

Capsules and Separation

Jean-Baptiste Jeannin
Computer Science Department
Cornell University
Ithaca, NY 14853-7501, USA

Dexter Kozen
Computer Science Department
Cornell University
Ithaca, NY 14853-7501, USA

Abstract—We study a formulation of separation logic using capsules, a representation of the state of a computation in higher-order programming languages with mutable variables. We prove soundness of the frame rule in this context and investigate alternative formulations with weaker side conditions.

I. INTRODUCTION

Capsules [8] give a simple and mathematically appealing approach to semantics of higher-order programs that is consistent with both the functional and imperative paradigms. They minimally extend the classical λ -calculus to allow mutable variables, thus enabling the construction of certain coterms (infinite terms) representing recursive functions without the need for fixpoint combinators. They have a well-defined statically-scoped evaluation semantics, are typable with simple types, and are Turing complete.

Perhaps the most important aspect of capsules is that their evaluation semantics captures static scoping without introducing cumbersome combinatorial machinery needed to implement closures. Closures have been well known since the earliest days of functional programming. They were introduced to rectify a bug¹ in the original evaluation semantics of LISP [12]. It is perhaps well understood that static scoping corresponds to β -reduction with safe substitution in the λ -calculus, and that closures correctly implement this, but there has

¹subsequently upgraded to a feature, now known as *dynamic scoping*

apparently never been a formal proof until quite recently [8]. Capsules were the key ingredient that made this proof mathematically tractable.

Capsules replace heaps, stores, stacks, and pointers with the single mathematical concept of variable binding, yet are equally expressive and represent the same data dependencies and liveness structure. In a sense, *capsules are to closures what graphs are to their adjacency list representations*. One would not reason mathematically about graphs in terms of their adjacency list representations; neither should one reason about the evaluation semantics of higher order programs in terms of their implementations. With no side effects, λ -calculus with β -reduction is the model of choice, but it does not allow for mutable variables.

Separation logic is a logic for the study of locality and shared data. Introduced by Reynolds in a series of lectures in the late 1990s, based on an earlier idea of Burstall, separation logic has been widely studied in the last decade [1], [6], [14], [16], [20], [21]. The difficulties of reasoning in the presence of heaps, stores, stacks, and pointers are no more apparent than in this literature. Several papers [4], [5], [15], [24] cite notoriously thorny issues from dangling pointers to arcane side conditions needed for soundness. Reynolds himself acknowledged that his original formulation of separation logic was flawed [21]. Chief among the difficulties is the issue of catastrophic failure due to the dereferencing of unbound variables or dangling pointers. There seems to be an unspoken belief that this is an

unavoidable aspect that must be confronted in any realistic model of computation.

On the contrary, we believe that the essential structure of separation is independent of these encumbrances. It is our thesis that freedom from catastrophic failure is the responsibility of the programming language, not the logic. Capsule semantics provides this assurance, even in the presence of higher-order constructs and mutable variables.

In this paper, we propose a semantics for separation logic based on capsules. The formulation works for higher-order programs with mutable variables. In §IV we give the semantics and prove the soundness of the frame rule in this context. In §IV-D, we study the motivation behind the nonstandard definition of partial correctness preferred in much of the literature of separation logic [5], [21] and investigate alternatives. It is here that the advantages of capsule semantics in the study of separation can best be seen.

A. Related Work

There are two main differences between our work and the previous work on separation logic. In previous work, the authors usually adopt either an imperative, C-style programming language with low-level heap operations, or a functional, ML-style programming language with immutable variables and explicit references. According to Mason and Talcott [11], in functional languages there are two approaches to introducing objects with memory: the LISP approach, where all variables are mutable, and the ML approach, where all variables are immutable and references are introduced. One of the reasons why the ML view is usually chosen for separation logic on functional programs is that having immutable variables is the only way to get a correct semantics based on closures [18]. By using capsules instead, we are able to use mutable variables in the style of LISP.

The second main difference in this work is that we insist that all capsule environments σ are closed, i.e., that every free variable appearing in a $\sigma(x)$ should be defined in σ . To us, this seems like a very natural thing. But as far as we know, none of the previous work requires anything like this. When using C-style languages with an environment and a heap, writing down a similar condition would require both the environment and the heap, whereas the separation logic definitions usually only use heaps. Even Neelakantan Krishnaswami et al. [10], though using an ML-style language, explicitly say that they permit dangling pointers as long as the pointers themselves are well typed. Note that, if trying to relate the semantics of capsules with, say, a more traditional semantics using closures and a heap, the capsule environment behaves like a heap rather than like an environment in the traditional sense [7].

The original work on separation logic, summarized by Reynolds [21], uses an imperative, C-style programming language with low-level commands and already gives a proof of a version of the frame rule.

Our work is most closely related to work by Krishnaswami, Birkedal, Aldrich and Reynolds [9], [10], who give a separation logic for ML. However, our system allows mutable variables in the style of LISP, whereas theirs uses explicit references allocated in an explicit heap.

Birkedal, Torp-Smith and Yang [3] also study the frame rule in the context of a higher-order language, idealized Algol extended with heaps, but their stack variables are immutable as well.

There has been some work on so-called higher-order stores [2], [19], [22], where some code can be stored in a heap cell. Because any λ -abstractions can be stored in the environment, and executing some of them can have side-effects, our setup naturally supports higher-order stores.

II. CAPSULE DEFINITIONS

In this section we briefly review the definition of capsules and their semantics from [8].

A. Syntax

Expressions $\text{Exp} = \{d, e, a, b, \dots\}$ contain both functional and imperative features. There is an unlimited supply of *variables* x, y, z, \dots of all (simple) types, as well as *constants* f, c, \dots for primitive values. $()$ is the only constant of type `unit`, and `true` and `false` are the only two constants of type `bool`. In the examples, $0, 1, 2, \dots$ are predefined constants of type `int`. In addition, there are functional features

- λ -abstraction $\lambda x. e$
- application $(d e),$

imperative features

- assignment $x := e$
- composition $d; e$
- conditional `if b then d else e`
- while loop `while b do $e,$`

and defined expressions

- let $x = d$ in e $(\lambda x. e) d$
- let `rec $x = d$ in e` `let $x = a$ in $x := d; e$`

where a is any expression of the appropriate type. The technique for formation of recursive functions in the last definition is known as *Landin's knot*.

Let `Var` be the set of variables, `Const` the set of constants, and $\lambda\text{-Abs}$ the set of λ -abstractions. Given an expression e , let $\text{FV}(e)$ denote the set of free variables of e . Given a partial function $h : \text{Var} \rightarrow \text{Var}$ such that $\text{FV}(e) \subseteq \text{dom } h$, let $h(e)$ be the expression e where every instance of a free variable $x \in \text{FV}(e)$ has been replaced by the variable $h(x)$. Thus $h : \text{Exp} \rightarrow \text{Exp}$ is the unique homomorphic extension of $h : \text{Var} \rightarrow \text{Var}$. Given two partial functions g and h , $g \circ h$ denotes their composition: $g \circ h(x) = g(h(x))$.

Given a function h , we write $h[x/v]$ the function such that $h[x/v](y) = h(y)$ for $y \neq x$ and $h[x/v](x) = v$. Given an expression e , we write $e[x/y]$ for the expression e with y substituted for all free occurrences of x .

Types α, β, \dots are ordinary simple types built inductively from an unspecified family of base types, including at least `unit` and `bool`, and the usual function type constructor \rightarrow . All constants c of the language have a type $\text{type}(c)$; by convention, we use c for a constant of a base type and f for a constant of a functional type. Γ is a type environment, a partial function $\text{Var} \rightarrow \text{Type}$. As is standard, we write $\Gamma, x : \alpha$ for the typing environment Γ where x has been bound or rebound to α . The typing rules are standard:

$$\begin{array}{c} \Gamma \vdash c : \alpha \text{ if } \text{type}(c) = \alpha \quad \Gamma, x : \alpha \vdash x : \alpha \\ \frac{\Gamma, x : \alpha \vdash e : \beta}{\Gamma \vdash \lambda x. e : \alpha \rightarrow \beta} \\ \frac{\Gamma \vdash d : \alpha \rightarrow \beta \quad \Gamma \vdash e : \alpha}{\Gamma \vdash (d e) : \beta} \\ \frac{\Gamma \vdash x : \alpha \quad \Gamma \vdash e : \alpha}{\Gamma \vdash x := e : \text{unit}} \quad \frac{\Gamma \vdash d : \text{unit} \quad \Gamma \vdash e : \alpha}{\Gamma \vdash d; e : \alpha} \\ \frac{\Gamma \vdash b : \text{bool} \quad \Gamma \vdash d : \alpha \quad \Gamma \vdash e : \alpha}{\Gamma \vdash \text{if } b \text{ then } d \text{ else } e : \alpha} \\ \frac{\Gamma \vdash b : \text{bool} \quad \Gamma \vdash e : \text{unit}}{\Gamma \vdash \text{while } b \text{ do } e : \text{unit}} \end{array}$$

Henceforth all the expressions we consider will be assumed to be well-typed with respect to these rules.

An expression is *irreducible* if it is either a constant or a λ -abstraction. Note that variables are not irreducible. Let $\text{Irred} = \text{Const} + \lambda\text{-Abs}$ denote the set of irreducible terms. (These are often called *values* in the λ -calculus literature, but we avoid this terminology here because it is misleading, as they are not values in the intuitive sense.)

A *capsule environment* is a partial function $\sigma : \text{Var} \rightarrow \text{Irred}$ satisfying the following closure

condition:

$$\forall x \in \text{dom } \sigma \quad \text{FV}(\sigma(x)) \subseteq \text{dom } \sigma.$$

This says that all free variables appearing in expressions $\sigma(x)$ must also be bound to an expression. Thus free variables are not really free; every variable in σ either occurs in the scope of a λ or is bound by σ to an expression. There may be circularities; this enables a representation of recursive functions.

The *closure* of a set $A \subseteq \text{dom } \sigma$ with respect to σ , denoted $\text{cl}_\sigma(A)$, is the smallest set B containing A such that if $x \in B$ then $\text{FV}(\sigma(x)) \subseteq B$. It is the domain of the least-defined capsule environment whose domain contains A and that agrees with σ on its domain.

A *capsule* is a pair $\langle e, \sigma \rangle$, where e is an expression and σ is a capsule environment, such that $\text{FV}(e) \subseteq \text{dom } \sigma$. As above, every variable appearing in $\langle e, \sigma \rangle$ either occurs in the scope of a λ or is bound by σ to an expression. These conditions preclude catastrophic failure due to access of unbound variables.

The term α -conversion refers to the renaming of bound variables. With a capsule $\langle e, \sigma \rangle$, this can happen in two ways. The traditional form maps a subterm $\lambda x.d$ to $\lambda y.d[x/y]$, provided y would not be captured in d . We call this α -conversion of the first kind. One can also rename a variable $x \in \text{dom } \sigma$ and all free occurrences of x in e and $\sigma(z)$ for $z \in \text{dom } \sigma$ to y , provided $y \notin \text{dom } \sigma$ already and y would not be captured. We call this α -conversion of the second kind.

B. Semantics

Capsule evaluation semantics looks very much like the original evaluation semantics of LISP, with the added twist that a fresh variable is substituted for the parameter in β -reductions. The relevant small-step rule is

$$\langle (\lambda x.e \ v), \sigma \rangle \rightarrow \langle e[x/y], \sigma[y/v] \rangle \quad (y \text{ fresh})$$

In the original evaluation semantics of LISP, the right-hand side is $\langle e, \sigma[x/v] \rangle$, which gives

dynamic scoping. As proved in [8], this simple change faithfully models β -reduction with safe substitution in the λ -calculus, providing static scoping without closures.

Another evaluation rule of particular note is the assignment rule:

$$\langle x := v, \sigma \rangle \rightarrow \langle (), \sigma[x/v] \rangle$$

where v is irreducible. The closure conditions on capsules ensure that x must already be bound in σ . The variable x is rebound to the irreducible expression v .

Other small-step rules are

$$\begin{aligned} \langle x, \sigma \rangle &\rightarrow \langle \sigma(x), \sigma \rangle \\ \langle f \ c, \sigma \rangle &\rightarrow \langle f(c), \sigma \rangle \\ \langle (); e, \sigma \rangle &\rightarrow \langle e, \sigma \rangle \\ \langle \text{if true then } d \ \text{else } e, \sigma \rangle &\rightarrow \langle d, \sigma \rangle \\ \langle \text{if false then } d \ \text{else } e, \sigma \rangle &\rightarrow \langle e, \sigma \rangle \\ \langle \text{while } b \ \text{do } e, \sigma \rangle &\rightarrow \\ &\langle \text{if } b \ \text{then } (e; \text{while } b \ \text{do } e) \ \text{else } (), \sigma \rangle \end{aligned}$$

There are also context rules that define a standard shallow applicative-order (leftmost innermost, call-by-value) evaluation strategy. The reduction rules preserve types and cannot fail catastrophically. Thus every computation either continues forever or terminates with a well-typed final capsule $\langle v, \sigma \rangle$, where v is irreducible.

The relation $\xrightarrow{*}$ is the reflexive transitive closure of \rightarrow .

See [8] for several examples that illustrate how the system manages recursive functions, static scoping, and garbage collection.

C. Assertions

Assertions P, Q, \dots are statements in some logical system, possibly with free variables in Var . We write $\text{FV}(P)$ for the set of free variables of P . These variables are subject to interpretation provided by a capsule environment σ .

The exact nature of the underlying logic is unimportant—it could be propositional, first order, second order or higher order—but we do require a few basic properties common to standard logical systems. There must be a well-defined satisfaction relation \models such that $\sigma \models P$ has a truth value when the free variables of P are interpreted by the capsule environment σ . Use of the metaexpression $\sigma \models P$ carries the tacit assumption that $\text{FV}(P) \subseteq \text{dom } \sigma$. We will augment the logic with the separation logic operators $*$ and $\text{--}*$ (defined later in §IV). Finally, we require the following (natural) property to hold: if $\sigma \models P$, and $z \in \text{dom } \sigma - \text{FV}(P)$, then z can be renamed via α -conversion of the second kind without affecting the truth of P . In examples, we will use first order logic with $*$ and $\text{--}*$, and equality on base types.

III. PARTIAL CORRECTNESS

The traditional definition of partial correctness and the definition used in the literature on separation logic diverge in a subtle and interesting way. The difference hinges on whether the precondition is required to assert the absence of catastrophic failure due to dangling pointers or lookup of unbound variables; this is the **abort** condition of Reynolds [21] or the **fault** condition of Calcagno, O’Hearn, and Yang [5]. Our view, however, is that avoidance of catastrophic failure is the responsibility of the programming language semantics, not the program logic, and capsules do just that. Can this condition then be eliminated? In this section we shed some light on this question.

Let P, Q be assertions and e a program. At issue is the meaning of the partial correctness assertion $\{P\} e \{Q\}$. Consider the following three metastatements, each parameterized by a closed environment σ :

- (A $_{\sigma}$) $\sigma \models P$
- (B $_{\sigma}$) $\text{FV}(e) \subseteq \text{dom } \sigma$
- (C $_{\sigma}$) if $\langle e, \sigma \rangle \xrightarrow{*} \langle v, \tau \rangle$ and v is irreducible, then $\tau \models Q$.

Statement (A $_{\sigma}$) entails $\text{FV}(P) \subseteq \text{dom } \sigma$, because the definition of \models does not make sense without it. More strongly, $\text{cl}_{\sigma}(\text{FV}(P)) \subseteq \text{dom } \sigma$, since σ is closed. Statement (B $_{\sigma}$) is equivalent to the assertion that $\langle e, \sigma \rangle$ is a valid capsule. Reynolds’s definition [21] uses a slightly different formulation

$$(B'_{\sigma}) \neg(\langle e, \sigma \rangle \xrightarrow{*} \mathbf{abort})$$

in place of (B $_{\sigma}$). Here $\langle e, \sigma \rangle$ need not be a valid capsule. The semantics of capsule evaluation already precludes **abort**, thus (B'_{\sigma}) is always true if $\langle e, \sigma \rangle$ is a capsule; that is, (B $_{\sigma}$) implies (B'_{\sigma}).

Now consider the following potential interpretations of $\{P\} e \{Q\}$.

$$\{P\} e \{Q\} \Leftrightarrow \forall \sigma (A_{\sigma}) \wedge (B_{\sigma}) \Rightarrow (C_{\sigma}) \quad (1)$$

$$\{P\} e \{Q\} \Leftrightarrow \forall \sigma (A_{\sigma}) \Rightarrow (B_{\sigma}) \wedge (C_{\sigma}) \quad (2)$$

Definition (1) says that if the precondition P holds of the input state σ and the evaluation of $\langle e, \sigma \rangle$ terminates normally, then the output state τ satisfies the postcondition Q . This is the naive interpretation used in traditional forms of Hoare logic. Alternatively, the version preferred in the literature on separation logic would be (2), the difference being that the precondition P must ensure that the evaluation of $\langle e, \sigma \rangle$ cannot terminate abnormally.

Reynolds’s version [21] is actually slightly weaker, using (B'_{\sigma}) instead of (B $_{\sigma}$):

$$\{P\} e \{Q\} \Leftrightarrow \forall \sigma (A_{\sigma}) \Rightarrow (B'_{\sigma}) \wedge (C_{\sigma}) \quad (3)$$

However, the difference is inconsequential: if $\{P\} e \{Q\}$ holds in the sense of (3) but not (2), then there exists a variable $x \in \text{FV}(e) - \text{dom } \sigma$ for some σ satisfying P , and consequently $x \in \text{FV}(e) - \text{cl}_{\sigma}(\text{FV}(P))$; but by (B'_{\sigma}), x can never be referenced or assigned in the evaluation of $\langle e, \sigma \rangle$. Thus the presence or absence of x in the domain of σ affects neither the truth of P nor the evaluation of $\langle e, \sigma \rangle$.

But there is a much more important benefit to (2) over (3). Consider the metastatement

(B) $\text{FV}(e) \subseteq \text{FV}(P)$.

A consequence of (2) is that (A_σ) implies (B_σ) for all σ . If P is satisfiable at all, say by some σ , then (B) must hold, since variables in $\text{dom } \sigma$ not occurring free in P can be renamed (by an α -conversion of the second kind—see §II-A) without affecting the truth of P . Thus (2) holds with (B) in place of (B_σ) . Moreover, since (B) is independent of σ , assuming P is satisfiable at all, (2) is equivalent to the definition

$$\{P\} e \{Q\} \Leftrightarrow (B) \wedge (\forall \sigma (A_\sigma \Rightarrow (C_\sigma))) \quad (4)$$

Note that, unlike (B_σ) and (B'_σ) , the condition (B) is syntactically checkable, thus suitable as a side condition in a rule of inference. If we like, we may remove the condition (B) in the definition of $\{P\} e \{Q\}$ and instead introduce it as a side condition in the frame rule. However, can it be eliminated entirely? That is, is the formulation (1) sound? We show in §IV-D that it is not. In fact, even only slightly weaker forms of the side condition (B) do not suffice for soundness.

IV. CAPSULES AND SEPARATION LOGIC

A. Definitions

Here is our semantics for separation logic in terms of capsules. Call closed environments σ and τ *independent* and write $\sigma \perp \tau$ if their domains are disjoint. Define $\sigma + \tau$ to be the join of σ and τ , provided they are independent. That is,

$$(\sigma + \tau)(x) = \begin{cases} \sigma(x), & \text{if } x \in \text{dom } \sigma, \\ \tau(x), & \text{if } x \in \text{dom } \tau, \\ \text{undefined}, & \text{otherwise.} \end{cases}$$

Define *separating conjunction* by

$$\sigma \models P * Q$$

if there exist σ_1 and σ_2 such that $\sigma = \sigma_1 + \sigma_2$, $\sigma_1 \models P$, and $\sigma_2 \models Q$. Define *separating implication* by

$$\sigma \models P -* Q$$

if $\sigma + \tau \models Q$ whenever $\tau \models P$ and $\sigma + \tau$ exists. It is easily seen that capsule environments form a *separation algebra* in the sense of [5] under these definitions. That is, the structure

$$(\{\text{capsule environments}\}, +, \emptyset)$$

is a cancellative partial commutative monoid. This means that $+$ is a commutative and associative partial binary operation with identity \emptyset satisfying the *cancellative property*: the partial function $+$ is injective in each variable. The relation $\sigma \perp \tau$ holds if and only if $\sigma + \tau$ is defined.

It follows from results of [5] that separating conjunction $*$ and separating implication $-*$ satisfy the usual intuitionistic relationship: For all closed σ such that $\text{FV}(P) \cup \text{FV}(Q) \cup \text{FV}(R) \subseteq \text{dom } \sigma$,

$$\sigma \models (P * Q) -* R \Leftrightarrow \sigma \models P -* (Q -* R).$$

Other axioms of separation logic mentioned in [21] are also easily checked:

$$\begin{aligned} (P \vee Q) * R &\Leftrightarrow (P * R) \vee (Q * R) \\ (P \wedge Q) * R &\Rightarrow (P * R) \wedge (Q * R) \\ (\exists x P) * Q &\Leftrightarrow \exists x (P * Q) \quad (x \notin \text{FV}(Q)) \\ (\forall x P) * Q &\Rightarrow \forall x (P * Q) \quad (x \notin \text{FV}(Q)). \end{aligned}$$

B. The Frame Rule

The soundness of the frame rule was first proved in [24] for the heap model of computation. Our proof is essentially the same as the one given in [21], but somewhat shorter due to the simplifications afforded by capsule semantics.

Lemma 4.1: If

$$\langle e, \sigma_1 + \sigma_2 \rangle \xrightarrow{*} \langle e, \tau \rangle$$

and $\text{FV}(e) \subseteq \text{dom } \sigma_1$ (that is, $\langle e, \sigma_1 \rangle$ is a capsule), then for some τ_1 , $\langle e, \sigma_1 \rangle \xrightarrow{*} \langle e, \tau_1 \rangle$ and $\tau = \tau_1 + \sigma_2$.

Proof: By induction on the derivation. None of the small-step evaluation rules listed

in §II-B access any variable outside the domain of σ_1 except for fresh variables introduced in the application rule. In particular, the environment σ_2 is not touched during the evaluation. ■

Theorem 4.2: Under capsule semantics, the frame rule

$$\frac{\{P\} e \{Q\}}{\{P * R\} e \{Q * R\}}$$

is sound with respect to definition (2) or (4) of partial correctness assertions. Equivalently, the frame rule is sound with respect to definition (1) of partial correctness assertions in the presence of the side condition $\text{FV}(e) \subseteq \text{FV}(P)$.

Proof: As argued in §III, in all cases we can assume $\text{FV}(e) \subseteq \text{FV}(P)$. Suppose $\{P\} e \{Q\}$. Let $\sigma \models P * R$. Then $\sigma = \sigma_1 + \sigma_2$ with $\sigma_1 \models P$ and $\sigma_2 \models R$. Then $\text{FV}(R) \subseteq \text{dom } \sigma_2$ and $\text{FV}(e) \subseteq \text{FV}(P) \subseteq \text{dom } \sigma_1$, therefore $\langle e, \sigma_1 \rangle$ is a valid capsule. Since $\langle e, \sigma \rangle \xrightarrow{*} \langle v, \tau \rangle$, by Lemma 4.1 there exists τ_1 such that $\langle e, \sigma_1 \rangle \xrightarrow{*} \langle v, \tau_1 \rangle$ and $\tau = \tau_1 + \sigma_2$, and $\tau_1 \models Q$ by the premise of the rule. Thus $\tau \models Q * R$. ■

C. Discussion

Calcagno, O’Hearn, and Yang [5] argue that the soundness of the frame rule for a given evaluation semantics is equivalent to the following two properties.

Safety Monotonicity: If $\langle e, \sigma_0 \rangle$ is safe and $\sigma = \sigma_0 + \sigma_1$, then $\langle e, \sigma \rangle$ is safe.

Frame Property: If $\langle e, \sigma_0 \rangle$ is safe, $\sigma = \sigma_0 + \sigma_1$, and $\langle e, \sigma \rangle \xrightarrow{*} \langle e, \sigma' \rangle$, then there exists σ'_0 such that $\sigma' = \sigma'_0 + \sigma_1$ and $\langle e, \sigma_0 \rangle \xrightarrow{*} \langle e, \sigma'_0 \rangle$.

(Here we are allowing $\langle e, \sigma \rangle$ to violate the closure conditions in the definition of capsules, and *safe* means that (B'_σ) holds.) In their words, “The first condition says that if a state has enough resources for safe execution of a command, then so do superstates. The second condition says that if a state has enough resources

for the command to execute safely, then execution on any bigger state can be tracked back to the small state.”

With capsules, the safety monotonicity property is vacuously true, and the frame property reduces to Lemma 4.1.

D. Alternative Conditions

Recall from §III the side condition

$$(B) \text{FV}(e) \subseteq \text{FV}(P),$$

for which the frame rule with semantics (1) for partial correctness assertions is sound. One might ask whether there is a weaker side condition that suffices for soundness. In this section we show that there is not much hope.

The frame rule as stated by Reynolds has a side condition, which says that “no variable occurring free in R is modified by e ” [21]. A literal formulation of the side condition in terms of capsules is

$$(C) \text{AV}(e) \cap \text{FV}(R) = \emptyset,$$

where $\text{AV}(e)$, the *assigned variables* of e , is the set of $x \in \text{FV}(e)$ having a free occurrence on the left-hand side of an assignment $:=$. This is a bit confusing, because (C) seems to serve no purpose:

Theorem 4.3:

- (a) Under semantics (2) of partial correctness assertions, the side condition (C) in the frame rule is redundant.
- (b) Under semantics (1) of partial correctness assertions, the frame rule with side condition (C) is not sound.

Proof: First (a). As argued in §III, semantics (2) is equivalent to semantics (1) with side condition (B), provided P is satisfiable at all. We show that in all such nontrivial instances, (C) is subsumed by (B).

Suppose $\sigma \models P * R$. Then

$$\sigma = \sigma_1 + \sigma_2 \quad \sigma_1 \models P \quad \sigma_2 \models R.$$

By (B), we have

$$\text{AV}(e) \subseteq \text{FV}(e) \subseteq \text{FV}(P) \subseteq \text{dom } \sigma_1$$

and also

$$\text{FV}(R) \subseteq \text{dom } \sigma_2 \quad \text{dom } \sigma_1 \cap \text{dom } \sigma_2 = \emptyset,$$

therefore (C) holds.

For (b), we give a counterexample to soundness. Let σ consist of the bindings

$$\sigma(f) = \lambda().x \quad \sigma(x) = 2.$$

Let $R = R(f)$ be the assertion $f() = 2$. Let e be the program $x := 3$. Let $P = Q = \text{true}$. The corresponding instance of the frame rule is

$$\frac{\{\text{true}\} x := 3 \{\text{true}\}}{\{\text{true} * f() = 2\} x := 3 \{\text{true} * f() = 2\}}$$

The premise $\{\text{true}\} x := 3 \{\text{true}\}$ holds, but the conclusion does not. We have

$$\sigma \models \text{true} * f() = 2,$$

where $\sigma = \emptyset + \sigma$, $\emptyset \models \text{true}$, and $\sigma \models f() = 2$ and \emptyset is the empty environment. The program e does not assign to f , the only variable free in R , yet it indirectly alters the value of f by assigning a new value to x , making R false. ■

We remark that Theorem 4.3(b) holds not just for capsules, but for any programming language with records, arrays, objects, pointers, or any form of aliasing whatsoever.

The problem at first seems to be that it is not enough to say that no variable in $\text{FV}(R)$ may be modified by e ; we must ensure that no variable in the closure of $\text{FV}(R)$ may be modified by e , so that e cannot even indirectly alter R . This is the condition

$$(B_1) \quad \forall \sigma \quad \sigma \models P * R \Rightarrow \text{AV}(e) \cap \text{cl}_\sigma(\text{FV}(R)) = \emptyset$$

which is not expressible by any syntactic property of e , P , Q , and R .

But even this is not enough for soundness. Condition (B₁) is implied by the very strong syntactic property

$$(B_2) \quad \text{AV}(e) \subseteq \text{FV}(P)$$

which is only slightly weaker than (B). It asserts that all free variables assigned by e are mentioned by P . Nevertheless, even (B₂) is not enough for soundness. At first this may seem quite counterintuitive, because (B₂) implies that starting in any state satisfying $P * R$, e cannot change any variable in the closure of $\text{FV}(R)$, therefore cannot affect the truth of R . We state it as a theorem.

Theorem 4.4: The frame rule under semantics (1) for partial correctness assertions with side condition (B₂) is not sound.

Proof: Let σ consist of the bindings

$$\sigma(g) = \lambda x.2 \quad \sigma(f) = \lambda().3.$$

Let $R = R(f)$ be the assertion $f() = 3$. Let e be the program

$$g := \lambda x.\text{if } x = 1 \text{ then } f() \text{ else } 2.$$

Let P and Q both be the assertion $g(0) = 2$. The corresponding instance of the frame rule is

$$\frac{\{g(0) = 2\} e \{g(0) = 2\}}{\{g(0) = 2 * f() = 3\} e \{g(0) = 2 * f() = 3\}}$$

The premise $\{g(0) = 2\} e \{g(0) = 2\}$ holds, as does the side condition (B₂), since

$$\text{AV}(e) = \{g\} = \text{FV}(P).$$

However, the conclusion does not. We have

$$\sigma \models g(0) = 2 * f() = 3,$$

where $\sigma = \sigma_1 + \sigma_2$, $\text{dom } \sigma_1 = \{g\}$, $\text{dom } \sigma_2 = \{f\}$, $\sigma_1 \models g(0) = 2$, and $\sigma_2 \models f() = 3$. However, after execution of the program e , the resulting environment binds g to a term containing a free occurrence of f , so g and f cannot be separated. ■

V. CONCLUSION AND FUTURE WORK

We were motivated to undertake this study in response to an anonymous review of [8] claiming that capsules “contradict the insights of separation logic which has been extensively researched for the last decade.” We hope that we have convinced the reader that there is no contradiction whatsoever—in fact quite the opposite! Capsules provide a novel perspective on separation logic, because they capture the same locality and persistence structure as traditional heap models, but in a simpler, more mathematically tractable framework. We feel that this has great potential for enhancing the understanding of separation by allowing research to focus on the essentials.

We have only begun to scratch the surface in this work. We would like to investigate other structures that have arisen in the study of separation logic to see whether capsules can contribute there as well. The preliminary results of this paper leave us optimistic.

In particular, higher-order separation logic [1] proposes to use the much more powerful higher-order logic in predicates. Nested Hoare triples [22] are a neat idea to specify code stored in the heap. The anti-frame rule [17], [23] presents a very interesting way of modeling hidden state. Finally, we would like to study the concurrency rule [13] in the context of capsules.

ACKNOWLEDGMENTS

We would like to thank Neelakantan Krishnaswami for suggesting that we look at the relation between capsules and separation logic after attending a presentation of [7]. We would also like to thank Mark Bickford, Bob Constable, and François Pottier for many useful discussions.

REFERENCES

- [1] Bodil Biering, Lars Birkedal, and Noah Torp-Smith. BI-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Trans. Program. Lang. Syst.*, 29, August 2007.
- [2] Lars Birkedal, Bernhard Reus, Jan Schwinghammer, and Hongseok Yang. A simple model of separation logic for higher-order store. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part II, ICALP '08*, pages 348–360, Berlin, Heidelberg, 2008. Springer-Verlag.
- [3] Lars Birkedal, Noah Torp-Smith, and Hongseok Yang. Semantics of separation-logic typing and higher-order frame rules for algol-like languages. *CoRR*, abs/cs/0610081, 2006.
- [4] Richard Bornat, Cristiano Calcagno, and Hongseok Yang. Variables as resource in separation logic. In *Proc. 21st Conf. Math. Found. Programming Semantics*, pages 247–276, 2005.
- [5] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local action and abstract separation logic. In *Proc. 22nd Annual IEEE Symp. Logic in Computer Science (LICS07)*, pages 366–378. IEEE, 2007.
- [6] Samin S. Ishtiaq and Peter W. O’Hearn. Bi as an assertion language for mutable data structures. In *Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL ’01, pages 14–26, New York, NY, USA, 2001. ACM.
- [7] Jean-Baptiste Jeannin. Capsules and closures. In Michael Mislove and Joël Ouaknine, editors, *Proc. 27th Conf. Math. Found. Programming Semantics (MFPS XXVII)*, Pittsburgh, PA, May 2011. Elsevier Electronic Notes in Theoretical Computer Science.
- [8] Jean-Baptiste Jeannin and Dexter Kozen. Computing with capsules. Technical Report <http://hdl.handle.net/1813/22082>, Computing and Information Science, Cornell University, January 2011.
- [9] Neelakantan R. Krishnaswami. *Verifying Higher-Order Imperative Programs with Higher-Order Separation Logic*. PhD thesis, Carnegie Mellon University, 2010.
- [10] Neelakantan R. Krishnaswami, Lars Birkedal, Jonathan Aldrich, and John C. Reynolds. Idealized ML and its separation logic. <http://www.cs.cmu.edu/~neelk/>, 2007.
- [11] Ian Mason and Carolyn Talcott. Axiomatizing operational equivalence in the presence of side effects. In *Fourth Annual Symposium on Logic in Computer Science. IEEE*, pages 284–293. IEEE Computer Society Press, 1989.
- [12] John McCarthy. History of LISP. In Richard L. Wexelblat, editor, *History of programming languages I*, pages 173–185. ACM, 1981.

- [13] P. O'Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375(1-3):271–307, May 2007.
- [14] Peter W. O'Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Proceedings of the 15th International Workshop on Computer Science Logic, CSL '01*, pages 1–19, London, UK, 2001. Springer-Verlag.
- [15] Matthew Parkinson, Richard Bornat, and Cristiano Calcagno. Variables as resource in hoare logics. In *Proceedings of the 21st Annual IEEE Symposium on Logic in Computer Science*, pages 137–146, Washington, DC, USA, 2006. IEEE Computer Society.
- [16] Matthew J. Parkinson and Gavin M. Bierman. Separation logic, abstraction and inheritance. In *Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages, POPL '08*, pages 75–86, New York, NY, USA, 2008. ACM.
- [17] François Pottier. Hiding local state in direct style: a higher-order anti-frame rule. In *Twenty-Third Annual IEEE Symposium on Logic In Computer Science (LICS'08)*, pages 331–340, Pittsburgh, Pennsylvania, June 2008.
- [18] François Pottier, 2012. Private communication.
- [19] Bernhard Reus and Jan Schwinghammer. Separation logic for higher-order store. In *In Proc. CSL*, pages 575–590. Springer, 2006.
- [20] John C. Reynolds. Intuitionistic reasoning about shared mutable data structures. In J. Davies, B. Roscoe, and J. Woodcock, editors, *Millennial Perspectives in Computer Science*, pages 303–321. Palgrave, 2000.
- [21] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th IEEE Symp. Logic in Computer Science (LICS'02)*, pages 55–74. IEEE, 2002.
- [22] Jan Schwinghammer, Lars Birkedal, Bernhard Reus, and Hongseok Yang. Nested Hoare triples and frame rules for higher-order store. In *In Proceedings of the 18th EACSL Annual Conference on Computer Science Logic*, 2009.
- [23] Jan Schwinghammer, Hongseok Yang, Lars Birkedal, François Pottier, and Bernhard Reus. A semantic foundation for hidden state. In *FOSSACS*, pages 2–17, 2010.
- [24] H. Yang and P. W. O'Hearn. A semantic basis for local reasoning. In M. Nielsen and U. Engberg, editors, *Proc. 5th Foundations of Software Science and Computation Structures (FOSSACS02)*, volume 2303 of *Lecture Notes in Computer Science*, pages 402–416. Springer-Verlag, 2002.