o in the premise, thereby giving a simple specification to discharge. The derivation of this specification itself is omitted; it involves only routine applications of standard separation logic proof rules. \square

Example 2. Next, we illustrate the typical use of the three rules for stored procedures with program C_n 's below:

$$\begin{aligned} F_1 &\stackrel{\text{def}}{=} \text{let } j=[i] \text{ in (if } j=0 \text{ then skip else } ([i]:=j-1; \text{eval } [i+1])) \\ F_2 &\stackrel{\text{def}}{=} \text{let } j=[i] \text{ in let } f=[i+1] \text{ in } ([i]:=f; \text{if } j=0 \text{ then } [i]:=0 \text{ else } ([i]:=j-1; \text{eval } [i+1])) \\ F_3 &\stackrel{\text{def}}{=} \text{let } j=[i] \text{ in (if } j=0 \text{ then } ([i+1]:='skip') \text{ else } ([i]:=j-1; \text{eval } [i+1])) \\ C_n &\stackrel{\text{def}}{=} [i+1]:=F_n; \text{eval } [i+1] \end{aligned}$$

All of the C_n 's decrease the value of i to zero (rather inefficiently), using recursion through the store. Additionally, C_2 dereferences cell $i+1$ to get the stored procedure F_2 and copy it to cell i temporarily. C_3 replaces the stored procedure in $i+1$ by `skip` at the end of the execution. For these programs, we want to prove:

$$i \vdash \{i \mapsto -, -\} C_1 \{i \mapsto 0, 'F_1'\} \quad i \vdash \{i \mapsto -, -\} C_2 \{i \mapsto 0, 'F_2'\} \quad i \vdash \{i \mapsto -, -\} C_3 \{i \mapsto 0, 'skip'\}.$$

The major step of the proof of C_1 is the use of the first rule for stored procedures:

$$\frac{i \vdash \{i \mapsto -\} \text{eval } [i+1] \{i \mapsto 0\} \Rightarrow \{i \mapsto -\} F_1 \{i \mapsto 0\}}{i \vdash \{i \mapsto -, 'F_1'\} \text{eval } [i+1] \{i \mapsto 0, 'F_1'\}}$$

which shows a property of the stored procedure F_1 . Note that the first rule successfully hides cell $i+1$ in the premise, giving us a simple subgoal to discharge. Similarly, the application of rules for stored procedures form the major steps of the proofs of the remaining triples for C_2 and C_3 :

$$\frac{i \vdash \forall x. \{i \mapsto -, x\} \text{eval } [i+1] \{i \mapsto 0, x\} \Rightarrow \{i \mapsto -, x\} F_2 \{i \mapsto 0, x\}}{i \vdash \{i \mapsto -, 'F_2'\} \text{eval } [i+1] \{i \mapsto 0, 'F_2'\}}$$

$$\frac{i \vdash \forall x. \{i \mapsto -, x\} \text{eval } [i+1] \{i \mapsto 0, 'skip'\} \Rightarrow \{i \mapsto -, x\} C_3 \{i \mapsto 0, 'skip'\}}{i \vdash \{i \mapsto -, 'F_3'\} C_3 \{i \mapsto 0, 'skip'\}}$$

Since F_2 directly accesses cell $i+1$, which stores the procedure, and F_3 updates the storing cell, we have used the second rule for C_2 and the third for C_3 . \square

3 Semantics of programs and assertions

Our interpretation of the programming language is based on a solution of a recursive domain equation, which is defined in the category **Cppo** of directed complete pointed partial orders (in short, cppo) and strict continuous functions.

Let $Nats^+$ be the set of positive natural numbers, ranged over by ℓ and n , and for $n \in Nats^+$, write $[n]$ for the set $\{1, \dots, n\}$. For a cppo A , we consider a

cpo of $Nats^+$ -labelled records with entries from A (i.e. a labelled smash product of arbitrary finite arity), which will be used to model *heaps*. Its underlying set is $Rec(A) = (\sum_{N \subseteq_{fin} Nats^+} (N \rightarrow A_{\perp}))_{\perp}$, where $(N \rightarrow A_{\perp})$ denotes the cpo of maps from the finite address set N to the cpo $A_{\perp} = A - \{\perp\}$ of non-bottom elements of A . For $\perp \neq \iota_N(r) \in Rec(A)$ we write $\mathbf{dom}(r) = N$ and use record notation $\{\ell_1 = a_1, \dots, \ell_n = a_n\}$ if $N = \{\ell_1, \dots, \ell_n\}$ and $r(\ell_i) = a_i$ for all $i \in [n]$. Note that field selection is actually application if the label is in the domain of the record (for our semantic definitions this restricted form of field selection will be sufficient). We shall also write $r[\ell \mapsto a]$ for the record that maps ℓ to a and all other $\ell' \in \mathbf{dom}(r)$ to $r(\ell')$ (assuming $\ell' \in \mathbf{dom}(r)$). In case that r is \perp , we define $r[\ell \mapsto a]$ to be \perp . The ordering on $Rec(A)$ is given by

$$r \sqsubseteq r' \stackrel{\text{def}}{\iff} r \neq \perp \implies (\mathbf{dom}(r) = \mathbf{dom}(r') \wedge \forall \ell \in \mathbf{dom}(r). r(\ell) \sqsubseteq r'(\ell)).$$

The *disjointness predicate* $r \# r'$ on records holds if $r, r' \neq \perp$ and $\mathbf{dom}(r) \cap \mathbf{dom}(r') = \emptyset$, and a continuous (*partial*) *combining operation* $r \bullet r'$ is defined by $r \bullet r' \stackrel{\text{def}}{=} \text{if } (r \# r') \text{ then } (r \cup r') \text{ else (if } (r = \perp \vee r' = \perp) \text{ then } \perp \text{ else undefined)}$.

The semantics of the programming language is given by a solution for the following domain equation:

$$Val = Integers_{\perp} \oplus Com_{\perp} \quad Heap = Rec(Val) \quad Com = Heap \multimap T_{err}(Heap)$$

where $T_{err}(D) = D \oplus \{error\}_{\perp}$ is the error monad. We usually omit the tags and (for $h \in Heap$) will simply write $h \in T_{err}(Heap)$ and $error \in T_{err}(Heap)$, resp. Recall that a solution $i : F_{Com}(Com, Com) \cong Com$ can be obtained by the usual inverse limit construction, where F_{Com} is the evident locally continuous functor obtained by separating negative and positive occurrences of Com in the right-hand sides of the three equations above.² Moreover, such a solution is a *minimal invariant*, in the sense that $id_{Com} = \text{lfp}(\lambda e : Com \multimap Com. i \circ F_{Com}(e, e) \circ i^{-1})$ [14]. The soundness proof of the rules for stored procedures exploits this fact.

Interpretation of the programming language Fig. 3 gives the interpretation $\llbracket C \rrbracket_{\eta}$ of commands in $Heap \multimap T_{err}(Heap)$ (which is isomorphic to Com), where $\eta \in Env \stackrel{\text{def}}{=} (Var \rightarrow Val_{\perp})$ is an environment mapping identifiers to (non-bottom) values in Val . An interpretation function for expressions $\llbracket e \rrbracket_{\eta}^{\mathcal{E}} \in Val_{\perp}$ is assumed, where the only non-standard cases are quoted commands. $\llbracket 'C' \rrbracket_{\eta}^{\mathcal{E}}$ is defined to be $i(\llbracket C \rrbracket_{\eta})$ (which implicitly makes use of the embedding of Com into Val). In the defining equations in Fig. 3 we assume that $h \neq \perp$, and set $\llbracket C \rrbracket_{\eta} \perp = \perp$ for all C and η . Note that the conditional only permits restricted comparison of expressions, so that commands denote continuous functions.

Interpretation of assertions Let \mathcal{P} be the set of predicates $p \subseteq Heap$ that contain \perp . The separating conjunction for these predicates, known from separation logic [17], is defined by: $h \in p_1 * p_2 \stackrel{\text{def}}{\iff} \exists h_1, h_2. h = h_1 \bullet h_2 \wedge h_1 \in p_1 \wedge h_2 \in p_2$. Note that $p * q \in \mathcal{P}$ whenever $p \in \mathcal{P}$ and $q \in \mathcal{P}$. Clearly $*$ is associative and commutative, since \bullet is, and if $p \subseteq p'$ and $q \subseteq q'$ then $p * q \subseteq p' * q'$.

² Formally, $F_{Com}(X, Y)$ is $Rec(Integers_{\perp} \oplus X_{\perp}) \multimap T_{err}(Rec(Integers_{\perp} \oplus Y_{\perp}))$.

$$\begin{aligned}
& \llbracket \text{skip} \rrbracket_\eta h \stackrel{\text{def}}{=} h \\
& \llbracket C_1; C_2 \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } \llbracket C_1 \rrbracket_\eta h \in \{\perp, \text{error}\} \text{ then } \llbracket C_1 \rrbracket_\eta \text{ else } \llbracket C_2 \rrbracket_\eta (\llbracket C_1 \rrbracket_\eta h) \\
& \llbracket \text{if } e=e' \text{ then } C_1 \text{ else } C_2 \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } \{\llbracket e_1 \rrbracket_\eta^\varepsilon, \llbracket e_2 \rrbracket_\eta^\varepsilon\} \subseteq \text{Com} \text{ then } \perp \\
& \quad \text{else if } (\llbracket e \rrbracket_\eta^\varepsilon = \llbracket e' \rrbracket_\eta^\varepsilon) \text{ then } \llbracket C_1 \rrbracket_\eta h \text{ else } \llbracket C_2 \rrbracket_\eta h \\
& \llbracket \text{let } x=\text{new } e_1, \dots, e_n \text{ in } C \rrbracket_\eta h \stackrel{\text{def}}{=} \text{let } \ell = \min\{\ell \mid \forall \ell'. (\ell \leq \ell' < \ell+n) \Rightarrow \ell' \notin \text{dom}(h)\} \\
& \quad \text{in } \llbracket C \rrbracket_{\eta[x \mapsto \ell]} (h \bullet \{\ell = \llbracket e_1 \rrbracket_\eta^\varepsilon, \dots, \ell+n-1 = \llbracket e_n \rrbracket_\eta^\varepsilon\}) \\
& \llbracket \text{free } e \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } \llbracket e \rrbracket_\eta^\varepsilon \notin \text{dom}(h) \text{ then } \text{error} \\
& \quad \text{else (let } h' \text{ s.t. } h = h' \bullet \{\llbracket e \rrbracket_\eta^\varepsilon = h(\llbracket e \rrbracket_\eta^\varepsilon)\}) \text{ in } h' \\
& \llbracket [e_1] := e_2 \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } \llbracket e_1 \rrbracket_\eta^\varepsilon \notin \text{dom}(h) \text{ then } \text{error} \text{ else } (h[\llbracket e_1 \rrbracket_\eta^\varepsilon \mapsto \llbracket e_2 \rrbracket_\eta^\varepsilon]) \\
& \llbracket \text{let } x=[e] \text{ in } C \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } \llbracket e \rrbracket_\eta^\varepsilon \notin \text{dom}(h) \text{ then } \text{error} \text{ else } \llbracket C \rrbracket_{\eta[x \mapsto h(\llbracket e \rrbracket_\eta^\varepsilon)]} h \\
& \llbracket \text{eval } [e] \rrbracket_\eta h \stackrel{\text{def}}{=} \text{if } (\llbracket e \rrbracket_\eta^\varepsilon \notin \text{dom}(h) \vee h(\llbracket e \rrbracket_\eta^\varepsilon) \notin \text{Com}) \text{ then } \text{error} \\
& \quad \text{else } i^{-1}(h(\llbracket e \rrbracket_\eta^\varepsilon))(h)
\end{aligned}$$

Fig. 3. Interpretation of commands $\llbracket C \rrbracket_\eta \in \text{Heap} \multimap T_{\text{err}}(\text{Heap})$

$$\begin{aligned}
& \llbracket e_1 \leq e_2 \rrbracket_\eta^A \stackrel{\text{def}}{=} \{h \in \text{Heap} \mid h \neq \perp \Rightarrow \llbracket e_i \rrbracket_\eta^\varepsilon \in \text{Integers} \wedge \llbracket e_1 \rrbracket_\eta^\varepsilon \leq \llbracket e_2 \rrbracket_\eta^\varepsilon\} \\
& \llbracket e_1 = e_2 \rrbracket_\eta^A \stackrel{\text{def}}{=} \{h \in \text{Heap} \mid h \neq \perp \Rightarrow \llbracket e_1 \rrbracket_\eta^\varepsilon = \llbracket e_2 \rrbracket_\eta^\varepsilon\} \\
& \llbracket \forall x. P \rrbracket_\eta^A \stackrel{\text{def}}{=} \bigcap \{\llbracket P \rrbracket_{\eta[x \mapsto v]}^A \mid v \in \text{Val}\} \quad \llbracket \text{emp} \rrbracket_\eta^A \stackrel{\text{def}}{=} \{\{\}, \perp\} \quad \llbracket P * Q \rrbracket_\eta^A \stackrel{\text{def}}{=} \llbracket P \rrbracket_\eta^A * \llbracket Q \rrbracket_\eta^A \\
& \llbracket e \mapsto e' \rrbracket_\eta^A \stackrel{\text{def}}{=} \{h \in \text{Heap} \mid h \neq \perp \Rightarrow \text{dom}(h) = \{\llbracket e \rrbracket_\eta^\varepsilon\} \wedge h(\llbracket e \rrbracket_\eta^\varepsilon) = \llbracket e' \rrbracket_\eta^\varepsilon\}
\end{aligned}$$

Fig. 4. Interpretation $\llbracket P \rrbracket_\eta^A : \text{Env} \rightarrow \mathcal{P}$ of assertions

The poset (\mathcal{P}, \subseteq) forms a complete boolean BI algebra.³ Thus, we get a canonical BI hyperdoctrine $\mathbf{Set}(-, \mathcal{P})$, which soundly models classical (higher-order) predicate BI [1]. In particular this yields an interpretation for the quantifiers. Some cases of this interpretation of assertions are spelled out in Fig. 4.

4 Semantics of specifications

We now define the interpretation of specifications, and show how it addresses the two key challenges described in the introduction. The most interesting components of our interpretation are semantic Hoare triples, which we will use to interpret (syntactic) Hoare triples. For each predicate $p \in \mathcal{P}$, let $\text{Ad}(p)$ be the admissible, downward closure of p in $T_{\text{err}}(\text{Heap})$ (i.e., the smallest admissible, downward-closed subset of $T_{\text{err}}(\text{Heap})$ that includes p ; it may be obtained as the intersection of all admissible, downward-closed subsets of Heap that include p).

³ The negation and false of this boolean algebra are slightly unusual, and are defined by $\neg p \stackrel{\text{def}}{=} (\text{Heap} - p) \cup \{\perp\}$ and $\text{false} \stackrel{\text{def}}{=} \{\perp\}$. Conjunction, disjunction and true are defined as in the usual powerset boolean algebra.

Definition 1 (Semantic triple). A semantic Hoare triple is a triple of predicates $p, q \in \mathcal{P}$ and function $c \in F_{Com}(Com, Com)$, written $\{p\}c\{q\}$. A semantic triple $\{p\}c\{q\}$ is valid, denoted $\models \{p\}c\{q\}$, if and only if, for all $r \in \mathcal{P}$ and all $h \in Heap$, we have that $h \in p * r \Rightarrow c(h) \in Ad(q * r)$.

Intuitively, a semantic triple $\{p\}c\{q\}$ specifies that c should transform an input state in p to an output state in q . Furthermore, the triple says that this transformation should modify only the portion of memory for p (because, otherwise, it would not preserve some invariant r when r was $*$ -attached to the precondition p). Note that $\models \{p\}c\{q\}$ ensures the absence of memory errors for inputs in $p * r$ for all r , because $Ad(q * r)$ cannot contain *error*.

We point out two important aspects of valid semantic Hoare triples and their relationships to the points raised in the introduction. First, the definition of validity includes a universal quantification over $*$ -added invariants r . Since we will interpret (syntactic) Hoare triples using the validity of semantic triples, this universal quantification means that Hoare triples in our logic impose a stronger requirement on commands than the ones in standard separation logic. In particular, the requirement is strong enough to imply the frame rule:

Lemma 1 (Frame rule). If $\models \{p\}c\{q\}$, then $\models \{p * r\}c\{q * r\}$ for all $r \in \mathcal{P}$.

In this way, our model addresses the first challenge in the introduction regarding the soundness of the frame rule. Second, the definition of the validity takes the admissible, downward closure $Ad(q * r)$ of post-conditions. As a result, whenever we define a subset of $F_{Com}(Com, Com)$ using a semantic Hoare triple, it is guaranteed that the resulting set is admissible and downward-closed:

Lemma 2. For all $p, q \in \mathcal{P}$, the subset $\{c \mid \{p\}c\{q\} \text{ is valid}\}$ is an admissible, downward-closed subset of $F_{Com}(Com, Com)$.

It is this property that lets us prove the soundness of the proof rules for stored procedures, without requiring any additional conditions, such as a syntactic restriction on assertions [15].

We interpret specifications following the usual Kripke semantics of intuitionistic logic as classical logic is unsound in presence of higher-order frame rules. Our interpretation uses a particular Kripke structure that lets us validate all the higher-order frame rules, i.e., rules for invariant extension $\varphi \otimes P$. Concretely, the Kripke structure is the preorder $(\mathcal{P}, \sqsubseteq)$ where the relation \sqsubseteq is defined by: $p \sqsubseteq q \stackrel{\text{def}}{\iff} \exists r \in \mathcal{P}. p * r = q$. Each world p in this Kripke structure should be thought of as an invariant to be added by (higher-order) frame rules, and the preorder $p \sqsubseteq q$ denotes that q is obtained by extending p with some disjoint invariant r , see also [3] from which we use some results.

Some cases of the definition of the satisfaction relation \models are shown in Fig. 5. Note that Hoare triples are interpreted using the validity of semantic triples.

Soundness We recall one consequence of our semantics, which is discussed in more detail in [3]. It is the soundness of the generalized frame rule: $\varphi \Rightarrow \varphi \otimes P$. Since the interpretation follows the standard Kripke semantics, every formula φ

$$\begin{aligned}
\eta, p \models \varphi \wedge \psi &\stackrel{\text{def}}{\iff} \eta, p \models \varphi \text{ and } \eta, p \models \psi \\
\eta, p \models \varphi \Rightarrow \psi &\stackrel{\text{def}}{\iff} \text{for all } r \in \mathcal{P}, \text{ if } p \sqsubseteq r \text{ and } \eta, r \models \varphi, \text{ then } \eta, r \models \psi \\
\eta, p \models \varphi \otimes P &\stackrel{\text{def}}{\iff} \eta, p * \llbracket P \rrbracket_\eta^A \models \varphi \\
\eta, p \models \{P\}C\{Q\} &\stackrel{\text{def}}{\iff} \models \{\llbracket P \rrbracket_\eta^A * p\} \llbracket C \rrbracket_\eta \{\llbracket Q \rrbracket_\eta^A * p\}
\end{aligned}$$

Fig. 5. Interpretation $\eta, p \models \varphi$ of specifications

satisfies the usual Kripke monotonicity: $\forall \eta, r, q. (\eta, r \models \varphi) \wedge (r \sqsubseteq q) \Rightarrow (\eta, q \models \varphi)$. Since $r \sqsubseteq q$ just means that $q = r * p$ for some p , the above monotonicity condition is equivalent to $\forall \eta, r, p. (\eta, r \models \varphi) \Rightarrow (\eta, r * p \models \varphi)$. This just means that adding an invariant p for each specification maintains the truth of a specification, and explains why the generalized frame rule is sound in our semantics.

Lemma 3 (Invariants, [3]). *All the axioms for invariant extensions are sound.*

Our semantics validates all the proof rules for specifications. In the following, we focus on the second rule for stored procedures.

Lemma 4 (Recursion). *The second rule for stored procedures is sound.*

Proof. For each $\eta \in \llbracket \Gamma \rrbracket$ and $r \in \mathcal{P}$, define a predicate $A_{\eta, r}$ on $Com \times Com$ by

$$A_{\eta, r}(c, d) \stackrel{\text{def}}{\iff} \forall \vec{v} \in Val^n. \models \{ \llbracket P * e \mapsto x \rrbracket_{\eta_1}^A * r \} i^{-1}(d) \{ \llbracket Q * e \mapsto x \rrbracket_{\eta_1}^A * r \}$$

where $\eta_1 = \eta[\vec{y} \mapsto \vec{v}, x \mapsto c]$. Pick any $\eta \in \llbracket \Gamma \rrbracket$ and $r \in \mathcal{P}$. By the definition of $\llbracket \text{eval } [e] \rrbracket$ and the usual substitution lemma (which holds for our interpretation), the soundness of the rule boils down to proving the following implication.

$$(\forall c \in Com. \forall r' \sqsupseteq r. A_{\eta, r'}(c, c) \Rightarrow A_{\eta, r'}(c, \llbracket C' \rrbracket_\eta)) \Rightarrow A_{\eta, r}(\llbracket C' \rrbracket_\eta, \llbracket C' \rrbracket_\eta).$$

Suppose that there is a predicate $S_{\eta', r'}$ on Com parameterized by (η', r') , such that (1) $S_{\eta', r'}(c) \iff (\forall d \in Com. S_{\eta', r'}(d) \Rightarrow A_{\eta', r'}(d, c))$. Then, we have that (2) $\forall c. S_{\eta, r}(c) \Rightarrow A_{\eta, r}(c, c)$. Hence, assuming the precondition $\forall c. \forall r' \sqsupseteq r. A_{\eta, r'}(c, c) \Rightarrow A_{\eta, r'}(c, \llbracket C' \rrbracket_\eta)$, we obtain $\forall c. S_{\eta, r}(c) \Rightarrow A_{\eta, r}(c, \llbracket C' \rrbracket_\eta)$ and therefore $S_{\eta, r}(\llbracket C' \rrbracket_\eta)$ by (1). But then (2) shows $A_{\eta, r}(\llbracket C' \rrbracket_\eta, \llbracket C' \rrbracket_\eta)$, as required. It remains to establish the existence of a predicate $S_{\eta', r'}$ satisfying (1). This is done in the following Lemma 5. \square

Lemma 5 (Existence). *For all η, r , there exists $S_{\eta, r} \subseteq Com$ such that $S_{\eta, r}(c)$ holds iff $\forall d. S_{\eta, r}(d) \Rightarrow A_{\eta, r}(d, c)$, where $A_{\eta, r}$ is as in the proof of Lemma 4.*

Proof. The proof builds on the same technique as used in [16], but many details have changed. Let \mathcal{C} denote the set of admissible subsets of Com , which forms a complete lattice when ordered by \sqsubseteq . Pick η and $r \in \mathcal{P}$. We define an operation $\Phi: \mathcal{C}^{op} \rightarrow \mathcal{C}$, by $S \mapsto \{c \in Com \mid \forall d. d \in S \Rightarrow A_{\eta, r}(d, c)\}$. That $\Phi(S)$ is admissible follows from the admissibility of $A_{\eta, r}(d, -)$, which itself comes from

Lemma 2. The symmetrisation $\Phi^{\S}(S, T) \stackrel{\text{def}}{=} \langle \Phi(T), \Phi(S) \rangle$ of Φ is a monotonic map on the complete lattice $\mathcal{C}^{op} \times \mathcal{C}$ and thus has a least (pre-) fixed point (S^-, S^+) , by Tarski's fixed point theorem. Then (S^+, S^-) is also a fixed point of Φ^{\S} , so one obtains $S^+ \subseteq S^-$. A predicate $S_{\eta, r} \in \mathcal{C}$ with the required property $S_{\eta, r} = \Phi(S_{\eta, r})$ is obtained by proving the opposite inclusion.

To this end, for $l \sqsubseteq id_{Com}$ and $S_1, S_2 \in \mathcal{C}$, define $l : S_1 \subset S_2$ to mean that $\forall c \in S_1. l(c) \in S_2$. Note that from

$$(1) \quad l : S_1 \subset S_2 \Rightarrow (i \circ F_{Com}(l, l) \circ i^{-1}) : \Phi(S_2) \subset \Phi(S_1)$$

for all $l \sqsubseteq id_{Com}$, it follows by fixed point induction that $lfp(\lambda l. i \circ F_{Com}(l, l) \circ i^{-1}) : S^- \subset S^+$. This is equivalent to $id_{Com} : S^- \subset S^+$, i.e., $S^- \subseteq S^+$, because $lfp(\dots)$ is id_{Com} by the minimal invariant property of Com .

It remains to prove (1). For this, one needs only prove the following two properties. Let $Cl^\downarrow(p)$ be the downward closure of a predicate p . For all environments η' , heaps h and functions l with $l \sqsubseteq id_{Com}$, if $j \stackrel{\text{def}}{=} Rec(\hat{l})$,

- (a) $h \in \llbracket P * e \mapsto x \rrbracket_{\eta'}^A$ implies $j(h) \in Cl^\downarrow \llbracket P * e \mapsto x \rrbracket_{\eta'[\hat{l}(x)]}^A$,
- (b) $h \in \llbracket Q * e \mapsto x \rrbracket_{\eta'[\hat{l}(x)]}^A$ implies $j(h) \in Cl^\downarrow \llbracket Q * e \mapsto x \rrbracket_{\eta'}^A$.

To see why it suffices to prove (a) and (b), suppose $l \sqsubseteq id_{Com}$ satisfies $l : S_1 \subset S_2$. Pick $c \in \Phi(S_2)$. We have to show $(i \circ F_{Com}(l, l) \circ i^{-1})(c) \in \Phi(S_1)$. Thus, for all $d \in S_1$, we must show that $A_{\eta, r}(d, (i \circ F_{Com}(l, l) \circ i^{-1})(c))$ holds, i.e., for all $\vec{v} \in Val^n$

$$(2) \quad \models \{ \llbracket P * e \mapsto x \rrbracket_{\eta[\vec{y} \mapsto \vec{v}, x \mapsto d]}^A * r \} F_{Com}(l, l)(i^{-1}(c)) \{ \llbracket Q * e \mapsto x \rrbracket_{\eta[\vec{y} \mapsto \vec{v}, x \mapsto d]}^A * r \}.$$

For this, pick $d \in S_1$ and $\vec{v} \in Val^n$. Since $l : S_1 \subset S_2$, we have that $l(d) \in S_2$, and since $c \in \Phi(S_2)$, it must be the case that

$$(3) \quad \models \{ \llbracket P * e \mapsto x \rrbracket_{\eta[\vec{y} \mapsto \vec{v}, x \mapsto l(d)]}^A * r \} i^{-1}(c) \{ \llbracket Q * e \mapsto x \rrbracket_{\eta[\vec{y} \mapsto \vec{v}, x \mapsto l(d)]}^A * r \}.$$

We will now prove that (3) implies (2).

To simplify notation, we assume without loss of generality that η is such that $\eta(\vec{y}) = \vec{v}$. Pick $r' \in \mathcal{P}$ and $h \in \llbracket P * e \mapsto x \rrbracket_{\eta[x \mapsto d]}^A * r * r'$. Let j be $Rec(\hat{l})$. Then, we have to show the set membership below:

$$F_{Com}(l, l)(i^{-1}(c))(h) = T_{err}(j)(i^{-1}(c)(j(h))) \in Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto d]}^A * r * r').$$

By property (a) and definition of j , $j(h)$ is in $Cl^\downarrow(\llbracket P * e \mapsto x \rrbracket_{\eta[x \mapsto l(d)]}^A * r * r')$. So, we have (4) $i^{-1}(c)(j(h)) \in Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto l(d)]}^A * r * r')$, because of (3) and the monotonicity of $i^{-1}(c)$. Note that by the property (b) and the definition of j , $T_{err}(j)$ should map heaps in $(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto l(d)]}^A * r * r')$ to those in $Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto d]}^A * r * r')$. Furthermore, for all continuous functions f on $T_{err}(Heap)$, if f maps every heap in a predicate p into $Ad(q)$, it also maps all heaps in $Ad(p)$ into $Ad(q)$. Thus, since $T_{err}(j)$ is continuous, it maps heaps in $Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto l(d)]}^A * r * r')$ into $Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto d]}^A * r * r')$. By (4), this means that $T_{err}(j)(i^{-1}(c)(j(h)))$ belongs to $Ad(\llbracket Q * e \mapsto x \rrbracket_{\eta[x \mapsto d]}^A * r * r')$. \square

5 Conclusion and future work

We have developed a simple model of separation logic for a language with higher-order store. The model validates proof rules for recursion through the store and a wide range of higher-order frame rules. Future work includes extending the model to richer programming languages, in particular to languages with higher-order functions. In order to obtain modularity it is also necessary to develop a version of the logic where assertions do not contain code explicitly but rather abstract specifications of its behaviour. We are confident that the simplicity of the present model will make that possible. In future work we also plan to extend the relationally parametric model of separation logic in [3] to higher-order store.

References

1. B. Biering, L. Birkedal, and N. Torp-Smith. BI-hyperdoctrines, higher-order separation logic, and abstraction. *ACM TOPLAS*, 29(5), 2007.
2. L. Birkedal, N. Torp-Smith, and H. Yang. Semantics of separation-logic typing and higher-order frame rules for algol-like languages. *LMCS*, 2(5:1), 2006.
3. L. Birkedal and H. Yang. Relational parametricity and separation logic. In *Proc. FOSSACS'07*, volume 4423 of *LNCS*, 2007.
4. H. Cai, Z. Shao, and A. Vaynberg. Certified self-modifying code. In *Proc. PLDI'07*, pages 66–77, 2007.
5. J. Corbet, A. Rubini, and G. Kroah-Hartman. *Linux Device Drivers*. O'Reilly, 3rd edition, 2005.
6. K. Honda, N. Yoshida, and M. Berger. An observationally complete program logic for imperative higher-order functions. In *Proc. LICS'05*, pages 270–279, 2005.
7. N. Krishnaswami, J. Aldrich, and L. Birkedal. Modular verification of the subject-observer pattern via higher-order separation logic. In *FTfJP'07*, 2007.
8. A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *Proc. ESOP'07*, volume 4421 of *LNCS*, pages 189–204, 2007.
9. A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and separation in Hoare type theory. In *Proc. ICFP'06*, pages 62–73, 2006.
10. P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Proc. CSL*, volume 2142 of *LNCS*, pages 1–18, 2001.
11. P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *Proc. of 31st POPL*, pages 268–280, 2004.
12. M. Parkinson. When separation logic met Java. In *FTfJP'06*, 2006.
13. M. Parkinson and G. Bierman. Separation logic, abstraction and inheritance. In *Proc. 35th POPL*, 2008.
14. A. M. Pitts. Relational properties of domains. *Information and Computation*, 127:66–90, 1996.
15. B. Reus and J. Schwinghammer. Separation logic for higher-order store. In *Proc. CSL'06*, volume 4207 of *LNCS*, pages 575–590, 2006.
16. B. Reus and T. Streicher. About Hoare logics for higher-order store. In *Proc. ICALP'05*, *LNCS*, pages 1337–1348, 2005.
17. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. LICS'02*, pages 55–74, 2002.