

Weak updates and separation logic

Gang Tan¹, Zhong Shao², Xinyu Feng³, and Hongxu Cai⁴

¹Lehigh University, ²Yale University

³Toyota Technological Institute at Chicago, ⁴Google Inc.

Abstract. Separation Logic (SL) provides a simple but powerful technique for reasoning about imperative programs that use shared data structures. Unfortunately, SL supports only “strong updates”, in which mutation to a heap location is safe only if a unique reference is owned. This limits the applicability of SL when reasoning about the interaction between many high-level languages (e.g., ML, Java, C#) and low-level ones since these high-level languages do not support strong updates. Instead, they adopt the discipline of “weak updates”, in which there is a global “heap type” to enforce the invariant of type-preserving heap updates. We present SL^W , a logic that extends SL with reference types and elegantly reasons about the interaction between strong and weak updates. We also describe a semantic framework for reference types; this framework is used to prove the soundness of SL^W .

1 Introduction

Reasoning about mutable, aliased heap data structures is essential for proving properties or checking safety of imperative programs. Two distinct approaches perform such kind of reasoning: Separation Logic, and a type-based approach employed by many high-level programming languages.

Extending Hoare Logic, the seminal work of Separation Logic (SL [11, 14]) is a powerful framework for proving properties of low-level imperative programs. Through its separating conjunction operator and frame rule, SL supports local reasoning about heap updates, storage allocation, and explicit storage deallocation.

SL supports “strong updates”: as long as a unique reference to a heap cell is owned, the heap-update rule of SL allows the cell to be updated with any value:

$$\frac{\{(e \mapsto -) * p\}[e] := e' \{ (e \mapsto e') * p \}}{(1)}$$

In the above heap-update rule, there is no restriction on the new value e' . Hereafter, we refer to heaps with strong updates as *strong heaps*. Heap cells in strong heaps can hold values of different types at different times of program execution.

Most high-level programming languages (e.g., Java, C#, and ML), however, support only “weak updates”. In this paradigm, programs can perform only type-preserving heap updates. There is a global “heap type” that tells the type of every allocated heap location. The contents at a location have to obey the prescribed type of the location in the heap type, at any time. Managing heaps with weak updates is a simple and type-safe mechanism for programmers to access memory. As an example, suppose an ML

variable has type “ τ ref” (i.e., it is a reference to a value of type τ). Then any update through this reference with a new value of type τ is type safe and does not affect other types, even in the presence of aliases and complicated points-to relations. Hereafter, we refer to heaps with weak updates as *weak heaps*.

This paper is concerned with the interaction between strong and weak updates. Strong-update techniques are more precise and powerful, allowing destructive memory updates and explicit deallocation. But aliases and uniqueness have to be explicitly tracked. Weak-update techniques allow type-safe management of memory without tracking aliases, but types of memory cells can never change. A framework that mixes strong and weak updates enables a trade-off between precision and scalability.

Such a framework is also useful for reasoning about *multilingual programs*. Most real-world programs are developed in multiple programming languages. Almost all high-level languages provide foreign function interfaces for interfacing with low-level C code (for example, the OCaml/C FFI, and the Java Native Interface). Real-world programs consist of a mixture of code in both high-level and low-level languages. A runtime state for such a program conceptually contains a union of a weak heap and a strong heap. The weak heap is managed by a high-level language (e.g., Java), accepts type-preserving heap updates, and is garbage-collected. The strong heap is managed by a low-level language, accepts strong updates, and its heap cells are manually recollected. To check the safety and correctness of multilingual programs, it is of practical value to have one framework that accommodates both strong and weak updates.

Since Separation Logic (SL) supports strong heaps, one natural thought to mix strong and weak updates is to extend SL with types so that assertions can also describe weak heaps. That is, in addition to regular SL assertions, we add $\{e \mapsto \tau\}$, which specifies a heap with a single cell and the cell holds a value of type τ . This scheme, however, would encounter two challenges.

First, allowing general reference types in $\{e \mapsto \tau\}$ would make SL unsound. An example demonstrating this point is as follows:

$$\{\{x \mapsto 4\} * \{y \mapsto \text{even ref}\}\} [x] := 3 \{\{x \mapsto 3\} * \{y \mapsto \text{even ref}\}\} \quad (2)$$

The example is an instantiation of the heap-update rule in (1) and uses the additional assertion $\{e \mapsto \tau\}$. The precondition states that y points to a heap cell whose contents are of type “even ref” (i.e., a reference to an even integer). Therefore, the precondition is met on a heap where y points to x . However, the postcondition will not hold on the new heap after the update because x will point to an odd number. Therefore, the above rule is sound only if y does not point to x .

The second challenge of adding types to SL is how to prove its soundness with mixed SL assertions and types. Type systems are usually proved sound following a syntactic approach [20], where types are treated as syntax. Following the tradition of Hoare Logic, SL’s soundness is proved through a denotational model, and SL assertions are interpreted semantically. There is a need to resolve the conflict between syntactic and semantic soundness proofs.

In this paper, we propose a hybrid logic, SL^W , which mixes SL and a type system. Although the logic is described in a minimal language and type system, it makes a solid step toward a framework that reasons about the interaction between high-level and low-level languages. The most significant technical aspects of the logic are as follows:

$$\begin{aligned}
(\text{Command}) \quad c &::= \dots \mid x := [e] \mid [x] := e \mid x := \text{alloc}(e) \mid \text{free}(e) \\
(\text{Expression}) \quad e &::= x \mid v \mid \text{op}(e_1, \dots, e_n) \\
(\text{Value}) \quad v &::= n \mid \ell
\end{aligned}$$

Fig. 1. Language syntax

- SL^{W} extends SL with a simple type system. It employs SL for reasoning about strong updates, and employs the type system for weak updates. Most interestingly, SL^{W} mixes SL assertions and types, and accommodates cross-boundary pointers (from weak to strong heaps and vice versa). This is achieved by statically maintaining the distinction between pointers to weak heaps and pointers to strong heaps. SL^{W} is presented in Section 2.
- To resolve the conflict between syntactic types and semantic assertions, we propose a semantic model of types. Our model of reference types follows a fixed-point approach and allows us to define a denotational model of SL^{W} and prove its soundness. The model of SL^{W} is presented in Section 3.

2 SL^{W} : Separation logic with weak updates

We next describe SL^{W} , an extension of SL that incorporates reasoning over weak heaps. In Section 2.1, we describe a minimal language that enables us to develop SL^{W} . Rules of SL^{W} are presented in Section 2.2 and examples of using the logic in Section 2.3.

We first describe some common notations. For a map f , we write $f[x \rightsquigarrow y]$ for a new map that agrees with f except it maps x to y . For two finite maps f_1 and f_2 , $f_1 \uplus f_2$ is the union of f_1 and f_2 when their domains are disjoint, and undefined otherwise. We write $f \setminus X$ for a new map after removing elements in X from the domain of f . We write \vec{x} for a sequence of x s and ε for an empty sequence.

2.1 Language syntax and semantics

Figure 1 presents the syntax of the programming language in which we will develop SL^{W} . The language is the imperative language used by Hoare [7], augmented with a set of commands for manipulating heap data structures. It is similar to the one used in Reynolds’ presentation of SL [14]. Informally, the command “ $x := [e]$ ” loads the contents at location e into variable x ; “[x] := e ” updates the location at x with the value e ; “ $x := \text{alloc}(e)$ ” allocates a new location, initializes it with e , and assigns the new location to x ; “ $\text{free}(e)$ ” deallocates the location e .

In the syntax, we use n for integers, x for variables, ℓ for heap locations, and op for arithmetic operators. We assume there is an infinite number of variables and locations.

Figure 2 presents a formal operational semantics of the language. A state consists of a map \mathbf{r} from variables to values, a heap h , and a sequence of commands. Commands bring one state to another state and their semantics is formally defined by a step relation \mapsto . We write \mapsto^* for the reflexive and transitive closure of \mapsto .

A state may have no next state, i.e., “getting stuck”. For example, a state whose next instruction to execute is $[x] := e$ gets stuck when x does not represent a location or the

(State) $s ::= (\mathbf{r}, h, \vec{c})$
(Local) $\mathbf{r} ::= \text{Var} \rightarrow \text{Value}$
(Heap) $h ::= \{\ell_1 \mapsto v_1, \dots, \ell_n \mapsto v_n\}$

$(\mathbf{r}, h, c \cdot \vec{c}_1) \mapsto (\mathbf{r}_2, h_2, \vec{c}_2)$	
if $c =$	then $(\mathbf{r}_2, h_2, \vec{c}_2) =$
...	...
$x := [e]$	$(\mathbf{r}[x \rightsquigarrow h(\ell)], h, \vec{c}_1)$ when $\mathbf{r}(e) = \ell$ and $\ell \in \text{dom}(h)$
$[x] := e$	$(\mathbf{r}, h[\ell \rightsquigarrow \mathbf{r}(e)], \vec{c}_1)$ when $\mathbf{r}(x) = \ell$ and $\ell \in \text{dom}(h)$
$x := \text{alloc}(e)$	$(\mathbf{r}[x \rightsquigarrow \ell], h \uplus \{\ell \mapsto \mathbf{r}(e)\}, \vec{c}_1)$ when $\ell \notin \text{dom}(h)$
$\text{free}(e)$	$(\mathbf{r}, h \setminus \{\ell\}, \vec{c}_1)$ when $\mathbf{r}(e) = \ell$ and $\ell \in \text{dom}(h)$

$$\text{where } \mathbf{r}(e) = \begin{cases} \mathbf{r}(x) & \text{when } e = x \\ v & \text{when } e = v \\ \text{op}(\mathbf{r}(e_1), \dots, \mathbf{r}(e_n)) & \text{when } e = \text{op}(e_1, \dots, e_n) \end{cases}$$

Fig. 2. Operational semantics

location is not in the domain of the state's heap. A state is a terminal state when the sequence of commands is empty.

Definition 1. (*Stuck and terminal states*)

$$\begin{aligned} \text{stuck}(s) &\triangleq \neg(\exists s'. s \mapsto s') \\ \text{terminal}(\mathbf{r}, h, \vec{c}) &\triangleq \vec{c} = \varepsilon \end{aligned}$$

Below we define the usual notions of safety and termination:

Definition 2. (*Safety and termination*)

$$\begin{aligned} \text{safe}(s) &\triangleq \forall s'. ((s \mapsto^* s') \wedge \neg \text{terminal}(s')) \Rightarrow \exists s''. s' \mapsto s'' \\ \text{terminate}(s) &\triangleq \forall s'. s \mapsto^* s' \Rightarrow \exists s''. s' \mapsto^* s'' \wedge \text{terminal}(s'') \end{aligned}$$

2.2 The logic SL^{W}

Figure 3 presents assertions and types used in SL^{W} . Assertions in SL^{W} include all formulas in predicate calculus (not shown in the figure), and all SL formulas. The only additional assertion form in SL^{W} is $\{e : \tau\}$, which denotes that e has type τ .

SL^{W} is equipped with a simple type system that classifies integers and locations. Although the type system does not include many types in high-level languages, by including reference types it is already sufficient to show interesting interactions between strong and weak heaps. Reference types are the most common types when high-level languages interoperate with low-level languages because in this setting most data are passed by references.

$$\begin{aligned}
(\text{Assertion}) \quad p &::= \dots \mid \text{emp} \mid \{e_1 \mapsto e_2\} \mid p_1 * p_2 \mid p_1 \text{--} * p_2 \mid \boxed{\{e : \tau\}} \\
(\text{Type}) \quad \tau &::= \text{int} \mid \text{ref} \mid \text{wref } \tau \\
(\text{HeapType}) \quad \Psi &::= \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\} \\
(\text{LocalVarType}) \quad \Gamma &::= \{x_1 : \tau_1, \dots, x_n : \tau_n\}
\end{aligned}$$

Fig. 3. Assertions and types

$$\boxed{\Psi, \Gamma \vdash e : \tau}$$


$$\begin{array}{c}
\frac{x \in \text{dom}(\Gamma)}{\Psi, \Gamma \vdash x : \Gamma(x)} \qquad \frac{}{\Psi, \Gamma \vdash n : \text{int}} \qquad \frac{}{\Psi, \Gamma \vdash \ell : \text{ref}} \\
\frac{\Psi(\ell) = \tau}{\Psi, \Gamma \vdash \ell : \text{wref } \tau} \qquad \frac{\forall i \in [1..n]. \Psi, \Gamma \vdash e_i : \text{int}}{\Psi, \Gamma \vdash \text{op}(e_1, \dots, e_n) : \text{int}}
\end{array}$$

Fig. 4. Typing rules for expressions

Type `int` is for all integers and `ref` for all locations. Type “`wref τ` ” is for locations in a weak heap, but not in a strong heap. A heap type Ψ tells the type of every location in a weak heap; mathematically, it is a finite map from locations to types. Given heap type Ψ , location ℓ has type “`wref τ` ” if $\Psi(\ell)$ equals τ . A local variable type, Γ , tells the type of local variables.

Figure 4 presents typing rules for expressions, which are unsurprising. Notice that the typing rule for “`wref τ` ” requires that the location ℓ is in the domain of the heap type Ψ and $\Psi(\ell)$ has to be the same as τ . This rule and the later weak-update rule enforce type-preserving updates on weak heaps.

The following schematic diagram helps to understand the relationship between weak heaps, strong heaps, local variables, assertions and various kinds of types in SL^W :



As shown in the diagram, SL^W conceptually divides a heap into a weak heap h_w and a strong heap h_s . The weak heap is specified by a heap type Ψ , and the strong heap by SL formula p . Pointers to weak-heap cells (in solid lines) have type “`wref τ` ” or `ref`. Pointers to strong heap cells (in dotted lines) can have only type `ref`.

Figure 5 presents rules for checking commands. These rules use the judgment $\Psi \vdash \{\Gamma, p\} \vec{c} \{\Gamma', p'\}$. In this judgment, Ψ , Γ and p are preconditions and specify conditions on the weak heap, local variables, and the strong heap respectively. Postcon-

$$\boxed{\Psi \vdash \{\Gamma, \mathbf{p}\} \vec{c} \{\Gamma', \mathbf{p}'\}}$$

(Well-formed statements)

$$\frac{\Psi, \Gamma \vdash e : \text{ref} \quad \Psi, \Gamma \vdash y : \tau}{\Psi \vdash \{\Gamma, \{e \mapsto y\}\} x := [e] \{\Gamma[x \rightsquigarrow \tau], x = y \wedge \{e \mapsto x\}\}} \text{ (S-LOAD)}$$

where $x \notin \text{FV}(e)$

$$\frac{\Psi, \Gamma \vdash x : \text{ref}}{\Psi \vdash \{\Gamma, \{x \mapsto -\}\} [x] := e \{\Gamma, \{x \mapsto e\}\}} \text{ (S-UPDATE)}$$

$$\frac{}{\Psi \vdash \{\Gamma, \text{emp}\} x := \text{alloc}(e) \{\Gamma[x \rightsquigarrow \text{ref}], \{x \mapsto e\}\}} \text{ (S-ALLOC)}$$

where $x \notin \text{FV}(e)$

$$\frac{\Psi, \Gamma \vdash e : \text{ref}}{\Psi \vdash \{\Gamma, \{e \mapsto -\}\} \text{free}(e) \{\Gamma, \text{emp}\}} \text{ (S-FREE)}$$

$$\frac{\Psi \vdash \{\Gamma, \mathbf{p}\} \vec{c} \{\Gamma', \mathbf{p}'\}}{\Psi \vdash \{\Gamma, \mathbf{p} * \mathbf{p}_1\} \vec{c} \{\Gamma', \mathbf{p}' * \mathbf{p}_1\}} \text{ (FRAME)}$$

where no variable occurring free in \mathbf{p}_1 is modified by \vec{c}

↑ THE WORLD OF STRONG HEAPS

 ↓ THE WORLD OF WEAK HEAPS

$$\frac{\Psi, \Gamma \vdash e : \text{wref } \tau}{\Psi \vdash \{\Gamma, \text{emp}\} x := [e] \{\Gamma[x \rightsquigarrow \tau], \text{emp}\}} \text{ (W-LOAD)}$$

$$\frac{\Psi, \Gamma \vdash x : \text{wref } \tau \quad \Psi, \Gamma \vdash e : \tau}{\Psi \vdash \{\Gamma, \text{emp}\} [x] := e \{\Gamma, \text{emp}\}} \text{ (W-UPDATE)}$$

$$\frac{\Psi, \Gamma \vdash e : \tau}{\Psi \vdash \{\Gamma, \text{emp}\} x := \text{alloc}(e) \{\Gamma[x \rightsquigarrow \text{wref } \tau], \text{emp}\}} \text{ (W-ALLOC)}$$

Fig. 5. Rules for commands (Rules for assignments, conditional statements, loops, and sequencing are the same as the ones in Hoare Logic and are omitted.)

ditions are Γ' and \mathbf{p}' ; they specify conditions on local variables and the strong heap of the state after executing \vec{c} . Readers may wonder why there is no postcondition specification of the weak heap. As common in mutable-reference type systems, the implicit semantics of the judgment is that there exists an extended heap type $\Psi' \supseteq \Psi$ and the weak heap of the poststate should satisfy Ψ' . In terms of type checking, the particular Ψ' does not matter. The formal semantics of the judgment will be presented in Section 3.

Rules in Figure 5 are divided into two groups. One group is for the world of strong heaps, and another for the world of weak heaps. The rules for strong heaps are almost the same as the corresponding ones in standard SL, except that they also update Γ when necessary.

The rules for weak heaps are the ones that one would usually find in a type system for mutable-reference types. The weak-update rule W-UPDATE requires that the

$$\text{Operational semantics: } (\mathbf{r}, h, \text{s2w}(x) \cdot \vec{c}) \mapsto (\mathbf{r}, h, \vec{c})$$

$$\text{Rule: } \frac{\Psi, \Gamma \vdash x : \text{ref} \quad \Psi, \Gamma \vdash e : \tau}{\Psi \vdash \{\Gamma, \{x \mapsto e\}\} \text{s2w}(x) \{\Gamma[x \rightsquigarrow \text{wref } \tau], \text{emp}\}} \text{s2w}$$

Fig. 6. Rule for converting a location from the strong heap to the weak heap

$$\frac{\vdash \{\Gamma'_1, \mathbf{p}'_1\} \Rightarrow \{\Gamma_1, \mathbf{p}_1\} \quad \Psi \vdash \{\Gamma_1, \mathbf{p}_1\} \vec{c} \quad \Gamma_2, \mathbf{p}_2 \vdash \{\Gamma_2, \mathbf{p}_2\} \Rightarrow \{\Gamma'_2, \mathbf{p}'_2\}}{\Psi \vdash \{\Gamma'_1, \mathbf{p}'_1\} \vec{c} \quad \Gamma'_2, \mathbf{p}'_2} \text{WEAKENING}$$

$$\boxed{\vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma', \mathbf{p}'\}}$$

$$\frac{}{\vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma, \mathbf{p} \wedge \{x : \Gamma(x)\}\}} \text{w1} \qquad \frac{\vdash \mathbf{p} \Rightarrow \mathbf{p}'}{\vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma, \mathbf{p}'\}} \text{w2}$$

Fig. 7. Weakening rules

pointer be of type “wref τ ”, and that the new value be of type τ . This rule enforces type-preserving updates. Once these conditions hold, Γ remains unchanged after the update. Notice in this rule there is no need to understand separation and aliases as the S-UPDATE rule does. The W-ALLOC rule does not need to extend the heap type Ψ because Ψ is only a precondition. When proving the soundness of the rule, we need to find a new Ψ' that extends Ψ and is also satisfied by the new weak heap after the allocation. Finally, there is no rule for $\text{free}(e)$ in the world of weak heaps. Weak heaps should be garbage-collected.¹

Figures 6 and 7 present some rules that show the interaction between weak and strong heaps. Figure 6 adds a new instruction “s2w(x)” for converting a location from a strong heap to a weak heap. Operationally, this instruction is a no-op (so it is an annotation, rather than a “real” instruction). Its typing rule, however, involves transforming the ownership in the strong heap to a pointer of weak-reference types. Notice that there is no rule for converting a location from the weak heap to the strong heap; this is similar to deallocation in weak heaps and requires the help of garbage collectors.

Figure 7 presents weakening rules. Rule w1 converts type information in Γ to information in assertion \mathbf{p} . This is useful since information in Γ might be overwritten due to assignments to variables. One of examples in later sections will show the use of this rule. Rule w2 uses the premise $\vdash \mathbf{p} \Rightarrow \mathbf{p}'$; any valid SL formula $\mathbf{p} \Rightarrow \mathbf{p}'$ is acceptable.

2.3 Examples

We now show a few examples that demonstrate the use of SL^{W} . In these examples, we assume an additional type `even` for even integers. For clarity, we will also annotate

¹ We do not formally consider the interaction between garbage collectors and weak heaps. When considering a garbage collector, SL^{W} has to build in an extra level of indirection for cross-boundary references from strong heaps to weak heaps as objects in weak heaps may get moved (this is how the JNI implements Java references in native code). We leave this as future work.

the allocation instruction to indicate whether the allocation happens in the strong heap or in the weak heap. We write $x := \text{alloc}_s(e)$ for a strong-heap allocation. We write $x := \text{alloc}_{w,\tau}(e)$ for a weak-heap allocation, and the intended type for e is τ . These annotations help in guiding the type checking of SL^W .

The first example shows how the counterexample in the introduction (formula (2) on page 2) plays out in SL^W . The following program first initializes the heap to a form such that y points to a location of type “wref even” and x points to 4, and then performs a heap update through x . The whole program is checkable in SL^W with respect to any heap type (remember the heap type specifies the *initial* weak heap). Below we also include conditions of the form “ Γ, p ” between instructions.

$$\begin{array}{l}
\{\}, \text{emp} \\
z := \text{alloc}_{w,\text{even}}(2) \\
\{z : \text{wref even}\}, \text{emp} \\
y := \text{alloc}_s(z) \\
\{y : \text{ref}, z : \text{wref even}\}, \{y \mapsto z\} \quad \text{by rule (w1)} \\
\{y : \text{ref}, z : \text{wref even}\}, \{y \mapsto z\} \wedge \{z : \text{wref even}\} \quad \text{by rule (w2)} \\
\{y : \text{ref}, z : \text{wref even}\}, \exists v. \{y \mapsto v\} \wedge \{v : \text{wref even}\} \\
z := 0 \\
\{y : \text{ref}, z : \text{int}\}, \exists v. \{y \mapsto v\} \wedge \{v : \text{wref even}\} \\
x := \text{alloc}_s(4) \\
\{x : \text{ref}, y : \text{ref}, z : \text{int}\}, \exists v. (\{y \mapsto v\} \wedge \{v : \text{wref even}\}) * \{x \mapsto 4\} \\
[x] := 3 \\
\{x : \text{ref}, y : \text{ref}, z : \text{int}\}, \exists v. (\{y \mapsto v\} \wedge \{v : \text{wref even}\}) * \{x \mapsto 3\}
\end{array}$$

Different from the counterexample, the condition before “[x] := 3” limits where y can point to. In particular, y cannot point to x because (1) by the type of v , variable y must point to a weak-heap location; (2) x represents a location in the strong heap. Therefore, the update through x does not invalidate the type of v . We could easily construct an example where y indeed points to x . But in that case the type of v would be `ref`, which would also not be affected by updates through x .

One of the motivations of SL^W is to reason about programs where code in high-level languages interacts with low-level code. Prior research [5, 15] has shown that it is error prone when high-level code interoperates with low-level code. All kinds of errors may occur. One common kind of errors occurs when low-level code makes type misuses of references that point to objects in the weak heap. For instance, in the JNI, types of all references to Java objects are conflated into one type in native code—`jobject`. Consequently, there is no static checking of whether native code uses these Java references in a type-safe way. Type misuses of these Java references can result in silent memory corruption or unexpected behavior.

The first example already demonstrates how SL^W enables passing pointers from high-level to low-level code. In the example, the first allocation is on the weak heap and can be thought of as an operation by high-level code. Then, the location is passed to the low level by being stored in the strong heap. Unlike foreign function interfaces where types of cross-boundary references are conflated into a single type in low-level code, SL^W can track the accurate types of those references and enable type safety.

The next example demonstrates how low-level code can initialize a data structure in the strong heap, and then transfer that structure to the weak heap so that the structure is usable by high-level code.

```

    {}, emp
  x := allocs(4)
    {x : ref}, {x ↦ 4}
  y := allocs(x)
    {x : ref, y : ref}, {x ↦ 4} * {y ↦ x}
  s2w(x)
    {x : wref even, y : ref}, {y ↦ x}
  s2w(y)
    {x : wref even, y : wref (wref even)}, emp

```

3 Soundness of SL^W

Soundness of SL^W is proved by a semantic approach. We first describe a semantic model for weak-reference types. Based on this model, semantics of various concepts in SL^W are defined. Every rule in SL^W is then proved as a lemma according to the semantics.

3.1 Modeling weak-reference types

Intuitively, a type is a set of values. This suggests that a semantic type should be a predicate of the metatype “ $Value \rightarrow Prop$ ”. However, this idea would not support weak-reference types. To see why, let us examine a naïve model where “wref τ ” in a heap h would denote a set of locations ℓ such that $h(\ell)$ is of type τ . This simple model is unfortunately unsound, which is illustrated by the following example:

1. Create a reference of type “wref even”, and let the reference be x .
2. Copy x to y . By the naïve model, a reference of type “wref even” also has type “wref int” (because an even number is also an integer). Let “wref int” be the type of y .
3. Update the reference through y with an odd integer, say 3. As y has the type “wref int”, updating it with an odd integer is legal.
4. Dereference x . Alas, the dereference returns 3, although the type of x implies a result of an even number!

The problem with the naïve model is that, with aliases, it allows inconsistent views of memory. In the foregoing example, x and y have inconsistent views on the same memory cell. To address this problem, SL^W uses a heap type Ψ to type check a location. This follows the approach of Tofte [17] and Harper [6]. An example Ψ is as follows:

$$\Psi = \{\ell_0 : \text{even}, \ell_1 : \text{int}, \ell_2 : \text{wref even}, \ell_3 : \text{wref int}\} \quad (3)$$

A heap type Ψ helps to define two related concepts, informally stated below (their formal semantic definitions will be presented in a moment):

- (i) A location ℓ is of type “wref τ ” if and only if $\Psi(\ell)$ equals τ .
- (ii) A heap h is consistent with Ψ if for every ℓ , the value $h(\ell)$ has type $\Psi(\ell)$. For the example Ψ , it means that $h(\ell_0)$ should be an even number, $h(\ell_1)$ should be an integer, $h(\ell_2)$ should be of type “wref even”, ...

The heap type Ψ prevents aliases from having inconsistent views of the heap. Aliases have to agree on their types because the types have to agree with the type in Ψ . In particular, the example showing the unsoundness of the naïve model would not work in the above model because, in step 3 of the example, y cannot be cast from type “wref even” to “wref int”: type “wref even” implies that $\Psi(y) = \text{even}$, which is a different type from int .

A subtlety of the above model is the denotation of “wref τ ” depends on the heap type Ψ , but is *independent* of the heap h . A weak-reference type is connected to the heap h only indirectly, through the consistency relation between h and Ψ .

Example 3. Let $h = \{\ell_0 \mapsto 4, \ell_1 \mapsto 3, \ell_2 \mapsto \ell_0, \ell_3 \mapsto \ell_1\}$. It is consistent with the example Ψ in (3). To see this, 4 at location ℓ_0 is an even number and 3 at location ℓ_1 is an integer. At location ℓ_2 , ℓ_0 is of type “wref even” because, by (i), this is equivalent to $\Psi(\ell_0) = \text{even}$ —a true statement. Similarly, the value ℓ_1 at location ℓ_3 is of type “wref int”. \square

Formalizing a set of semantic predicates following (i) and (ii) directly, however, would encounter difficulties because of a circularity in the model: by (ii), Ψ is a map from locations to types; by (i), the model of types takes Ψ as an argument— Ψ is necessary to decide if a location belongs to “wref τ ”. If defined naïvely, the model would result in inconsistent cardinality, as described by Ahmed [1].

We next propose a fixed-point approach. We rewrite the heap type Ψ as a recursive equation. After adding Ψ as an argument to types, the example in (3) becomes:

$$\Psi = \{\ell_0 : \text{even}(\Psi), \ell_1 : \text{int}(\Psi), \ell_2 : (\text{wref even})(\Psi), \ell_3 : (\text{wref int})(\Psi)\} \quad (4)$$

Notice that Ψ appears on both the left and the right side of the equation. Once Ψ is written as a recursive equation, it follows that any fixed point of the following functional is a solution to the equation (4):

$$\lambda\Psi.\{\ell_0 : \text{even}(\Psi), \ell_1 : \text{int}(\Psi), \ell_2 : (\text{wref even})(\Psi), \ell_3 : (\text{wref int})(\Psi)\} \quad (5)$$

To get a fixed point of (5), we follow the indexed model of recursive types by Appel and McAllester [2]. We first introduce some domains:

$$\begin{aligned} (\text{SemHeapType}) \quad \mathbb{F} &\in \text{Loc} \rightarrow \text{SemIType} \\ (\text{SemIType}) \quad \mathbb{t} &\in \text{SemHeapEnv} \rightarrow \text{Nat} \rightarrow \text{Value} \rightarrow \text{Prop} \\ (\text{SemHeapEnv}) \quad \phi &\in \text{Loc} \rightarrow \text{Nat} \rightarrow \text{Value} \rightarrow \text{Prop} \end{aligned}$$

We use \mathbb{F} for a semantic heap type (it is the metatype of the denotation of heap types, as we will see). It maps locations to indexed types. An important point is that from \mathbb{F} we can define $\lambda\phi.\ell. \mathbb{F}(\ell) \phi$, which has the metatype $\text{SemHeapEnv} \rightarrow \text{SemHeapEnv}$.

Therefore, a semantic heap type is effectively a functional similar to the one in (5), and a fixed point of F is of the metatype $SemHeapEnv$.

A semantic type τ is a predicate over the following arguments: ϕ is a semantic heap environment; k is a natural-number index; v is a value. The heap environment $\phi \in SemHeapEnv$ is used in our model of $WRef(\tau)$ to constrain reference types. The index k comes from the indexed model and is a technical device that enables us to define the fixed point of a semantic heap type F .

Following the indexed model, we introduce a notion of contractiveness.

Definition 4. (*Contractiveness*)

$$\begin{aligned} \text{contractive}(F) &\triangleq \forall \ell \in \text{dom}(F). \text{contractive}(F(\ell)) \\ \text{contractive}(t) &\triangleq \forall \phi, k, j \leq k, v. (t \phi j v) \leftrightarrow (t (\text{approx}(k, \phi)) j v) \\ \text{approx}(k, \phi) &\triangleq \lambda \ell, j, v. j < k \wedge \phi \ell j v. \end{aligned}$$

We define $(\wp F) = \lambda \phi, \ell. F(\ell) \phi$. That is, it turns F into a functional of type $SemHeapEnv \rightarrow SemHeapEnv$.

Theorem 5. *If $\text{contractive}(F)$, then the following μF is the least fixed point² of the functional $(\wp F)$:*

$$\mu F \triangleq \lambda \ell, k, v. (\wp F)^{k+1}(\perp) \ell k v,$$

where $\perp = \lambda \ell, k, v. \text{false}$, and $(\wp F)^{k+1}$ applies the functional $k+1$ times.

The theorem is proved by following the indexed model of recursive types [2]. We present the proof in our technical report [16].

The following lemma is an immediate corollary of Theorem 5.

Lemma 6. *For any contractive F , any ℓ, k, v , we have $(F(\ell) (\mu F) k v) \leftrightarrow ((\mu F)(\ell) k v)$*

Most of the semantic types ignore the ϕ argument. For example,

$$\text{Even} \triangleq \lambda \phi, k, v. \exists u. v = 2 \times u.$$

We use capitalized `Even` to emphasize that it is a predicate, instead of the syntactic type `even`. The model of weak-reference types uses the argument ϕ .

Definition 7. $WRef(t) \triangleq \lambda \phi, k, \ell. \forall j < k, v. \phi \ell j v \leftrightarrow t \phi j v$

In words, a location ℓ is of type $WRef(\tau)$ under heap environment ϕ , if $\phi(\ell)$ equals τ approximately, with index less than k .

Example 8. Let $F_0 = \{\ell_0 : \text{Even}, \ell_1 : WRef(\text{Even})\}$. Then “ $WRef(\text{Even}) (\mu F_0) k \ell_0$ ” holds for any k . To see this, for any $j < k$ and v , we have

$$(\mu F_0) \ell_0 j v \leftrightarrow F_0(\ell_0)(\mu F_0) j v \leftrightarrow \text{Even} (\mu F_0) j v$$

The first step is by lemma 6, and the second is by the definition of F_0 at location ℓ_0 . We can similarly show “ $WRef(WRef(\text{Even})) (\mu F_0) k \ell_1$ ” holds. \square

Note that the definition of $WRef(\tau)$ is more general than the “`wref τ` ” type in SL^W , as τ is syntactically defined, while τ can be any (contractive) semantic predicate.

² Since F is contractive in the sense that “ $F(\ell) \phi k w$ ” performs only calls to ϕ on arguments smaller than k , it is easy to show by induction that any two fixed points of F are identical; therefore, the least fixed point of F is also its greatest fixed point.

Heap allocation. We need an additional idea to cope with heap allocation in the weak heap. Our indexed types take the fixed point of a semantic heap type F as an argument. But F changes after heap allocation. For example, from

$$F = \{\ell_0 : \text{Even}, \ell_1 : \text{WRef}(\text{Even})\} \text{ to } F' = \{\ell_0 : \text{Even}, \ell_1 : \text{WRef}(\text{Even}), \ell_2 : \text{Even}\},$$

after ℓ_2 is allocated and initialized with an even number.

After a new heap location is allocated, any value that has type τ before allocation should still have the same type after allocation. This is the monotonicity condition maintained by type systems. To model it semantically, our idea is to quantify explicitly outside of the model of types over all future semantic heap types and assert that the type in question is true over the fixed point of any future semantic heap type.

First is a semantic notion of type-preserving heap extension from F to F' :

Definition 9. $F' \geq F \triangleq$
 $\text{contractive}(F') \wedge \text{contractive}(F) \wedge \forall \ell \in \text{dom}(F), \phi, k, v. F'(\ell) \phi k v \leftrightarrow F(\ell) \phi k v$

Lemma 10. *The relation $F' \geq F$ is reflexive, anti-symmetric, and transitive (thus a partial order).*

Next, we define the consistency relation between h and F , and also a relation that states a value v is of type τ under F . Both relations quantify over all future semantic heap types, and require that the type in question be true over the fixed point of any future semantic heap type.

Definition 11. $\models h : F \triangleq \text{dom}(h) \subseteq \text{dom}(F) \wedge \forall \ell \in \text{dom}(h). F \models h(\ell) : F(\ell)$
 $F \models v : t \triangleq \forall F' \geq F. \forall k. t (\mu F') k v$

With our model, the following theorem for memory operations can be proved (please see our technical report [16] for proofs).

Theorem 12.

- (i) (*Read*) If $\models h : F$, and $\ell \in \text{dom}(h)$, and $F \models \ell : \text{WRef}(t)$, then $F \models h(\ell) : t$.
- (ii) (*Write*) If $\models h : F$, and $\ell \in \text{dom}(h)$, and $F \models \ell : \text{WRef}(t)$, and $F \models v : t$, then $\models h[\ell \rightsquigarrow v] : F$.
- (iii) (*Allocation*) If $\models h : F$, and $F \models v : t$, and $\text{contractive}(t)$, and $\ell \notin \text{dom}(F)$, then $\models h \uplus \{\ell \mapsto v\} : F \uplus \{\ell \mapsto t\}$.

3.2 Semantic model of SL^{W}

To show the soundness of SL^{W} , we define semantics for judgments in SL^{W} and then prove each rule as a lemma according to the semantics. Figure 8 presents definitions that are used in the semantics.

The semantics of types is unsurprising. In particular, the semantics of $\llbracket \text{wref } \tau \rrbracket$ is defined in terms of the predicate $\text{WRef}(t)$ in Definition 7. All these types are contractive. The semantics of Ψ and Γ is just the point-wise extension of the semantics of types.

The predicate “ $F, r, h \models p$ ” interprets the truth of assertion p . When p is a standard SL formula, the interpretation is the same as the one in SL . When p is $\{e : \tau\}$, the

$$\boxed{\llbracket \tau \rrbracket \in \text{SemType}}$$

$$\llbracket \text{int} \rrbracket \triangleq \lambda \phi, k, v. \exists n. v = n. \quad \llbracket \text{ref} \rrbracket \triangleq \lambda \phi, k, v. \exists \ell. v = \ell. \quad \llbracket \text{wref } \tau \rrbracket \triangleq \text{WRef}(\llbracket \tau \rrbracket)$$

$$\boxed{\llbracket \Psi \rrbracket \in \text{Loc} \rightarrow \text{SemType}}$$

$$\llbracket \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\} \rrbracket \triangleq \{\ell_1 : \llbracket \tau_1 \rrbracket, \dots, \ell_n : \llbracket \tau_n \rrbracket\}$$

$$\boxed{\llbracket \Gamma \rrbracket \in \text{Var} \rightarrow \text{SemType}}$$

$$\llbracket \{x_1 : \tau_1, \dots, x_n : \tau_n\} \rrbracket \triangleq \{x_1 : \llbracket \tau_1 \rrbracket, \dots, x_n : \llbracket \tau_n \rrbracket\}$$

$$\boxed{\mathbb{F}, \mathbf{r}, h \models \mathbf{p}}$$

$$\mathbb{F}, \mathbf{r}, h \models \{e : \tau\} \triangleq \mathbb{F} \models \mathbf{r}(e) : \llbracket \tau \rrbracket$$

$$\mathbb{F}, \mathbf{r}, h \models \text{emp} \triangleq \text{dom}(h) = \emptyset$$

$$\mathbb{F}, \mathbf{r}, h \models \{e_1 \mapsto e_2\} \triangleq \text{dom}(h) = \mathbf{r}(e_1) \wedge h(\mathbf{r}(e_1)) = \mathbf{r}(e_2)$$

$$\mathbb{F}, \mathbf{r}, h \models \mathbf{p}_1 * \mathbf{p}_2 \triangleq \exists h_1, h_2. (h = h_1 \uplus h_2) \wedge (\mathbb{F}, \mathbf{r}, h_1 \models \mathbf{p}_1) \wedge (\mathbb{F}, \mathbf{r}, h_2 \models \mathbf{p}_2)$$

$$\mathbb{F}, \mathbf{r}, h \models \mathbf{p}_1 \multimap \mathbf{p}_2 \triangleq \forall h_1. ((\text{dom}(h_1) \cap \text{dom}(h) = \emptyset) \wedge (\mathbb{F}, \mathbf{r}, h_1 \models \mathbf{p}_1)) \Rightarrow (\mathbb{F}, \mathbf{r}, h_1 \uplus h \models \mathbf{p}_2)$$

$$\mathbb{F} \models \mathbf{r} : \Gamma \triangleq \forall x \in \text{dom}(\Gamma). \mathbb{F} \models \mathbf{r}(x) : \llbracket \Gamma(x) \rrbracket$$

$$\mathbf{r}, h \models \mathbb{F} * \mathbf{p} \triangleq \exists h_1, h_2. (h = h_1 \uplus h_2) \wedge (\text{dom}(h_1) = \text{dom}(\mathbb{F})) \wedge (h_1 \models \mathbb{F}) \wedge (\mathbb{F}, \mathbf{r}, h_2 \models \mathbf{p})$$

Fig. 8. Semantic definitions

interpretation depends on \mathbb{F} . Notice that the interpretation of $\{e : \tau\}$ is independent of the heap; it is a pure assertion (that is, it does not depend on the strong heap).

The definition of $\mathbb{F} \models \mathbf{r} : \Gamma$ is the point-wise extension of $\mathbb{F} \models v : t$ to local variable types. The definition of “ $\mathbf{r}, h \models \mathbb{F} * \mathbf{p}$ ” splits the heap into two parts. One for the weak heap, which should satisfy \mathbb{F} , and the other for the strong heap, which is specified by \mathbf{p} .

With the above definitions, we are ready to define the semantics of the judgments in SL^{W} . The following definitions interpret “ $\Psi, \Gamma \vdash e : \tau$ ”, “ $\vdash \mathbf{p} \Rightarrow \mathbf{p}'$ ”, and “ $\vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma', \mathbf{p}'\}$ ”.

Definition 13.

$$\Psi, \Gamma \vdash e : \tau \triangleq \forall \mathbb{F} \geq \llbracket \Psi \rrbracket. \forall \mathbf{r}. \mathbb{F} \models \mathbf{r} : \Gamma \Rightarrow \mathbb{F} \models \mathbf{r}(e) : \llbracket \tau \rrbracket.$$

$$\vdash \mathbf{p} \Rightarrow \mathbf{p}' \triangleq \forall \mathbb{F}, \mathbf{r}, h. (\mathbb{F}, \mathbf{r}, h \models \mathbf{p}) \Rightarrow (\mathbb{F}, \mathbf{r}, h \models \mathbf{p}')$$

$$\begin{aligned} \vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma', \mathbf{p}'\} &\triangleq \\ &\forall \mathbb{F}, \mathbf{r}, h. (\mathbb{F} \models \mathbf{r} : \Gamma \wedge \mathbf{r}, h \models \mathbb{F} * \mathbf{p}) \Rightarrow (\mathbb{F} \models \mathbf{r} : \Gamma' \wedge \mathbf{r}, h \models \mathbb{F} * \mathbf{p}') \end{aligned}$$

Now we are ready to interpret $\Psi \vdash \{\Gamma, \mathbf{p}\} \vec{c} \{\Gamma', \mathbf{p}'\}$. Following Hoare Logic, we define both partial and total correctness:

Definition 14. (*Partial and total correctness*)

$$\begin{aligned} \Psi \models_p \{\Gamma, p\} \vec{c} \{\Gamma', p'\} &\triangleq \\ &\forall F \geq \llbracket \Psi \rrbracket, r, h. ((F \models r : \Gamma) \wedge (r, h \models F * p)) \Rightarrow \\ &\text{safe}(r, h, \vec{c}) \wedge \\ &(\forall r', h'. (r, h, \vec{c}) \mapsto^* (r', h', \varepsilon) \Rightarrow \exists F' \geq F. (F' \models r' : \Gamma') \wedge (r', h' \models F' * p')) \\ \Psi \models_t \{\Gamma, p\} \vec{c} \{\Gamma', p'\} &\triangleq \\ (\Psi \models_p \{\Gamma, p\} \vec{c} \{\Gamma', p'\}) \wedge & \\ (\forall F \geq \llbracket \Psi \rrbracket, r, h. ((F \models r : \Gamma) \wedge (r, h \models F * p)) \Rightarrow \text{terminate}(r, h, \vec{c})) & \end{aligned}$$

In the partial-correctness interpretation, it assumes a state that satisfies the condition $\{\Gamma, p\}$ and requires that the state be safe (see Definition 2 on page 4 for safety). In addition, it requires that, for any terminal state after the execution of \vec{c} , we must be able to find a new semantic heap type F' so that $F' \geq F$ and the new state satisfies $\{\Gamma', p'\}$. Note that F' may be larger than F due to allocations in \vec{c} . The total-correctness interpretation requires termination in addition to the requirements of partial correctness.

Theorem 15. *All rules in Figures 5, 6 and 7 are sound for both partial and total correctness.*

The proof uses Theorem 12. It is largely straightforward and omitted. We refer interested readers to our technical report [16] for the proof.

4 Related work

We discuss related work in three categories: (1) work related to language interoperation; (2) work related to integrating SL with type systems; and (3) work related to semantic models of types.

Most work in language interoperation focuses on the design and implementation of foreign function interfaces. Examples are plenty. Given a multilingual program, one natural question is how to reason about the program as a whole. This kind of reasoning requires models, program analyzers, and program logics that can work across language boundaries. Previous work has addressed the question of how to model the interoperation between dynamically typed languages and statically typed languages [10], and the interoperation between two safe languages when they have different systems of computational effects [19]. By integrating SL and type systems, SL^W can elegantly reason about properties of heaps that are shared by high-level and low-level code.

Previous systems of integrating SL with type systems [12, 9] assume that programs are well-typed according to a syntactic type system, and SL is then used as an add-on to reason about more properties of programs. Honda *et al*'s program logic [8, 21] for higher-order languages supports reference types but also requires a separate type system (in addition to the Hoare assertions); Reus *et al* [13] presented an extension of separation logic for supporting higher-order store (i.e., references to higher-order functions), but their logic does not support weak heaps which we believe embodies the key feature of reference types (i.e., the ability to perform safe updates without knowing the

exact aliasing relation). Compared to previous systems, SL^W targets the interoperation between high-level and low-level code. It allows cross-boundary references and mixes SL formulas and types.

The soundness of SL^W is justified by defining a semantic model, notably for types. Ahmed [1] and Appel *et al* [3] presented a powerful index-based semantic model for a rich type system with ML-style references. They rely on constructing a “dependently typed” global heap type to break the circularity discussed in Section 3. Our current work, in contrast, simply takes a fixed point of the recursively defined heap type predicate and avoids building any dependently typed data structures. Our work also differs from theirs in that we are reasoning about reference types in a program logic. Appel *et al.* [3] can also support impredicative polymorphism which is not addressed in our current work. Birkedal *et al* [4] recently presented a category-theoretic model that accommodates reference types as well as impredicative polymorphism. Similar to our model, their model also finds a fixed point and there is no need to work with approximation information. On the other hand, it appears that an implementation of their model in Coq requires a stratification of types and the use of dependent types, which our model avoids.

5 Discussion and future work

This work aims toward a framework for reasoning about language interoperation, but a lot remains to be done. A realistic high-level language contains many more language features and types. We do not foresee much difficulty in incorporating language features and types at the logic level as their modeling is largely independent from the interaction between weak and strong heaps. One technical concern is how to extend our semantic model to cover a complicated type system, including function types and OO classes.

SL^W does not formally consider the effect of a garbage collector. A garbage collector would break the crucial monotonicity condition of the weak heap that our semantic model relies on. We believe a possible way to overcome this problem is to use a region-based type system [18]. A garbage collector would also imply that there cannot be direct references from strong heaps to weak heaps; an extra level of indirection has to be added.

6 Conclusion

In his survey paper of Separation Logic [14], Reynolds asked “*whether the dividing line between types and assertions can be erased*”. This paper adds evidence that the type-based approach has its unique place when ensuring safety in weak heaps and when reasoning about the interaction between weak and strong heaps. The combination of types and SL provides a powerful framework for checking safety and verifying properties of multilingual programs.

Acknowledgments

We thank anonymous referees for suggestions and comments on an earlier version of this paper. Gang Tan is supported in part by NSF grant CCF-0915157. Zhong Shao is supported in part by a gift from Microsoft and NSF grants CCF-0524545 and CCF-0811665. Xinyu Feng is supported in part by NSF grant CCF-0524545 and National Natural Science Foundation of China (grant No. 90818019)

References

- [1] A. J. Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, 2004.
- [2] A. W. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. on Prog. Lang. and Sys.*, 23(5):657–683, 2001.
- [3] A. W. Appel, P.-A. Mellies, C. D. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *POPL '07*, pages 109–122. ACM Press, Jan. 2007.
- [4] L. Birkedal, K. Støvring, and J. Thamsborg. Realizability semantics of parametric polymorphism, references, and recursive types. In *FoSSaCS*, pages 456–470. Springer-Verlag, April 2009.
- [5] M. Furr and J. S. Foster. Checking type safety of foreign function calls. *ACM Trans. Program. Lang. Syst.*, 30(4):1–63, 2008.
- [6] R. Harper. A simplified account of polymorphic references. *Information Processing Letters*, 57(1):15–16, 1996.
- [7] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):578–580, October 1969.
- [8] K. Honda, N. Yoshida, and M. Berger. An observationally complete program logic for imperative higher-order frame rules. In *LICS '05*, pages 270–279, June 2005.
- [9] N. Krishnaswami, L. Birkedal, J. Aldrich, and J. Reynolds. Idealized ML and its separation logic. Unpublished manuscript, July 2007.
- [10] J. Matthews and R. B. Findler. Operational semantics for multi-language programs. In *Proc. 34th ACM Symp. on Principles of Prog. Lang.*, pages 3–10, 2007.
- [11] P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Computer Science Logic*, pages 1–19, 2001.
- [12] M. Parkinson. *Local reasoning for Java*. PhD thesis, University of Cambridge Computer Laboratory, Oxford, Nov. 2005. Tech Report UCAM-CL-TR-654.
- [13] B. Reus and J. Schwinghammer. Separation logic for higher-order store. In *20th International Workshop on Computer Science Logic (CSL)*, pages 575–590, 2006.
- [14] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. LICS'02*, pages 55–74, July 2002.
- [15] G. Tan and J. Croft. An empirical security study of the native code in the JDK. In *17th Usenix Security Symposium*, pages 365–377, 2008.
- [16] G. Tan, Z. Shao, X. Feng, and H. Cai. Weak updates and separation logic. <http://www.cse.lehigh.edu/~gtan/paper/WUSL-tr.pdf>, June 2009.
- [17] M. Tofte. Type inference for polymorphic references. *Inf. and Comp.*, 89(1):1–34, 1990.
- [18] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [19] V. Trifonov and Z. Shao. Safe and principled language interoperation. In *8th European Symposium on Programming (ESOP)*, pages 128–146, 1999.
- [20] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

- [21] N. Yoshida, K. Honda, and M. Berge. Logical reasoning for higher-order functions with local state. In H. Seidl, editor, *FoSSaCS*, pages 361–377, March 2007.