

Weak updates and separation logic

Gang Tan¹, Zhong Shao², Xinyu Feng³, Hongxu Cai⁴

¹*Lehigh University*, ²*Yale University*,

³*University of Science and Technology of China*, ⁴*Google Inc.*

Abstract Separation logic provides a simple but powerful technique for reasoning about low-level imperative programs that use shared data structures. Unfortunately, separation logic 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 separation logic when reasoning about the interaction between many high-level languages (e.g., ML, Java, C#) and low-level ones since the 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 separation logic with reference types and elegantly reasons about the interaction between strong and weak updates. We describe a semantic framework for reference types, which is used to prove the soundness of SL^W . Finally, we show how to extend SL^W with concurrency.

Keywords Weak Updates, Separation Logic, Type Systems, Mutable References, Language Interoperation

§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 ^{12, 17} (abbreviated to SL hereafter) 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:

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

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 in 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 article 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 updates, one natural thought to mix strong and weak updates is to extend SL with types so that assertions can also describe values in 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, when general reference types are allowed in “ $e \mapsto \tau$ ”, care must be taken to avoid unsoundness. An example demonstrating this point 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), assuming “ $e \mapsto \tau$ ” is a valid assertion. The precondition states that y points to a heap cell whose contents are of type “even ref”. If “ τ ref” is interpreted as a set of locations whose contents are of type τ , then 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 the model of “even ref” does not permit y to 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 approach²⁴⁾, where types are treated as syntax. Following the tradition of Hoare Logic, SL’s soundness is proved through a semantic model, and SL assertions are interpreted semantically. There is a need to resolve the differences between syntactic and semantic soundness proofs.

In this article, 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:

- 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. It accommodates cross-boundary pointers (from weak to strong heaps and vice versa) by distinguishing between pointers to weak heaps and pointers to strong heaps through the type system. SL^{W} is presented in Section 2.
- To resolve the differences 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.

- We extend SL^{W} with concurrency in Section 4. We add an atomic-block command and a parallel-execution command to the language, extend SL^{W} with new rules for concurrency, and adjust the semantic model for soundness. The extension to the logic demonstrates an important difference between weak and strong heaps: weak heaps can be accessed safely without synchronization, while safe access to strong heaps requires synchronization.

A preliminary version of this article was published in the Proceedings of the Seventh Asian Symposium on Programming Languages and Systems (APLAS 2009)¹⁹. The differences between the conference version and this article is described as follows. First, we streamlined the presentation and added proofs to the most important lemmas and theorems. Second, in this article we show how to extend the syntax and semantics of SL^{W} to accommodate concurrency. The change to a concurrent setting requires some changes to the original SL^{W} . For instance, there is a need to treat variables as resources (similar to the way how memory cells are treated in separation logic). We reuse local variable types for this purpose.

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

We next describe SL^{W} , an extension of SL that incorporates reasoning about 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.

Notation convention. 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 , we write $f_1 \perp f_2$ when their domains are disjoint. The notation $f_1 \uplus f_2$ is the union of f_1 and f_2 when $f_1 \perp f_2$, and undefined otherwise. The notation $f_1 \upharpoonright f_2$, implicitly carries the restriction $f_1 \perp f_2$. We write $f \setminus X$ for a new map resulting from removing elements in the set X from the domain of f .

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 Hoare⁸),

(Command) $c ::= \dots \mid x := [e] \mid [x] := e \mid x := \text{alloc}(e) \mid \text{free}(e) \mid c_1; c_2 \mid \epsilon$
(Expression) $e ::= x \mid v \mid \text{op}(e_1, \dots, e_n)$
(Value) $v ::= n$

Fig. 1 Language syntax

augmented with a set of commands for manipulating heap data structures. It is similar to the one used in Reynolds' presentation of SL¹⁷). Informally, the command “ $x := [e]$ ” loads the contents at location e into variable x ; “ $[x] := e$ ” updates the location at x with the value of e ; “ $x := \text{alloc}(e)$ ” allocates a new location, initializes it with the value of e , and assigns the new location to x ; “ $\text{free}(e)$ ” deallocates the location e . “ $c_1; c_2$ ” is the usual syntax for sequencing two commands. ϵ is the empty command, used in the operational semantics to signal the termination of a program.

An expression is either a variable x , a value v , or an operator op with a list of operands. We assume there is an infinite number of variables. All values, including heap locations, are natural numbers. Treating heap locations as natural numbers allows address arithmetic (i.e., adding a number to a location). We sometimes write ℓ for a heap location, but it is just a number.

Figure 2 presents the formal operational semantics of the language. A state consists of a store s (a map from variables to values), a heap h (a map from locations to values), and a command. Commands bring one state to another state and their semantics is formally defined by a step relation \longrightarrow . We write \longrightarrow^* for the reflexive and transitive closure of \longrightarrow .

A state may step to a **wrong** state. For instance, a state whose next instruction to execute is $[x] := e$ goes to a **wrong** state when its store does not contain one of the variables in $\{x\} \cup \text{fvar}(e)$ or the value of x is not in the domain of its heap.

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 (the semantics is formally defined in Section 3).

SL^W is equipped with a simple type system for tracking weak-heap locations. Although the type system does not include many types in high-level languages, by including reference types it is already sufficient to show interest-

$$\begin{array}{l}
\text{(State)} \quad \sigma ::= (s, h, c) \\
\text{(Store)} \quad s ::= \{x_1 \rightsquigarrow v_1, \dots, x_n \rightsquigarrow v_n\} \\
\text{(Heap)} \quad h ::= \{\ell_1 \rightsquigarrow v_1, \dots, \ell_n \rightsquigarrow v_n\}
\end{array}$$

$$\frac{s(e) = \ell \in \text{dom}(h)}{(s, h, x := [e]) \longrightarrow (s[x \rightsquigarrow h(\ell)], h, \epsilon)} \quad \frac{\{x\} \cup \text{fvar}(e) \not\subseteq \text{dom}(s) \text{ or } s(e) \notin \text{dom}(h)}{(s, h, x := [e]) \longrightarrow \text{wrong}}$$

$$\frac{s(x) = \ell \in \text{dom}(h)}{(s, h, [x] := e) \longrightarrow (s, h[\ell \rightsquigarrow s(e)], \epsilon)} \quad \frac{\{x\} \cup \text{fvar}(e) \not\subseteq \text{dom}(s) \text{ or } s(x) \notin \text{dom}(h)}{(s, h, [x] := e) \longrightarrow \text{wrong}}$$

$$\frac{\ell \notin \text{dom}(h)}{(s, h, x := \text{alloc}(e)) \longrightarrow (s[x \rightsquigarrow \ell], h \uplus \{\ell \rightsquigarrow s(e)\}, \epsilon)}$$

$$\frac{\{x\} \cup \text{fvar}(e) \not\subseteq \text{dom}(s)}{(s, h, x := \text{alloc}(e)) \longrightarrow \text{wrong}}$$

$$\frac{s(e) = \ell \in \text{dom}(h)}{(s, h, \text{free}(e)) \longrightarrow (s, h \setminus \{\ell\}, \epsilon)} \quad \frac{\text{fvar}(e) \not\subseteq \text{dom}(s) \text{ or } s(e) \notin \text{dom}(h)}{(s, h, \text{free}(e)) \longrightarrow \text{wrong}}$$

$$\frac{(s, h, c_1) \longrightarrow (s', h', c'_1)}{(s, h, c_1; c_2) \longrightarrow (s', h', c'_1; c_2)} \quad \frac{}{(s, h, \epsilon; c_2) \longrightarrow (s, h, c_2)}$$

$$\frac{(s, h, c_1) \longrightarrow \text{wrong}}{(s, h, c_1; c_2) \longrightarrow \text{wrong}}$$

$$\text{where } s(e) = \begin{cases} s(x) & \text{when } e = x \text{ and } x \in \text{dom}(s) \\ v & \text{when } e = v \\ \text{op}(s(e_1), \dots, s(e_n)) & \text{when } e = \text{op}(e_1, \dots, e_n) \end{cases}$$

Fig. 2 Operational semantics

ing interactions between strong and weak heaps. Reference types are the most common types when high-level languages interoperate with low-level languages because most data across language boundaries are passed by references.

Type `int` is for all numbers and it is essentially a top type since all values in SL^{W} are natural numbers. Type “`wref τ` ” is for locations in a weak heap, but not in a strong heap. We have removed the `ref` type in the APLAS version ¹⁹. The `ref` type was the type for heap locations. But since in this paper locations are treated as natural numbers, it is safe to remove `ref` and give `int` to heap locations. 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 a heap type Ψ , location ℓ has

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

Fig. 3 Assertions and types

$\Psi, \Gamma \vdash e : \tau$

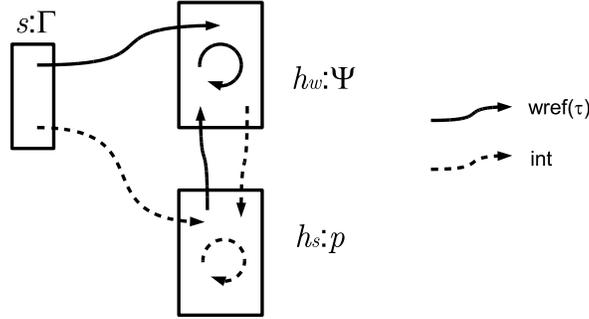
$$\begin{array}{c}
\frac{x \in \textit{dom}(\Gamma)}{\Psi, \Gamma \vdash x : \Gamma(x)} \qquad \frac{}{\Psi, \Gamma \vdash n : \textit{int}} \\
\frac{\Psi(\ell) = \tau}{\Psi, \Gamma \vdash \ell : \textit{wref} \tau} \qquad \frac{\forall i \in [1..n]. \Psi, \Gamma \vdash e_i : \textit{int}}{\Psi, \Gamma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

Fig. 4 Typing rules for expressions

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 \textit{SL}^W :



As shown in the diagram, \textit{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 ”. Pointers to strong heap cells (in dotted lines) can have only

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

$$\frac{x \in \text{dom}(\Gamma) \quad \Psi, \Gamma \vdash e : \text{int} \quad \Psi, \Gamma \vdash e' : \tau}{\Psi \vdash \{\Gamma.(e \mapsto e')\} x := [e] \{\Gamma[x \rightsquigarrow \tau].x = e' \wedge (e \mapsto x)\}} \text{ (S-LOAD)}$$

where $x \notin \text{fvar}(e) \cup \text{fvar}(e')$

$$\frac{\Psi, \Gamma \vdash x : \text{int} \quad \text{fvar}(e) \subseteq \text{dom}(\Gamma)}{\Psi \vdash \{\Gamma.(x \mapsto -)\} [x] := e \{\Gamma.(x \mapsto e)\}} \text{ (S-UPDATE)}$$

$$\frac{\{x\} \cup \text{fvar}(e) \subseteq \text{dom}(\Gamma)}{\Psi \vdash \{\Gamma.\text{emp}\} x := \text{alloc}(e) \{\Gamma[x \rightsquigarrow \text{int}].(x \mapsto e)\}} \text{ (S-ALLOC)}$$

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

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

Fig. 5 Rules for heap-manipulating commands in the world of strong heaps

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

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

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

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

Fig. 6 Rules for heap-manipulating commands in the world of weak heaps

type `int`.

Figures 5, 6, and 7 present inference rules for checking commands. These rules use the judgment $\Psi \vdash \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$. In this judgment, Ψ , Γ and \mathbf{p} are preconditions and specify conditions on the weak heap, local variables, and the strong heap respectively. Postconditions are Γ' and \mathbf{p}' ; they specify conditions on local variables and the strong heap of the state after executing 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.

In anticipation of the development for concurrency, rules for commands treat variables as resources. They satisfy the following lemma:

Lemma 2.1

If $\Psi \vdash \{\Gamma.p\} c \{\Gamma'.p'\}$, then $\text{fvar}(c) \subseteq \text{dom}(\Gamma)$, and $\text{dom}(\Gamma) = \text{dom}(\Gamma')$.

That is, all variables read and written by c should be in the domain of the local-variable type Γ and Γ has the same domain as Γ' . This is the reason why the S-LOAD rule requires $x \in \text{dom}(\Gamma)$. It also implicitly requires $\text{fvar}(e) \subseteq \text{dom}(\Gamma)$, which is implied by the assumption $\Psi, \Gamma \vdash e : \text{int}$. As another note, $\Gamma.p$ is equivalent to false when $\text{fvar}(p) \not\subseteq \text{dom}(\Gamma)$ (this is evident from the semantics of $\Gamma.p$ in Section 3).

Rules for commands are divided into two groups. One group is for the world of strong heaps (in Figure 5), and another for the world of weak heaps (in Figure 6). 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 the pointer be of type “wref τ ”, and 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

Figure 7 presents additional rules. In the FRAME rule, the notation $(\Gamma.p) * (\Gamma'.p')$ stands for $\Gamma \uplus \Gamma'.(p * p')$. Rule w1 converts type information in Γ to information in assertion 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 $\vdash p \Rightarrow p'$; any formula that is valid according to the semantics of $\vdash p \Rightarrow p'$ is acceptable (the semantics is defined in Sec. 3).

Figure 8 presents a rule that shows the interaction between weak and

*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.

$$\frac{\Psi \vdash \{\Gamma.\mathbf{p}\} c_1 \{\Gamma'.\mathbf{p}'\} \quad \Psi \vdash \{\Gamma'.\mathbf{p}'\} c_2 \{\Gamma''.\mathbf{p}''\}}{\Psi \vdash \{\Gamma.\mathbf{p}\} c_1; c_2 \{\Gamma''.\mathbf{p}''\}} \text{ (SEQ)}$$

$$\frac{}{\Psi \vdash \{\Gamma.\text{emp}\} \epsilon \{\Gamma.\text{emp}\}} \text{ (EMP)}$$

$$\frac{\Psi \vdash \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}}{\Psi \vdash \{(\Gamma.\mathbf{p}) * (\Gamma_1.\mathbf{p}_1)\} c \{(\Gamma'.\mathbf{p}') * (\Gamma_1.\mathbf{p}_1)\}} \text{ (FRAME)}$$

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

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

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

$$\frac{}{\vdash \{\Gamma.\mathbf{p}\} \Rightarrow \{\Gamma.(\mathbf{p} \wedge (x : \Gamma(x)))\}} \text{ (w1)}$$

$$\frac{\vdash \mathbf{p} \Rightarrow \mathbf{p}'}{\vdash \{\Gamma.\mathbf{p}\} \Rightarrow \{\Gamma.\mathbf{p}'\}} \text{ (w2)}$$

Fig. 7 Sequencing, frame, and weakening rules (Rules for assignments, conditional statements, and loops are similar to the ones in Hoare Logic and are omitted.)

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

Fig. 8 A rule for converting a location from the strong heap to the weak heap

strong heaps. The command “ $\epsilon_{\text{s2w}(x)}$ ” is the empty command with a type annotation. It instructs the type checker to transform the ownership of the cell referenced by x in the strong heap to the weak heap. 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.

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 and an obvious rule for the `even` type.

$$\frac{n \text{ is an even number}}{\Psi, \Gamma \vdash n : \text{even}}$$

For clarity, we will also annotate 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 guide the type checking of SL^W .

The first example shows how the counterexample in the introduction (formula (2) on page 3) 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 “ $\Gamma.p$ ” between instructions.

```

    {x : int, y : int, z : int}. emp
z := alloc_{w,even}(2);
    {x : int, y : int, z : wref even}. emp
y := alloc_s(z);
    {x : int, y : int, z : wref even}. (y ↦ z)    // by rule (w1)
    {x : int, y : int, z : wref even}. (y ↦ z) ∧ {z : wref even}    // by rule (w2)
    {x : int, y : int, z : wref even}. ∃v. (y ↦ v) ∧ (v : wref even)
z := 0;
    {x : int, y : int, z : int}. ∃v. (y ↦ v) ∧ (v : wref even)
x := alloc_s(4);
    {x : int, y : int, z : int}. ∃v. ((y ↦ v) ∧ (v : wref even)) * (x ↦ 4)
[x] := 3
    {x : int, y : int, z : int}. ∃v. ((y ↦ v) ∧ (v : wref even)) * (x ↦ 3)

```

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 `int`, 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^{6, 18)} has shown that the interface code between high-level programs and low-level programs is error-prone. 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.

$$\begin{aligned} & \{x : \text{int}, y : \text{int}\}. \text{emp} \\ x & := \text{alloc}_s(4); \\ & \{x : \text{int}, y : \text{int}\}. (x \mapsto 4) \\ y & := \text{alloc}_s(x); \\ & \{x : \text{int}, y : \text{int}\}. (x \mapsto 4) * (y \mapsto x) \\ & \epsilon_{\text{s2w}(x)}; \\ & \{x : \text{wref even}, y : \text{int}\}. (y \mapsto x) \\ & \epsilon_{\text{s2w}(y)} \\ & \{x : \text{wref even}, y : \text{wref (wref even)}\}. \text{emp} \end{aligned}$$

§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 “ $\text{Value} \rightarrow \text{Prop}$ ”. 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 named 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 dereferencing 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 Tofte²⁰⁾ and Harper⁷⁾. 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 , condition (ii) 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 3.1

Let $h = \{\ell_0 \rightsquigarrow 4, \ell_1 \rightsquigarrow 3, \ell_2 \rightsquigarrow \ell_0, \ell_3 \rightsquigarrow \ell_1\}$. 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:

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

We use \mathbf{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 \mathbf{F} we can define $\lambda\phi, \ell. \mathbf{F}(\ell) \phi$, 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 \mathbf{F} is of the metatype SemHeapEnv .

A semantic type \mathbf{t} 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 $\phi \in \text{SemHeapEnv}$ is used in our model of $\text{WRef}(\mathfrak{t})$ 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 \mathbf{F} .

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

Definition 3.1 (Contractiveness)

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

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

Theorem 3.1

If $\text{contractive}(\mathbf{F})$, then the following $\mu \mathbf{F}$ is the least fixed point^{*2} of the functional $(\wp \mathbf{F})$:

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

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

The theorem is proved by following the indexed model of recursive types²⁾.

The following lemma is an immediate corollary of Theorem 3.1.

Lemma 3.1

For any contractive \mathbf{F} , any ℓ, k, v , we have $(\mathbf{F}(\ell) (\mu \mathbf{F}) k v) \leftrightarrow ((\mu \mathbf{F})(\ell) k v)$

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 3.2

$$\text{WRef}(t) \triangleq \lambda \phi, k, \ell. \forall j < k, v. \phi \ell j v \leftrightarrow t \phi j v$$

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

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

Example 3.2

Let $F_0 = \{\ell_0 : \text{Even}, \ell_1 : \text{WRef}(\text{Even})\}$. Then “ $\text{WRef}(\text{Even}) (\mu F_0) k \ell_0$ ” 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 3.1, and the second is by the definition of F_0 at location ℓ_0 . We can similarly show “ $\text{WRef}(\text{WRef}(\text{Even})) (\mu F_0) k \ell_1$ ” holds. \square

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

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 $F = \{\ell_0 : \text{Even}, \ell_1 : \text{WRef}(\text{Even})\}$ 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 \mathfrak{t} 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 notion of type-preserving heap extension from F to F' :

Definition 3.3

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

Lemma 3.2

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 \mathfrak{t} 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 3.4

$$\begin{aligned} \models h:F &\triangleq \text{dom}(h) \subseteq \text{dom}(F) \wedge \forall \ell \in \text{dom}(h). F \models h(\ell):F(\ell) \\ F \models v:\mathfrak{t} &\triangleq \forall F' \geq F. \forall k. \mathfrak{t} (\mu F') k v \end{aligned}$$

With our model, the following theorem for memory operations can be proved.

Theorem 3.2

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

Proof We prove (i) and (iii); the proof of (ii) is similar.

(i) By the definition of $F \models h(\ell):\mathfrak{t}$, the goal is to prove “ $\mathfrak{t} (\mu F') k (h(\ell))$ ”, for all $F' \geq F$, and k .

By $\models h:F$ and $\ell \in \text{dom}(h)$, we have $F \models h(\ell):F(\ell)$. Unfold its definition and use the assumption $F' \geq F$, we get “ $F(\ell) (\mu F') k (h(\ell))$ ”. With this and the assumption $F' \geq F$, we derive “ $F'(\ell) (\mu F') k (h(\ell))$ ”, which is the same as the following by Lemma 3.1.

$$(\mu F') \ell k (h(\ell)). \tag{6}$$

Now by the premise $F \models \ell:\text{WRef}(\mathfrak{t})$, we derive “ $(\text{WRef}(\mathfrak{t})) (\mu F') (k + 1) \ell$ ”. By the definition of $\text{WRef}(\mathfrak{t})$, we further get

$$\forall v. (\mu F') \ell k v \leftrightarrow \mathfrak{t} (\mu F') k v \tag{7}$$

By (6) and (7), we get “ $\mathfrak{t} (\mu F') k (h(\ell))$ ”, which is our goal.

(iii) Let $h' = h \uplus \{\ell \rightsquigarrow v\}$, and $F' = F \uplus \{\ell \rightsquigarrow \mathfrak{t}\}$. To show $\models h':F'$, we

need to show

$$\text{dom}(h') \subseteq \text{dom}(\mathbf{F}'), \quad (8)$$

$$\forall \ell' \in \text{dom}(h'). \mathbf{F}' \models h'(\ell') : \mathbf{F}'(\ell'). \quad (9)$$

(8) is immediate from $\text{dom}(h) \subseteq \text{dom}(\mathbf{F})$. To prove (9), we show that for all $\ell' \in \text{dom}(h')$, for all $\mathbf{F}'' \geq \mathbf{F}'$, for all k , “ $\mathbf{F}'(\ell') (\mu\mathbf{F}'') k (h'(\ell'))$ ”. We prove it by a case analysis of ℓ' .

- (a) $\ell' \neq \ell$. The goal becomes “ $\mathbf{F}'(\ell') (\mu\mathbf{F}'') k (h'(\ell'))$ ”. First, it is easy to show $\mathbf{F}' \geq \mathbf{F}$. By this and $\mathbf{F}'' \geq \mathbf{F}'$, we get $\mathbf{F}'' \geq \mathbf{F}$. From $\models h : \mathbf{F}$, we get $\mathbf{F} \models h(\ell') : \mathbf{F}(\ell')$. By this and $\mathbf{F}'' \geq \mathbf{F}$, we derive the goal.
- (b) $\ell' = \ell$. The goal becomes “ $\mathbf{t} (\mu\mathbf{F}'') k v$ ”. This is immediate from $\mathbf{F} \models v : \mathbf{t}$ and $\mathbf{F}'' \geq \mathbf{F}$.

□

3.2 Semantic model of $\text{SL}^{\mathbf{W}}$

To show the soundness of $\text{SL}^{\mathbf{W}}$, we define semantics for judgments in $\text{SL}^{\mathbf{W}}$ and then prove each rule as a lemma according to the semantics. Figure 9 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 3.2. All these types are contractive. The semantics of Ψ and Γ is just the point-wise extension of the semantics of types.

The predicate “ $\mathbf{F}, s, h \models \mathbf{p}$ ” interprets the truth of assertion \mathbf{p} . When \mathbf{p} is a SL formula, the interpretation is similar to the one in SL . When \mathbf{p} is $\{e : \tau\}$, the interpretation depends on \mathbf{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 $\mathbf{F} \models s : \Gamma$ requires the domains of s and Γ match and every value in s has the specified type in Γ . The definition of “ $\models (s, h) \text{ sat } (\mathbf{F}, \Gamma, \mathbf{p})$ ” splits the heap into two parts. One for the weak heap, which should satisfy \mathbf{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 $\text{SL}^{\mathbf{W}}$. The following definitions interpret “ $\Psi, \Gamma \vdash e : \tau$ ”, “ $\vdash \mathbf{p} \Rightarrow \mathbf{p}'$ ”, and “ $\vdash \{\Gamma, \mathbf{p}\} \Rightarrow \{\Gamma', \mathbf{p}'\}$ ”.

$$\begin{array}{l}
\boxed{\llbracket \tau \rrbracket \in \text{SemIType}} \quad \llbracket \text{int} \rrbracket \triangleq \lambda \phi, k, v. \text{True.} \quad \llbracket \text{wref } \tau \rrbracket \triangleq \text{WRef}(\llbracket \tau \rrbracket) \\
\boxed{\llbracket \Psi \rrbracket \in \text{Loc} \rightarrow \text{SemIType}} \quad \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{SemIType}} \quad \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{\mathbf{F}, s, h \models \mathbf{p}} \\
\mathbf{F}, s, h \models \{e : \tau\} \triangleq \mathbf{F} \models s(e) : \llbracket \tau \rrbracket \\
\mathbf{F}, s, h \models \text{emp} \triangleq \text{dom}(h) = \emptyset \\
\mathbf{F}, s, h \models (e_1 \mapsto e_2) \triangleq \text{dom}(h) = s(e_1) \wedge h(s(e_1)) = s(e_2) \\
\mathbf{F}, s, h \models \mathbf{p}_1 * \mathbf{p}_2 \triangleq \exists h_1, h_2, s_1, s_2. \\
\quad (h = h_1 \uplus h_2) \wedge (s = s_1 \uplus s_2) \wedge (\mathbf{F}, s_1, h_1 \models \mathbf{p}_1) \wedge (\mathbf{F}, s_2, h_2 \models \mathbf{p}_2) \\
\mathbf{F}, s, h \models \mathbf{p}_1 - * \mathbf{p}_2 \triangleq \forall h_1. ((h_1 \perp h) \wedge (\{\}, s, h_1 \models \mathbf{p}_1)) \Rightarrow (\mathbf{F}, s, h_1 \uplus h \models \mathbf{p}_2) \\
----- \\
\mathbf{F} \models s : \Gamma \triangleq \text{dom}(s) = \text{dom}(\Gamma) \wedge \forall x \in \text{dom}(\Gamma). \mathbf{F} \models s(x) : \llbracket \Gamma(x) \rrbracket \\
\mathbf{F}, s, h \models \Gamma.\mathbf{p} \triangleq \text{fvar}(\mathbf{p}) \subseteq \text{dom}(\Gamma) \wedge (\mathbf{F} \models s : \Gamma) \wedge (\mathbf{F}, s, h \models \mathbf{p}) \\
\models (s, h) \text{ sat } (\mathbf{F}, \Gamma.\mathbf{p}) \triangleq \\
\quad \exists h_1, h_2. h = h_1 \uplus h_2 \wedge \text{dom}(h_1) = \text{dom}(\mathbf{F}) \wedge (\models h_1 : \mathbf{F}) \wedge (\mathbf{F}, s, h_2 \models \Gamma.\mathbf{p})
\end{array}$$

Fig. 9 Semantic definitions**Definition 3.5**

$$\begin{array}{l}
\Psi, \Gamma \models e : \tau \triangleq \text{fvar}(e) \subseteq \text{dom}(\Gamma) \wedge \forall \mathbf{F} \geq \llbracket \Psi \rrbracket. \forall s. \mathbf{F} \models s : \Gamma \Rightarrow \mathbf{F} \models s(e) : \llbracket \tau \rrbracket. \\
\models \mathbf{p} \Rightarrow \mathbf{p}' \triangleq \forall \mathbf{F}, s, h. (\mathbf{F}, s, h \models \mathbf{p}) \Rightarrow (\mathbf{F}, s, h \models \mathbf{p}') \\
\models \{\Gamma.\mathbf{p}\} \Rightarrow \{\Gamma'.\mathbf{p}'\} \triangleq \forall \mathbf{F}, s, h. (\mathbf{F}, s, h \models \Gamma.\mathbf{p}) \Rightarrow (\mathbf{F}, s, h \models \Gamma'.\mathbf{p}')
\end{array}$$

Now we are ready to interpret $\Psi \vdash \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$. Below we present a partial-correctness interpretation.

Definition 3.6 (Partial correctness)

$$\begin{array}{l}
\Psi \models \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\} \triangleq \\
\quad \forall \mathbf{F} \geq \llbracket \Psi \rrbracket, s, h. (\models (s, h) \text{ sat } (\mathbf{F}, \Gamma.\mathbf{p})) \Rightarrow \\
\quad \quad \neg((s, h, c) \longrightarrow^* \text{wrong}) \wedge \\
\quad \quad (\forall s', h'. (s, h, c) \longrightarrow^* (s', h', \epsilon) \Rightarrow \exists \mathbf{F}' \geq \mathbf{F}. \models (s', h') \text{ sat } (\mathbf{F}', \Gamma'.\mathbf{p}'))
\end{array}$$

In the interpretation, it assumes \mathbf{F} and a state that satisfies the condition $(\mathbf{F}, \Gamma.\mathbf{p})$

and requires that the state does not go wrong. In addition, it requires that, for any terminal state after the execution of c , we must be able to find a new semantic heap type F' so that $F' \geq F$ and the new state satisfies $(F', \Gamma'.p')$. Note that F' may be larger than F due to allocations in c .

Theorem 3.3 (Soundness)

If $\Psi \vdash \{\Gamma.p\} c \{\Gamma'.p'\}$, then $\Psi \models \{\Gamma.p\} c \{\Gamma'.p'\}$.

Proof The proof is by induction over the derivation of the assumption. We show the cases S-LOAD and W-ALLOC. Most other cases are similar. The case of the FRAME rule uses Lemma 3.3.

$$(i) \quad \frac{x \in \text{dom}(\Gamma) \quad \Psi, \Gamma \vdash e : \text{int} \quad \Psi, \Gamma \vdash e' : \tau}{\Psi \vdash \{\Gamma.(e \mapsto e')\} x := [e] \{\Gamma[x \rightsquigarrow \tau].x = e' \wedge (e \mapsto x)\}} \text{ (S-LOAD)}$$

where $x \notin \text{fvar}(e) \cup \text{fvar}(e')$

Pick $F \geq \llbracket \Psi \rrbracket, s, h$ and assume $\models (s, h) \text{ sat } (F, \Gamma.(e \mapsto e'))$. Therefore, for some h_1 and h_2 , we have

$$h = h_1 \uplus h_2 \wedge \text{dom}(h_1) = \text{dom}(F) \wedge \models h_1 : F \wedge F, s, h_2 \models \Gamma.(e \mapsto e') \quad (10)$$

The first subgoal is to show $(s, h, (x := [e]))$ does not step to a wrong state. It is sufficient to show “ $\{x\} \cup \text{fvar}(e) \subseteq \text{dom}(s)$ ” and “ $s(e) \in \text{dom}(h)$ ”. The first is true because “ $\{x\} \cup \text{fvar}(e) \subseteq \text{dom}(\Gamma)$ ” (by the assumptions of the S-LOAD rule), and “ $\text{dom}(s) = \text{dom}(\Gamma)$ ” (by $F, s, h_2 \models \Gamma.(e \mapsto e')$ in (10)).

The second subgoal is to show $\exists F' \geq F. \models (s', h) \text{ sat } (F', \Gamma'.p')$, where $s' = s[x \rightsquigarrow h(s(e))]$, $\Gamma' = \Gamma[x \rightsquigarrow \tau]$, and $p' = (x = e') \wedge (e \mapsto x)$. Let $F' = F$. Then by (10), we only need to show $F, s', h_2 \models \Gamma'.p'$, which is the same as

$$(\text{fvar}(p') \subseteq \text{dom}(\Gamma')) \wedge (F \models s' : \Gamma') \wedge (F, s', h_2 \models p')$$

All the above can be proved easily using the assumptions of the S-LOAD rule.

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

Pick $F \geq \llbracket \Psi \rrbracket, s, h$ and assume $\models (s, h) \text{ sat } (F, \Gamma.\text{emp})$. Therefore, we have

$$\text{dom}(h) = \text{dom}(F) \wedge \models h : F \wedge F, s, \{\} \models \Gamma.\text{emp} \quad (11)$$

The first subgoal is to show $(s, h, x := \text{alloc}(e))$ does not go to a wrong state. It is sufficient to show “ $\{x\} \cup \text{fvar}(e) \subseteq \text{dom}(s)$ ”. It is true because $x \in \text{dom}(\Gamma)$, $\text{fvar}(e) \subseteq \text{dom}(\Gamma)$ (from “ $\Psi, \Gamma \models e : \tau$ ”), and “ $\text{dom}(s) = \text{dom}(\Gamma)$ ”.

$$\begin{array}{c}
\text{(Command) } c ::= \dots \mid \text{atomic}\{c\} \mid (c_1 \parallel c_2) \\
\\
\frac{(s, h, c) \longrightarrow^* (s', h', \epsilon)}{(s, h, \text{atomic}\{c\}) \longrightarrow (s', h', \epsilon)} \quad \frac{(s, h, c) \longrightarrow^* \text{wrong}}{(s, h, \text{atomic}\{c\}) \longrightarrow \text{wrong}} \\
\\
\frac{(s, h, c_1) \longrightarrow (s', h', c'_1)}{(s, h, c_1 \parallel c_2) \longrightarrow (s', h', c'_1 \parallel c_2)} \quad \frac{(s, h, c_2) \longrightarrow (s', h', c'_2)}{(s, h, c_1 \parallel c_2) \longrightarrow (s', h', c_1 \parallel c'_2)} \\
\\
\frac{(s, h, c_1) \longrightarrow \text{wrong}}{(s, h, c_1 \parallel c_2) \longrightarrow \text{wrong}} \quad \frac{(s, h, c_2) \longrightarrow \text{wrong}}{(s, h, c_1 \parallel c_2) \longrightarrow \text{wrong}} \quad \frac{}{(s, h, \epsilon \parallel \epsilon) \longrightarrow (s, h, \epsilon)}
\end{array}$$

Fig. 10 Syntax and semantics of commands for concurrency.

The second subgoal is to show $\exists F' \geq F. \models (s', h') \text{ sat } (F', \Gamma'.\text{emp})$, where $s' = s[x \rightsquigarrow \ell]$, $h' = h \uplus \{\ell \rightsquigarrow s(e)\}$, and $\Gamma' = \Gamma[x \rightsquigarrow \text{wref } \tau]$.

Let $F' = F \uplus \{\ell \rightsquigarrow \llbracket \tau \rrbracket\}$; then $F' \geq F$. By Theorem 3.2(iii) and $F \models s(e) : \llbracket \tau \rrbracket$ (from “ $\Psi, \Gamma \models e : \tau$ ”), we can show $\models h' : F'$. It is also easy to show $F', s', \{\} \models \Gamma'.\text{emp}$. All these give us $\models (s', h') \text{ sat } (F', \Gamma'.\text{emp})$. \square

Lemma 3.3 (Locality)

If $s = s_1 \uplus s_2$, $h = h_1 \uplus h_2$, $\neg(s_1, h_1, c) \longrightarrow \text{wrong}$, then

1. $\neg(s, h, c) \longrightarrow \text{wrong}$; and
2. for all s' and h' , if $(s, h, c) \longrightarrow (s', h', c')$, then there exist s'_1 and h'_1 such that $s' = s'_1 \uplus s_2$, $h' = h'_1 \uplus h_2$, and $(s_1, h_1, c) \longrightarrow (s'_1, h'_1, c')$.

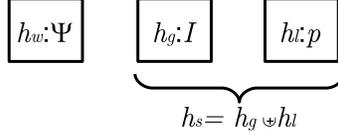
§4 Concurrency

In this section, we extend SL^W with concurrency. Figure 10 presents the syntax of extra commands for concurrency and their operational semantics. The atomic-block command $\text{atomic}\{c\}$ executes c atomically. Vafeiadis and Parkinson²³⁾ pointed out that such a command can be used to model synchronizations provided by locks or transactional memory. The operational semantics of $\text{atomic}\{c\}$ executes c in one step and is not interrupted by other threads. The parallel-composition command “ $c_1 \parallel c_2$ ” has the usual semantics of running c_1 and c_2 in parallel. It makes a step when either c_1 makes a step or c_2 makes a step. When both c_1 and c_2 reach the end (i.e., ϵ), the command “ $c_1 \parallel c_2$ ” terminates.

The judgment for checking commands is changed to

$$\Gamma_G, \mathbf{I}, \Psi \vdash_c \{\Gamma.p\} c \{\Gamma'.p'\}.$$

Compared to the sequential version, the concurrent one has an extra $\Gamma_c.I$ for specifying invariants in a shared strong heap. We use the symbol I for a global invariant that should always hold unless in an atomic block; invariant I is a separation-logic assertion and is required to be a *precise* assertion¹³). The following diagram depicts the relationship between various heaps and their specifications from the point view of one particular thread.



The weak heap h_w remains the same as that in the sequential setting and is specified by a heap type Ψ . The strong heap h_s , however, is conceptually split into a global one (h_g), shared by all threads, and a local one (h_l), private to one thread. The global one is specified by the global invariant I , and the local one by p .

Figure 11 presents inference rules for checking commands. The `ATOMIC` rule states that the global strong heap with invariant I is accessible for the command c in an atomic block. The command can temporarily break the invariant, but has to restore the invariant when it is finished. Since it is inside an atomic block, its temporary breaking of the invariant will not interfere with other threads.

The `PAR` rule conceptually breaks the store and the local strong heap into two disjoint parts, specified by $\Gamma_1.p_1$ and $\Gamma_2.p_2$, respectively. It then checks c_1 with $\Gamma_1.p_1$ as the precondition and c_2 with $\Gamma_2.p_2$. For simplicity, the rule does not allow concurrent reads of variables or heap locations. This could be enabled by fractional permissions⁵).

In addition, rules in Fig. 5, 6, and 7 have to be adjusted: every judgment of the form $\Psi \vdash \{\Gamma.p\} c \{\Gamma'.p'\}$ is changed to $\Gamma_c.I, \Psi \vdash_c \{\Gamma.p\} c \{\Gamma'.p'\}$.

It is worth mentioning that the rules allow the weak heap to be accessed concurrently without synchronization (in contrast to the strong heap). Although it is partly because there are only single-cell references and because access to these cells is atomic in the operational semantics, this demonstrates an important difference between strong and weak heaps in the concurrent setting.

$$\boxed{\Gamma_{\mathbf{G}}, \mathbf{I}, \Psi \vdash_c \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}}$$

$$\frac{\text{dom}(\Gamma) = \text{dom}(\Gamma') \quad \{\}. \text{emp}, \Psi \vdash_c \{(\Gamma_{\mathbf{G}}.I) * (\Gamma.\mathbf{p})\} c \{(\Gamma_{\mathbf{G}}.I) * (\Gamma'.\mathbf{p}')\}}{\Gamma_{\mathbf{G}}.I, \Psi \vdash_c \{\Gamma.\mathbf{p}\} \text{atomic}\{c\} \{\Gamma'.\mathbf{p}'\}} \quad (\text{ATOMIC})$$

$$\frac{\forall i \in [1, 2]. \Gamma_{\mathbf{G}}.I, \Psi \vdash_c \{\Gamma_i.\mathbf{p}_i\} c_i \{\Gamma'_i.\mathbf{p}'_i\}}{\Gamma_{\mathbf{G}}.I, \Psi \vdash_c \{(\Gamma_1.\mathbf{p}_1) * (\Gamma_2.\mathbf{p}_2)\} c_1 \| c_2 \{(\Gamma'_1.\mathbf{p}'_1) * (\Gamma'_2.\mathbf{p}'_2)\}} \quad (\text{PAR})$$

Fig. 11 Checking commands in the presence of concurrency.

An example. This example computes the largest Fibonacci number that is less than 101376. Below is the initialization code for the example.

```
r1 := allocw,int(0); r2 := r1; y := allocs(1); z := allocs(1)
```

The example simulates the situation where a high-level program invokes some low-level program for performing the actual computation. The first operation in the initialization allocates a weak-heap cell for storing the final result; this is to simulate an operation by the high-level program. The low-level program receives the reference $r1$, makes a copy ($r2 := r1$), and then initializes two cells in the strong heap (pointed to by y and z) for storing two adjacent numbers in the Fibonacci sequence. The two cells are initialized to be 1.

After initialization, the low-level code uses two threads for performing the computation: $c_1 \| c_2$. The program c_1 , presented below, contains a loop that retrieves the two numbers stored in the global strong heap, terminates if the second is greater than or equal to 101376, and computes the next two numbers in the Fibonacci sequence otherwise. We define c_2 to be the same as c_1 except that $r1$, t_{d1} , t_{y1} , and t_{z1} are replaced by $r2$, t_{d2} , t_{y2} , and t_{z2} , respectively.

```
td1 := 0;
while (td1 = 0) do {
  atomic {
    ty1 := [y]; tz1 := [z];
    if (tz1 ≥ 101376) {td1 := 1}
    else {[y] := tz1; [z] := ty1 + tz1}
  }
};
[r1] := ty1
```

After initialization, the heap satisfies the specification $(\Gamma_G.I) * (\Gamma_1.\text{emp}) * (\Gamma_2.\text{emp})$, where $\Gamma_G = \{y : \text{int}, z : \text{int}\}$, $I = (y \mapsto -) * (z \mapsto -)$, $\Gamma_1 = \{r1 : \text{wref int}, t_{d1} : \text{int}, t_{y1} : \text{int}, t_{z1} : \text{int}\}$, and $\Gamma_2 = \{r2 : \text{wref int}, t_{d2} : \text{int}, t_{y2} : \text{int}, t_{z2} : \text{int}\}$. It is easy to show that $\Gamma_G.I, \Psi \vdash_c \{\Gamma_1.\text{emp}\} c_1 \{\Gamma_1.\text{emp}\}$ holds for any Ψ , assuming the obvious rules for while and if:

$$\frac{\Gamma_G.I, \Psi \vdash_c \{\Gamma.p \wedge b\} c \{\Gamma.p\}}{\Gamma_G.I, \Psi \vdash_c \{\Gamma.p\} \text{ while } (b) \text{ do } c \{\Gamma.p \wedge \neg b\}} \text{ (WHILE)}$$

$$\frac{\Gamma_G.I, \Psi \vdash_c \{\Gamma.p \wedge b\} c_1 \{\Gamma'.p'\} \quad \Gamma_G.I, \Psi \vdash_c \{\Gamma.p \wedge \neg b\} c_2 \{\Gamma'.p'\}}{\Gamma_G.I, \Psi \vdash_c \{\Gamma.p\} \text{ if } (b) c_1 \text{ else } c_2 \{\Gamma'.p'\}} \text{ (IF)}$$

Similarly, $\Gamma_G.I, \Psi \vdash_c \{\Gamma_2.\text{emp}\} c_2 \{\Gamma_2.\text{emp}\}$ holds. Therefore, by the parallel-composition rule, we have

$$\Gamma_G.I, \Psi \vdash_c \{\Gamma_1.\text{emp} * \Gamma_2.\text{emp}\} c_1 \| c_2 \{\Gamma_1.\text{emp} * \Gamma_2.\text{emp}\}$$

Soundness. We next prove the soundness of the concurrent version of SL^W . We first define the semantics of the judgment for commands.

Definition 4.1 (Partial correctness)

$$\begin{aligned} \Gamma_G.I, \Psi \models_c \{\Gamma.p\} c \{\Gamma'.p'\} &\triangleq \\ &\forall F \geq \llbracket \Psi \rrbracket, s, h. (\models (s, h) \text{ sat } (F, (\Gamma_G.I) * (\Gamma.p))) \Rightarrow \\ &\quad \neg((s, h, c) \longrightarrow^* \text{wrong}) \wedge \\ &\quad (\forall s', h'. (s, h, c) \longrightarrow^* (s', h', \epsilon) \Rightarrow \\ &\quad \exists F' \geq F. \models (s', h') \text{ sat } (F', (\Gamma_G.I) * (\Gamma'.p'))) \end{aligned}$$

Theorem 4.1 (Soundness)

If $\Gamma_G.I, \Psi \vdash_c \{\Gamma.p\} c \{\Gamma'.p'\}$, then $\Gamma_G.I, \Psi \models_c \{\Gamma.p\} c \{\Gamma'.p'\}$.

Proving the soundness theorem directly based on the induction over the derivation of the assumption would not go through for the case of the PAR rule. We need a stronger induction principle given by Definition 4.3. Then the soundness follows from the following Lemma 4.2 and Lemma 4.3.

Definition 4.2

$\Gamma_G.I, F \models_c^0 (s, h, c) : \Gamma.p$ always holds; $\Gamma_G.I, F \models_c^{k+1} (s, h, c) : \Gamma.p$ holds iff

1. $\neg((s, h, c) \longrightarrow^* \text{wrong})$;

2. if $c = \epsilon$, then $\models (s, h) \text{ sat } (\mathbf{F}, (\Gamma_{\mathbf{G}}.\mathbf{I}) * (\Gamma.\mathbf{p}))$.
3. for all s', h' , and c' , if $(s, h, c) \longrightarrow (s', h', c')$, then there exists \mathbf{F}' such that $\mathbf{F}' \geq \mathbf{F}$ and $\forall j \leq k. \Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F}' \models_c^j (s', h', c') : \Gamma.\mathbf{p}$;
4. there exist h_1, h_2, s_1 and s_2 such that $h = h_1 \uplus h_2, s = s_1 \uplus s_2$, and
 - a. $\models (s_1, h_1) \text{ sat } (\mathbf{F}, \Gamma_{\mathbf{G}}.\mathbf{I})$;
 - b. for all h'_1, s'_1 and \mathbf{F}' , if $h'_1 \perp h_2, s'_1 \perp s_2, \mathbf{F}' \geq \mathbf{F}$, and $\models (s'_1, h'_1) \text{ sat } (\mathbf{F}', \Gamma_{\mathbf{G}}.\mathbf{I})$, then, $\forall j \leq k. \Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F}' \models_c^j (s'_1 \uplus s_2, h'_1 \uplus h_2, c) : \Gamma.\mathbf{p}$.

We also define $\Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F} \models_c (s, h, c) : \Gamma.\mathbf{p}$ as $\forall k. \Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F} \models_c^k (s, h, c) : \Gamma.\mathbf{p}$.

Lemma 4.1

If $\Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F} \models_c (s, h, c) : \Gamma.\mathbf{p}$, then (1) $\neg((s, h, c) \longrightarrow^* \text{wrong})$, and,

(2) for all s', h' , if $(s, h, c) \longrightarrow^* (s', h', \epsilon)$, then there exists \mathbf{F}' such that $\mathbf{F}' \geq \mathbf{F}$ and $\models (s', h') \text{ sat } (\mathbf{F}', (\Gamma_{\mathbf{G}}.\mathbf{I}) * (\Gamma.\mathbf{p}))$.

Proof. By induction over the number of execution steps. \square

Definition 4.3

$$\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \models_{\text{sc}} \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\} \triangleq \\ \forall \mathbf{F} \geq \llbracket \Psi \rrbracket, s, h. (\models (s, h) \text{ sat } (\mathbf{F}, (\Gamma_{\mathbf{G}}.\mathbf{I}) * (\Gamma.\mathbf{p}))) \Rightarrow (\Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F} \models_c (s, h, c) : \Gamma'.\mathbf{p}')$$

Lemma 4.2

If $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \models_{\text{sc}} \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$, then $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \models_c \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$.

Proof. Immediate from Lemma 4.1. \square

Lemma 4.3

If $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \vdash_c \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$, then $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \models_{\text{sc}} \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$.

Proof. By induction over the derivation of $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \vdash_c \{\Gamma.\mathbf{p}\} c \{\Gamma'.\mathbf{p}'\}$.

Case 1, The PAR rule. We know $c = c_1 \parallel c_2$ and there exist $\Gamma_1, \Gamma_2, \Gamma'_1, \Gamma'_2, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}'_1$ and \mathbf{p}'_2 such that $\Gamma.\mathbf{p} = \Gamma_1.\mathbf{p}_1 * \Gamma_2.\mathbf{p}_2, \Gamma'.\mathbf{p}' = \Gamma'_1.\mathbf{p}'_1 * \Gamma'_2.\mathbf{p}'_2, \Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \vdash_c \{\Gamma_1.\mathbf{p}_1\} c_1 \{\Gamma'_1.\mathbf{p}'_1\}$, and $\Gamma_{\mathbf{G}}.\mathbf{I}, \Psi \vdash_c \{\Gamma_2.\mathbf{p}_2\} c_2 \{\Gamma'_2.\mathbf{p}'_2\}$.

For all s, h and \mathbf{F} , if $\mathbf{F} \geq \llbracket \Psi \rrbracket$ and $\models (s, h) \text{ sat } (\mathbf{F}, \Gamma_{\mathbf{G}}.\mathbf{I} * \Gamma.\mathbf{p})$, show $\Gamma_{\mathbf{G}}.\mathbf{I}, \mathbf{F} \models_c (s, h, c) : \Gamma'.\mathbf{p}'$. We know there exist s_0, s_1, s_2, h_0, h_1 and h_2 such that $s = s_0 \uplus s_1 \uplus s_2, h = h_0 \uplus h_1 \uplus h_2, \models (s_0, h_0) \text{ sat } (\mathbf{F}, \Gamma_{\mathbf{G}}.\mathbf{I}), \mathbf{F}, s_1, h_1 \models \Gamma_1.\mathbf{p}_1$, and $\mathbf{F}, s_2, h_2 \models \Gamma_2.\mathbf{p}_2$.

By the induction hypothesis we know $\Gamma_G.I, \Psi \models_{sc} \{\Gamma_1.p_1\} c_1 \{\Gamma'_1.p'_1\}$, and $\Gamma_G.I, \Psi \models_{sc} \{\Gamma_2.p_2\} c_2 \{\Gamma'_2.p'_2\}$. Then we know

- (1) $\Gamma_G.I, F \models_c (s_0 \uplus s_1, h_0 \uplus h_1, c_1) : \Gamma'_1.p'_1$; and
- (2) $\Gamma_G.I, F \models_c (s_0 \uplus s_2, h_0 \uplus h_2, c_2) : \Gamma'_2.p'_2$.

Our final goal $\Gamma_G.I, F \models_c (s, h, c) : \Gamma'.p'$ can be derived from Lemma 4.4 below.

Other cases: Proof is similar and is omitted here. \square

Lemma 4.4

For all k , if $s = s_0 \uplus s_1 \uplus s_2$, $h = h_0 \uplus h_1 \uplus h_2$, $\models (s_0, h_0) \text{ sat } (F, \Gamma_G.I)$, $\Gamma_G.I, F \models_c^k (s_0 \uplus s_1, h_0 \uplus h_1, c_1) : \Gamma_1.p_1$, and $\Gamma_G.I, F \models_c^k (s_0 \uplus s_2, h_0 \uplus h_2, c_2) : \Gamma_2.p_2$, then we have $\Gamma_G.I, F \models_c^k (s, h, c_1 \parallel c_2) : (\Gamma_1.p_1) * (\Gamma_2.p_2)$.

Proof. By induction over k . It is trivial when $k = 0$. Suppose $k = i + 1$. We need to prove the four sub-goals in Definition 4.2.

The sub-goals 1, 2 and 4.a are trivial. 4.b can be proved based on the induction hypothesis. Now we prove the sub-goal 3.

Suppose $(s, h, c_1 \parallel c_2) \longrightarrow (s', h', c')$ for some s' , h' and c' . There are three cases.

Case 1, $c_1 = c_2 = c' = \epsilon$.

Then we know $s = s'$, and $h = h'$. By $\Gamma_G.I, F \models_c^{i+1} (s_0 \uplus s_1, h_0 \uplus h_1, c_1) : \Gamma_1.p_1$, we know $\models (s_0 \uplus s_1, h_0 \uplus h_1) \text{ sat } (F, (\Gamma_G.I) * (\Gamma_1.p_1))$. Similarly we can prove $\models (s_0 \uplus s_2, h_0 \uplus h_2) \text{ sat } (F, (\Gamma_G.I) * (\Gamma_2.p_2))$. Since $\models (s_0, h_0) \text{ sat } (F, \Gamma_G.I)$ and I is precise, we have $\models (s, h) \text{ sat } (F, (\Gamma_G.I) * (\Gamma_1.p_1) * (\Gamma_2.p_2))$. By Definition 4.2 it is easy to prove $\forall j \leq i. \Gamma_G.I, F \models_c^j (s', h', c') : (\Gamma_1.p_1) * (\Gamma_2.p_2)$.

Case 2, $c' = c'_1 \parallel c_2$ and $(s, h, c_1) \longrightarrow (s', h', c'_1)$.

By Lemma 3.3, we know there exist s'' , and h'' such that $s' = s'' \uplus s_2$, $h' = h'' \uplus h_2$, and $(s_0 \uplus s_1, h_0 \uplus h_1, c_1) \longrightarrow (s'', h'', c'_1)$. By $\Gamma_G.I, F \models_c^{i+1} (s_0 \uplus s_1, h_0 \uplus h_1, c_1) : \Gamma_1.p_1$ we know there exists $F' \geq F$ such that for all $j \leq i$, $\Gamma_G.I, F' \models_c^j (s'', h'', c'_1) : \Gamma_1.p_1$ (following 3 of Definition 4.2). Then by 4.a of Definition 4.2 we know there exist s'_0, s'_1, h'_0 and h'_1 such that $s'' = s'_0 \uplus s'_1$, $h'' = h'_0 \uplus h'_1$, and $\models (s'_0, h'_0) \text{ sat } (F', \Gamma_G.I)$.

Then, by $\Gamma_G.I, F \models_c^{i+1} (s_0 \uplus s_2, h_0 \uplus h_2, c_2) : \Gamma_2.p_2$ and 4.b of Definition 4.2 we know for all $j \leq i$, $\Gamma_G.I, F' \models_c^j (s'_0 \uplus s_2, h'_0 \uplus h_2, c_2) : \Gamma_2.p_2$.

By the induction hypothesis, we know for all $j \leq i$, $\Gamma_G.I, F' \models_c^j (s', h', c') : (\Gamma_1.p_1) * (\Gamma_2.p_2)$.

Case 3, $c' = c_1 \parallel c'_2$ and $(s, h, c_2) \longrightarrow (s', h', c'_2)$.

The proof is similar to the case above. \square

§5 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 languages¹¹⁾, and the interoperation between two safe languages when they have different systems of computational effects²²⁾. 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^{14, 10)} 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^{9, 25)} for higher-order languages supports reference types but also requires a separate type system (in addition to the Hoare assertions); Reus *et al*¹⁶⁾ 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.

Pottier showed how to use an anti-frame rule to encode weak references in separation logic¹⁵⁾. His encoding relies heavily on the information-hiding aspect. A client of a weak reference is supplied with a pair of a getter and a setter function and the very existence of a reference cell is hidden by the anti-frame rule. Most type systems (and also this article), in contrast, use a global heap type to ensure the soundness of weak references. The benefit is that a client of a weak heap can get direct references to objects in the weak heap. This style matches the setting of foreign function interfaces where references to weak-heap objects are passed to foreign code. The plus side of Pottier's encoding is that a single logic suffices while we mix separation logic with a type system.

The soundness of SL^W is justified by defining a semantic model, notably for types. Ahmed¹⁾ 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 requires a stratification of types and the use of dependent types, which our model avoids.

§6 Discussion and future work

This work aims toward a framework for reasoning about language inter-operation, 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 system²¹⁾. 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.

§7 Conclusion

In his survey paper of separation logic¹⁷⁾, Reynolds asked “*whether the dividing line between types and assertions can be erased*”. This article 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 Hongseok Yang for pointing out the connection between our work and Pottier’s work. We thank anonymous referees for suggestions and comments on an earlier version of this article. 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 *Proc. 34th ACM Symp. on Principles of Prog. Lang.*, 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) R. Bornat, C. Calcagno, P. O’Hearn, and M. Parkinson. Permission accounting in separation logic. In *Proc. 32nd ACM Symp. on Principles of Prog. Lang.*, pages 259–270, 2005.
- 6) M. Furr and J. S. Foster. Checking type safety of foreign function calls. *ACM Trans. Program. Lang. Syst.*, 30(4):1–63, 2008.
- 7) R. Harper. A simplified account of polymorphic references. *Information Processing Letters*, 57(1):15–16, 1996.
- 8) C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):578–580, October 1969.
- 9) K. Honda, N. Yoshida, and M. Berger. An observationally complete program logic for imperative higher-order frame rules. In *Proc. 20th IEEE Symposium on Logic in Computer Science*, pages 270–279, June 2005.
- 10) N. Krishnaswami, L. Birkedal, J. Aldrich, and J. Reynolds. Idealized ML and its separation logic. Unpublished manuscript, July 2007.
- 11) 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.
- 12) 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.
- 13) P. W. O’Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *Proc. 31th ACM Symp. on Principles of Prog. Lang.*, pages 268–280, Venice, Italy, Jan. 2004.

- 14) M. Parkinson. *Local reasoning for Java*. PhD thesis, University of Cambridge Computer Laboratory, Oxford, Nov. 2005. Tech Report UCAM-CL-TR-654.
- 15) F. Pottier. Hiding local state in direct style: a higher-order anti-frame rule. In *Proc. 23rd IEEE Symposium on Logic in Computer Science*, pages 331–340, June 2008.
- 16) B. Reus and J. Schwinghammer. Separation logic for higher-order store. In *20th International Workshop on Computer Science Logic (CSL)*, pages 575–590, 2006.
- 17) J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th IEEE Symposium on Logic in Computer Science*, pages 55–74, July 2002.
- 18) 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.
- 19) G. Tan, Z. Shao, X. Feng, and H. Cai. Weak updates and separation logic. In *Proceedings of the 7th Asian Symposium on Programming Languages and Systems (APLAS '09)*, pages 178–193, 2009.
- 20) M. Tofte. Type inference for polymorphic references. *Inf. and Comp.*, 89(1):1–34, 1990.
- 21) M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- 22) V. Trifonov and Z. Shao. Safe and principled language interoperation. In *8th European Symposium on Programming (ESOP)*, pages 128–146, 1999.
- 23) V. Vafeiadis and M. J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, pages 256–271, 2007.
- 24) A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.
- 25) 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.