

A Kripke Logical Relation Between ML and Assembly

Chung-Kil Hur* Derek Dreyer

Max Planck Institute for Software Systems (MPI-SWS)

{gil,dreyer}@mpi-sws.org

Abstract

There has recently been great progress in proving the correctness of compilers for increasingly realistic languages with increasingly realistic runtime systems. Most work on this problem has focused on proving the correctness of a particular compiler, leaving open the question of how to verify the correctness of assembly code that is hand-optimized or linked together from the output of multiple compilers. This has led Benton and other researchers to propose more abstract, compositional notions of when a low-level program correctly realizes a high-level one. However, the state of the art in so-called “compositional compiler correctness” has only considered relatively simple high-level and low-level languages.

In this paper, we propose a novel, extensional, compiler-independent notion of equivalence between high-level programs in an expressive, impure ML-like λ -calculus and low-level programs in an (only slightly) idealized assembly language. We define this equivalence by means of a biorthogonal, step-indexed, Kripke logical relation, which enables us to reason quite flexibly about assembly code that uses local state in a different manner than the high-level code it implements (*e.g.*, self-modifying code). In contrast to prior work, we factor our relation in a symmetric, language-generic fashion, which helps to simplify and clarify the formal presentation, and we also show how to account for the presence of a garbage collector. Our approach relies on recent developments in Kripke logical relations for ML-like languages, in particular the idea of possible worlds as state transition systems.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.3 [Programming Languages]: Language Constructs and Features; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms Languages, Theory, Verification

Keywords Step-indexed Kripke logical relations, biorthogonality, compositional compiler correctness, garbage collection, self-modifying code

*This work was undertaken while the first author was at PPS, Université Paris Diderot, supported by Digiteo/Ile-de-France project COLLODI (2009-28HD) and Engineering Research Center of Excellence Program of Korea Ministry of Education, Science and Technology (MEST) / National Research Foundation of Korea (NRF) Grant R11-2008-007-01002-0.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL'11, January 26–28, 2011, Austin, Texas, USA.
Copyright © 2011 ACM 978-1-4503-0490-0/11/01...\$10.00

1. Introduction

While compiler verification is an age-old problem, there has been remarkable progress in the last several years in proving the correctness of compilers for increasingly realistic languages with increasingly realistic runtime systems. Of particular note is Leroy’s Compcert project [18], in which he used the Coq proof assistant to both program and verify a multi-pass optimizing compiler from Cminor (a C-like intermediate language) to PowerPC assembly. Dargaye [13] has adapted the Compcert framework to a compiler for a pure mini-ML language, and McCreight *et al.* [19] have extended it to support interfacing with a garbage collector. Independently, Chlipala [10, 12] has developed verified compilers for both pure and impure functional core languages, the former garbage-collected, with a focus on using custom Coq tactics to provide significant automation of verification.

That said, all of the aforementioned work has focused on proving the correctness of a *particular* compiler, leaving open the question of how to verify the correctness of assembly code that is hand-optimized or linked together from the output of multiple compilers. The issue is that compiler correctness results are typically established by exhibiting a fairly close simulation relation between source and target code, but code produced by another compiler may obey an entirely different simulation relation with the source program, and hand-optimized code might not closely simulate the source program at all. Thus, existing correctness proofs depend fundamentally on the “closed-world” assumption that one has control over how the whole source program is compiled.

In order to lift the closed-world assumption, Benton and Hur [5] suggest that what is needed is a more abstract, extensional notion of what it means for a low-level program to correctly implement a high-level one—a notion that is not tied to a particular compiler and that, moreover, offers as much flexibility in the low-level representation of high-level features as possible. When reasoning strictly about high-level programs, the canonical extensional notion of when one program implements the same functionality as another is *observational* (or *contextual*) equivalence, which says that the two programs exhibit the same (termination) behavior when placed into the context of an arbitrary enclosing well-typed high-level program. However, it is not clear how to define such a contextual notion of equivalence between high- and low-level programs, because there is no way to run both programs under the *same* context—one would need to quantify over *equivalent* high- and low-level program contexts, but when are two contexts equivalent? We are back to the original question.

Benton and Hur’s solution is to define a *logical relation* between the high- and low-level languages (actually two relations, one for each direction of semantic approximation, employing a denotational semantics to represent the high-level side). Logical relations are inherently extensional—*e.g.*, two functions are logically related iff they map related arguments to related results, regardless of their private implementation details—and guarantee equivalent

termination behavior under arbitrary contexts that are themselves logically related. While not as canonical as contextual equivalence, logical equivalence is nevertheless useful as long as one can establish that the logical relations are sufficiently populated—*i.e.*, that they relate enough programs/contexts of interest.

In the traditional setting where one is defining equivalence of programs in the same language, this “sufficient population” property is ensured by the *fundamental theorem of logical relations*, which states that all well-typed programs (and thus all well-typed contexts) are logically self-related. For mixed high-low relations, Benton and Hur demonstrate sufficient population by providing a simple, one-pass compilation translation from their high- to their low-level language and proving that all well-typed high-level programs are logically equivalent to their compilations. They also use the logical relations to show the relatedness of some simple hand-optimized low-level code with corresponding high-level programs.

Benton and Hur present their work as the first step towards “compositional compiler correctness”. However, the source language they consider—the simply-typed λ -calculus with recursion—is purely functional, and the target language they consider—an SECD machine—is relatively high-level. Chlipala [11] has subsequently proposed a more syntactic approach to proving compositional compiler correctness, applicable to a richer, impure (albeit untyped) source language, but the target language he considers is also pretty high-level, namely a CPS variant of the source.

1.1 Contributions

In this paper, we study compositional equivalence of high- and low-level programs in a more realistic setting. Our high-level language is an expressive ML-like CBV λ -calculus, supporting abstract types, general recursive types, and general mutable references. Our low-level language is an (only slightly) idealized assembly language. Furthermore, our logical relation is designed to be sound in the presence of garbage collection, under some fairly abstract assumptions about the behavior of the garbage collector that are satisfied by both mark-and-sweep and copying collectors.

Following Benton and Hur, we define our equivalence using a *biorthogonal, step-indexed* logical relation. Biorthogonality is useful when reasoning about programs (such as low-level ones) whose behavior is *context-sensitive*, and step-indexing is useful in reasoning about semantically “cyclic” features like recursive types and higher-order state.

We depart from prior work, though, in that our relation is also a *Kripke* logical relation—*i.e.*, it is indexed by *possible worlds* that specify assumptions about the machine state. Possible worlds are useful in enforcing invariants about low-level data structures (*e.g.*, that a heap-allocated representation of a closure is immutable). They are also helpful in encoding a variety of runtime system invariants, such as the convention concerning callee-save registers and the notion of data liveness. Last but not least, possible worlds enable us to reason quite flexibly about assembly code that uses *local* state in a different manner than the high-level code it implements. An interesting example of this is *self-modifying* assembly code, whose correctness proof involves reasoning about low-level internal state changes—specifically, changes to the code itself—that clearly have no high-level counterpart. This is the essence of what we mean when we say that our relation is *extensional*.

Technically, our approach relies closely on recent developments in Kripke logical relations for ML-like languages, in particular the idea of possible worlds as *state transition systems*. This idea, which we review in Section 3.3, was proposed originally by Ahmed *et al.* [2] (in a somewhat different form) as a way to reason about representation independence for so-called “generative” abstract data types, whose private state undergoes a controlled series of state transitions during the execution of the program.

Dreyer *et al.* [14] have subsequently generalized the idea in order to reason about “well-bracketed” state changes. By basing our high-low logical relations on this most recent work, we are able to cleanly model a variety of state transition systems that arise naturally in low-level code (*e.g.*, in self-modifying code).

Lastly, a novel feature of our logical relation is that, while it is defined by induction on an ML-like type structure, it is also defined in a *language-generic* fashion. That is, it may be instantiated to form an equivalence relation between any two languages (high- or low-level) that are capable of implementing various relevant linguistic constructs—*e.g.*, function application, plugging a continuation with a value, etc. Factoring the relation in this way helps to simplify and clarify the formal presentation. Moreover, it has the advantage that the relation becomes inherently symmetric and thus easier to use in proving high-low equivalences than Benton and Hur’s asymmetric approximation relations.

2. HIGH and LOW

The high-level language **HIGH** is a System F-like polymorphic λ -calculus extended with existential, product and iso-recursive types, as well as ML-style general references (higher-order state). Figure 1 shows the syntax and the typing and evaluation judgments. The inference rules for typing and evaluation are standard, so we omit them (see the companion technical appendix [15] for details).

The low-level language **LOW** is an assembly language idealized in two ways: its word and memory sizes are infinite and its instructions are represented by abstract objects rather than by physical words. Its memory consists of four entities: code memory, register file, stack and heap. A code memory is a map from physical addresses (represented by natural numbers) to instructions. A register file is a map from registers to words. Words in turn are represented by natural numbers with an extra bit indicating whether the word is a pointer to a heap cell or not (useful for garbage collection purposes). There are 12 registers, half of which (sp, sv_0, \dots, sv_4) are specified as callee-save registers by our calling convention. Both a stack and a heap are random-access memories (*i.e.*, maps from addresses to words).

The instruction set includes standard instructions (`jmp`, `jnz`, `jneq`, `move`, `plus`, `minus`) supporting different addressing modes via `lvalues` and `rvalues`. The non-standard operations include `halt` for normal termination, `fail` for raising a runtime error, `jptr` for testing whether a value is a pointer or not, `setptr` for marking a pointer bit, and `isr`, `isw` for inspecting and updating instructions in code memory. Note that as instructions are not represented by words, we employ a bijection \mathbb{E} and its inverse \mathbb{D} (for \mathbb{E} encode and \mathbb{D} decode) to convert back and forth between instructions and words.

The dynamic semantics of **LOW** is standard and is given in Figure 1. A machine configuration (Φ, pc) is a pair consisting of a memory and a program counter; it halts, fails or evolves to another configuration (Φ', pc') by executing the instruction stored at `pc` in the code memory $\Phi.code$.

3. The Key Ideas

In this section, we present the key ideas behind our work through the lens of an illustrative, challenging example. We will walk through the code of this example, suggest intuitively how to reason about it, and then explain how our Kripke logical relation formalizes this intuition. While we will initially ignore the question of how garbage collection affects matters, we will return to this issue at the end of the section and discuss how our logical relation enables reasoning in the presence of a garbage collector.

This section is intended to be accessible to a broad programming-languages audience, and hopefully to serve as a useful guide to the

HIGH – Syntax & Semantics	LOW – Semantics
$\tau ::= \alpha \mid b \mid \tau_1 \times \tau_2 \mid \tau_1 \rightarrow \tau_2 \mid \forall \alpha. \tau \mid \exists \alpha. \tau \mid \mu \alpha. \tau \mid \text{ref } \tau$ $e ::= x \mid \ell \mid \langle e_1, e_2 \rangle \mid e.1 \mid e.2 \mid \lambda x:\tau. e \mid e_1 e_2 \mid \Lambda \alpha. e \mid e \tau \mid$ $\text{pack } \langle \tau_1, e \rangle \text{ as } \tau_2 \mid \text{unpack } e_1 \text{ as } \langle \alpha, x \rangle \text{ in } e_2 \mid$ $\text{roll}_\tau e \mid \text{unroll } e \mid \text{ref } e \mid e_1 := e_2 \mid !e \mid e_1 == e_2 \mid \dots$ $v ::= x \mid \ell \mid \langle v_1, v_2 \rangle \mid \lambda x:\tau. e \mid \Lambda \alpha. e \mid \text{pack } \langle \tau_1, v \rangle \text{ as } \tau_2 \mid \text{roll}_\tau v \mid \dots$ $K ::= \bullet \mid \langle K, e_2 \rangle \mid \langle v_1, K \rangle \mid K.1 \mid K.2 \mid K e_2 \mid v_1 K \mid K \tau \mid \text{roll}_\tau K \mid$ $\text{unroll } K \mid \text{pack } \langle \tau_1, K \rangle \text{ as } \tau_2 \mid \text{unpack } K \text{ as } \langle \alpha, x \rangle \text{ in } e_2 \mid$ $\text{ref } K \mid K := e_2 \mid v_1 := K \mid !K \mid K == e_2 \mid v_1 == K \mid \dots$ $\Sigma ::= \cdot \mid \Sigma, \ell:\tau \text{ with } \text{ftv}(\tau) = \emptyset \quad \Delta ::= \cdot \mid \Delta, \alpha \quad \Gamma ::= \cdot \mid \Gamma, x:\tau$ <p>Static semantics : $\Sigma; \Delta; \Gamma \vdash e : \tau$</p> <p>HCVal $\stackrel{\text{def}}{=} \{ v \mid \text{ftv}(v) = \emptyset \wedge \text{fv}(v) = \emptyset \}$</p> <p>HHeap $\stackrel{\text{def}}{=} \{ h \in \text{HLoc} \rightarrow_{\text{fin}} \text{HCVal} \}$</p> <p>Dynamic semantics : $(h, e) \hookrightarrow (h', e')$</p>	$\llbracket w \rrbracket \stackrel{\text{def}}{=} \pi_2(w) \quad \text{isptr}(w) \stackrel{\text{def}}{=} (\pi_1(w) = 1) \quad \underline{n} \stackrel{\text{def}}{=} (0, n) \quad \hat{a} \stackrel{\text{def}}{=} (1, a)$ $\Phi(w) \stackrel{\text{def}}{=} w \quad \Phi(\llbracket r \rrbracket) \stackrel{\text{def}}{=} \Phi.\text{reg}(r)$ $\Phi(\langle a \rangle_s) \stackrel{\text{def}}{=} \Phi.\text{stk}(a) \quad \Phi(\langle r - o \rangle_s) \stackrel{\text{def}}{=} \Phi.\text{stk}(\Phi(\llbracket r \rrbracket) - o)$ $\Phi(\langle a \rangle_h) \stackrel{\text{def}}{=} \Phi.\text{hp}(a) \quad \Phi(\langle r + o \rangle_h) \stackrel{\text{def}}{=} \Phi.\text{hp}(\Phi(\llbracket r \rrbracket) + o)$ $\Phi[\llbracket r \rrbracket \mapsto w] \stackrel{\text{def}}{=} (\Phi.\text{code}, \Phi.\text{reg}[\llbracket r \rrbracket \mapsto w], \Phi.\text{stk}, \Phi.\text{hp})$ $\Phi[\langle a \rangle_s \mapsto w] \stackrel{\text{def}}{=} (\Phi.\text{code}, \Phi.\text{reg}, \Phi.\text{stk}[\langle a \rangle_s \mapsto w], \Phi.\text{hp})$ $\Phi[\langle r - o \rangle_s \mapsto w] \stackrel{\text{def}}{=} (\Phi.\text{code}, \Phi.\text{reg}, \Phi.\text{stk}[\Phi(\llbracket r \rrbracket) - o \mapsto w], \Phi.\text{hp})$ $\Phi[\langle a \rangle_h \mapsto w] \stackrel{\text{def}}{=} (\Phi.\text{code}, \Phi.\text{reg}, \Phi.\text{stk}, \Phi.\text{hp}[\langle a \rangle_h \mapsto w])$ $\Phi[\langle r + o \rangle_h \mapsto w] \stackrel{\text{def}}{=} (\Phi.\text{code}, \Phi.\text{reg}, \Phi.\text{stk}, \Phi.\text{hp}[\Phi(\llbracket r \rrbracket) + o \mapsto w])$ $\llbracket \text{fail} \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} \text{fail}$ $\llbracket \text{halt} \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} \text{halt}$ $\llbracket \text{jmp } rv \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi, \Phi(rv))$ $\llbracket \text{jnz } rv_1 \text{ } rv_2 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} \text{if } \text{isptr}(\Phi(rv_2)) \neq 0$ $\llbracket \text{jneq } rv_1 \text{ } rv_2 \text{ } rv_3 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} \text{if } \Phi(rv_2) \neq \Phi(rv_3)$ $\llbracket \text{jptr } rv_1 \text{ } rv_2 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} \text{if } \text{isptr}(\Phi(rv_2))$ $\llbracket \text{move } lv \text{ } rv \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi, \Phi(rv) , \text{pc} + 1)$ $\llbracket \text{setptr } lv \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi[\llbracket lv \rrbracket \mapsto \Phi(lv)], \text{pc} + 1)$ $\llbracket \text{plus } lv \text{ } rv_1 \text{ } rv_2 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi[\llbracket lv \rrbracket \mapsto \Phi(rv_1) + \Phi(rv_2)], \text{pc} + 1)$ $\llbracket \text{minus } lv \text{ } rv_1 \text{ } rv_2 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi[\llbracket lv \rrbracket \mapsto \Phi(rv_1) - \Phi(rv_2)], \text{pc} + 1)$ $\llbracket \text{isr } lv \text{ } rv \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} (\Phi[\llbracket lv \rrbracket \mapsto \mathbb{E}(\Phi.\text{code}(\Phi(rv)))], \text{pc} + 1)$ $\llbracket \text{isw } rv_1 \text{ } rv_2 \rrbracket (\Phi, \text{pc}) \stackrel{\text{def}}{=} ((\Phi.\text{code}[\Phi(rv_1) \mapsto \mathbb{D}(\Phi(rv_2))], \Phi.\text{reg}, \Phi.\text{stk}, \Phi.\text{hp}, \text{pc} + 1)$ <p>where $\mathbb{E} : \text{Instruction} \rightarrow \mathbb{N}$ is a bijection, and $\mathbb{D} = \mathbb{E}^{-1}$.</p> <p>Dynamic semantics : $(\Phi, \text{pc}) \hookrightarrow \llbracket \Phi.\text{code}(\text{pc}) \rrbracket (\Phi, \text{pc})$</p>
<p style="text-align: center;">LOW – Syntax</p> <p>PConf $\stackrel{\text{def}}{=} \{ (\Phi, \text{pc}) \in \text{PMem} \times \text{PAddr} \}$</p> <p>PMem $\stackrel{\text{def}}{=} \{ \Phi = (\text{code}, \text{reg}, \text{stk}, \text{hp})$ $\in \text{PCode} \times \text{PRegFile} \times \text{PStack} \times \text{PHeap} \}$</p> <p>PCode $\stackrel{\text{def}}{=} \text{PAddr} \rightarrow \text{Instruction}$ PRegFile $\stackrel{\text{def}}{=} \text{Register} \rightarrow \text{PWord}$</p> <p>PStack $\stackrel{\text{def}}{=} \text{PAddr} \rightarrow \text{PWord}$ PHeap $\stackrel{\text{def}}{=} \text{PAddr} \rightarrow \text{PWord}$</p> <p>PAddr $\stackrel{\text{def}}{=} \{ a \in \mathbb{N} \}$ PWord $\stackrel{\text{def}}{=} \{ w \in \{0, 1\} \times \mathbb{N} \}$</p> <p>$r \in \text{Register} ::= \text{sp} \mid \text{sv}_0 \mid \dots \mid \text{sv}_4 \mid \text{wk}_0 \mid \dots \mid \text{wk}_5$</p> <p>$lv \in \text{PLvalue} ::= \llbracket r \rrbracket \mid \langle a \rangle_s \mid \langle r - o \rangle_s \mid \langle a \rangle_h \mid \langle r + o \rangle_h$</p> <p>$rv \in \text{PRvalue} ::= lv \mid w$</p> <p>$\iota \in \text{Instruction} ::= \text{fail} \mid \text{halt} \mid \text{jmp } rv \mid \text{jnz } rv_1 \text{ } rv_2 \mid \text{jneq } rv_1 \text{ } rv_2 \text{ } rv_3 \mid$ $\text{jptr } rv_1 \text{ } rv_2 \mid \text{setptr } lv \mid \text{move } lv \text{ } rv \mid \text{plus } lv \text{ } rv_1 \text{ } rv_2 \mid$ $\text{minus } lv \text{ } rv_1 \text{ } rv_2 \mid \text{isr } lv \text{ } rv \mid \text{isw } rv_1 \text{ } rv_2$</p>	

Figure 1. Syntax and Semantics for the **HIGH** and **LOW** Languages (excerpt)

vast majority of readers who are not intimately familiar with recent developments in Kripke logical relations.

3.1 A Motivating Example

Our motivating example is based on Pitts and Stark’s “awkward” example [22]. Their original example is almost blindingly simple—prove that the following **HIGH** terms are contextually equivalent:

$$e \stackrel{\text{def}}{=} \text{let } x = \text{ref } 0 \text{ in } \lambda f:\text{unit} \rightarrow \text{unit}. (x := 1; f \langle \rangle; !x)$$

$$e' \stackrel{\text{def}}{=} \lambda f:\text{unit} \rightarrow \text{unit}. (f \langle \rangle; 1)$$

The first term, e , allocates a fresh memory location x , initially set to 0, and then returns a higher-order function. When called, the latter will set x to 1, invoke its callback argument f , and then return the contents of x . The second term, e' , is similar, except that it does not bother allocating or updating any memory, and the function it defines always returns 1. Proving that e and e' are contextually equivalent is tantamount to showing that, in the former, whenever the callback invocation $f \langle \rangle$ returns, x points to 1.

After a moment’s thought, it should hopefully be intuitively clear why the equivalence holds. The pointer x in e is initially set to 0, but once the function that e returns is applied for the first time, x will be set to 1 and will never be set back to 0. This is because e “owns” x as a piece of *local state* and the only thing e ever does to x after first allocating it is to set it to 1. Thus, the “awkward” example serves as an elegant distillation of (1) the ability of code to control some local state and impose arbitrary constraints on it, and (2) the ability of that local state to evolve over time in an *irreversible* (or *monotone*) way. Such irreversible changes to local state arise in a variety of real-world situations—*e.g.*, in “generative” ADTs (whose sets of inhabitants grow over time), and in data structure initialization [2]. It is thus rather remarkable that only recently have

methods been developed for proving an equivalence as simple as the “awkward” example [8, 2, 23]. (More on that in Section 3.3.)

We are now ready to present our motivating example. We want to prove a variant of the awkward example, namely that the **HIGH** term e is implemented correctly by a **LOW** program p , where p ’s implementation follows the second **HIGH** term e' fairly closely. By itself, that would already be an interesting result—it would demonstrate the extensional equivalence of a high-level program with an “optimized” low-level program, where the optimization is based on the inability of high-level program contexts to observe e ’s manipulation of local state. But we will make it more interesting still with an added twist: the code of the function that p evaluates to will be *obfuscated* using a primitive form of encryption, and when first applied, the function will first decrypt and overwrite itself via *self-modifying code*.

The **LOW** program p is shown in Figure 2. Before we walk through the code, let us first note that p is parameterized by alloc , a code pointer to the memory allocation routine, and bg , the location in the code segment where p ’s code will be loaded and where its execution will begin—these parameters will be instantiated as part of linking and loading. (We will define the formal semantics of linking and loading for **LOW** programs in Section 7.)

bg: Create and return a closure. The evaluation of p is very simple: it does no interesting computation except to immediately create and return a closure value, just as the **HIGH** term e' does. The first 3 instructions starting at bg allocate a fresh one-word closure on the heap by invoking the alloc routine (passing it the size parameter 1 in register wk_5 and the return address $\text{bg} + 3$ in register wk_4). We only need one word for the closure because the function we’re implementing (e') is closed, so all we need to store in the closure is the naked code pointer. The alloc routine is assumed to return the pointer to a fresh one-word cell in wk_5 , without modifying the contents of any registers except wk_4 and

$e \stackrel{\text{def}}{=} \text{let } x = \text{ref } 0 \text{ in } \lambda f:\text{unit} \rightarrow \text{unit}. x := 1; f \langle \rangle; !x$			
$p \stackrel{\text{def}}{=} \lambda \text{alloc}, \text{bg}. [$			
bg	move	[wk ₄]	bg + 3
	move	[wk ₅]	<u>1</u>
	jmp	alloc	
bg + 3	move	$\langle \text{wk}_5 + 0 \rangle_h$	bg + 5
	jmp	[wk ₀]	
bg + 5	move	[wk ₃]	bg + 10
bg + 6	isr	[wk ₄]	[wk ₃]
	minus	[wk ₄]	[wk ₄]
	isw	[wk ₃]	[wk ₄]
	plus	[wk ₃]	[wk ₃]
bg + 10	$\mathbb{D}(\mathbb{E}(\text{jneq}$	bg + 6	[wk ₃]
bg + 11	$\mathbb{D}(\mathbb{E}(\text{isw}$	bg + 5	$\mathbb{E}(\text{jmp } \text{bg} + 12)$
bg + 12	$\mathbb{D}(\mathbb{E}(\text{move}$	$\langle \text{wk}_1 + 0 \rangle_h$	bg + 13
bg + 13	$\mathbb{D}(\mathbb{E}(\text{plus}$	[sp]	[sp]
	$\mathbb{D}(\mathbb{E}(\text{move}$	$\langle \text{sp} - 1 \rangle_s$	[wk ₀]
	$\mathbb{D}(\mathbb{E}(\text{move}$	[wk ₁]	[wk ₂]
	$\mathbb{D}(\mathbb{E}(\text{move}$	[wk ₀]	bg + 18
	$\mathbb{D}(\mathbb{E}(\text{jmp}$	$\langle \text{wk}_1 + 0 \rangle_h$	
bg + 18	$\mathbb{D}(\mathbb{E}(\text{move}$	[wk ₅]	<u>1</u>
	$\mathbb{D}(\mathbb{E}(\text{minus}$	[sp]	[sp]
bg + 20	$\mathbb{D}(\mathbb{E}(\text{jmp}$	$\langle \text{sp} - 0 \rangle_s$	
]			

Figure 2. Motivating Example

wk₅. We then store in that cell the code pointer bg + 5, before jumping to the return address, which we assume the linker/loader had passed to p originally in register wk₀. (The linker/loader has essentially the same calling convention as for ordinary functions, which is different from the one for alloc and is described below—see bg + 12.) Although our calling convention is that return values are passed back to the caller in wk₅, our return value was already in wk₅ after the call to alloc, so we need not explicitly move anything into wk₅ before returning.

We now describe the implementation of the closure returned (in wk₅) by the evaluation of p . Initially, this closure contains just a code pointer to bg + 5, but eventually that code pointer will be updated (see below).

bg + 5: Decrypt and overwrite the code. The code from bg + 10 on is obfuscated by the addition of 666 to the machine representation of each instruction. Eventually, once the code is decrypted, the function will be executable starting at the address bg + 13. Before that time, however, the function must begin executing at bg + 5, because the first step will be to decrypt and overwrite the obfuscated instructions. The reader can easily verify manually that the code starting at bg + 5 will use register wk₃ to loop through the instructions bg + 10 through bg + 20. For each, it will use isr to read the instruction stored at wk₃ into wk₄, subtract 666 from it, and then write the decrypted instruction back to the code segment at address wk₃. When this loop is finished, the program counter will be at bg + 11.

bg + 11: Redirect the first instruction of the function. Having decrypted the code, we do not want future calls to this function to perform the decryption again. We therefore overwrite the first instruction (at bg + 5) with a jump to bg + 12.

bg + 12: Update the code pointer. The LOW calling convention is that a function is passed its return address in wk₀, its argument in wk₂, and its own closure in wk₁. The decryption code starting at bg + 5 did not touch any of these registers, so at bg + 12 we know that wk₁ still stores a closure with a code pointer to bg + 5. Having decrypted the code, we can now safely update this code pointer— $\langle \text{wk}_1 + 0 \rangle_h$ —to point to bg + 13, the address where the function begins its computation in earnest.

An important point: Why did we bother *both* redirecting the bg + 5 instruction (to jump to bg + 12) *and* updating the code pointer in the function closure? Would the latter alone not have been sufficient? The answer is that it depends. If p were only evaluated once, in which case only one closure for this function were ever generated, then yes, just updating the code pointer would be sufficient because bg + 5 would become effectively dead code. But we would like our notion of high-low program equivalence to be preserved under a rich set of program contexts including those that evaluate the programs—here, p and e —more than once. Repeated evaluation of p will result in the creation of multiple closures for the function p defines, and merely updating the code pointer for one closure will not change the fact that other closures may still point to bg + 5, so it is necessary to redirect the bg + 5 instruction as well.

bg + 13: Implement $\lambda f. (f \langle \rangle; 1)$. This is the implementation of the function proper. We first push our return address wk₀ onto the stack. We then invoke the callback argument (f in the **HIGH** code) by moving a pointer to f 's closure into wk₁, moving the return address bg + 18 into wk₀, and jumping to f 's code pointer $\langle \text{wk}_1 + 0 \rangle_h$. (Note: f 's argument type is unit, so there is no need to pass anything in the argument register wk₂.) When control is returned to bg + 18, we store the result 1 in the result register wk₅, pop the return address off the stack, and jump to it.

3.2 Discussion of the Motivating Example

Why does p implement e ? Intuitively, the reason is that the self-modifying aspects of p are not visible to p 's clients because they do not affect its extensional behavior; and ultimately, once p has decrypted itself, it behaves essentially the same as the **HIGH** term e' , which we have already argued is equivalent to e . Of course, this begs the question: how exactly do we know that p 's clients cannot observe its self-modifications?

Interestingly, the answer is remarkably similar to the argument for why e and e' are equivalent, namely that what p does to its own code takes the form of irreversible changes to local state. Specifically, we take it as a given that p “owns” its own code, and since the evaluation of p will allocate a fresh memory cell for the closure it returns, p owns that closure as well. Thus, in reasoning about p , we can place restrictions on how its code and closure may evolve over time. Much as the local variable x in e starts out pointing to 0 and eventually points to 1, the code of p starts out in encrypted form, and if/when any closure it returns is first applied, it changes to decrypted form. In both cases, it is critical that we never revert to the earlier state. Similarly, the closure returned by the evaluation of p starts out with its code pointer set to bg + 5, but if/when the closure is ever applied, it will be set to bg + 13. In this case, it is not so essential for correctness that the code pointer never revert back to bg + 5, but it *is* essential that the closure's code pointer only be set to bg + 13 when the code is in the decrypted state.

Given these restrictions on how the local state of p and e may evolve, it is but a short distance to a *bona fide* proof. Before sketching that proof, let us first review the recent work on Kripke logical relations that will put our reasoning on a solid footing.

3.3 Kripke Logical Relations and State Transition Systems

Logical relations are a well-established technique for reasoning about equivalence of higher-order programs. A logical relation is defined inductively on the type structure of the language: at base type the logical relation coincides with observable equality—*e.g.*, two programs of type int are logically related if they produce the same integer—and at higher type the relation is defined by interpreting each type operator by the appropriate logical connective—*e.g.*, two functions are related at type $\tau_1 \rightarrow \tau_2$ if relatedness of their

arguments at type τ_1 *implies* relatedness of their results at type τ_2 . The important feature about logical relations for our purposes is that they give considerable leeway to how related functions are implemented, so long as they produce related results. (As Benton and Hur [5] channeling Machiavelli put it, “the ends justify the means.”)

In the presence of state, we cannot talk about the relatedness of two programs without making some assumptions and imposing some restrictions on how they manipulate state. This is where *Kripke logical relations* come in. Kripke logical relations are indexed by *possible worlds*, which represent a set of restrictions on the memories of the two programs under which the programs are guaranteed to behave equivalently. When we want to prove the relatedness of two programs under a world W , we suppose we are given arbitrary initial memories that are related by (*i.e.*, satisfy the restrictions of) W , and we proceed typically by showing that when evaluated under those memories the programs either (1) both diverge (don’t terminate), or else (2) produce values and final memories that are related under some “future” world W' of W .

What does it mean for W' to be a future world of W ? If in the course of evaluation the programs allocate fresh pieces of memory, W' may extend the initial world W with new restrictions governing the use of the freshly allocated memory. This approach allows us to establish whatever constraints we want on freshly allocated state that is kept *local*. (If the state is made globally accessible—*e.g.*, by being passed to the context at a ref type—then the state will have to obey the usual invariants dictated by the ref type.)

In traditional Kripke logical relations, such as those of Pitts and Stark [22], possible worlds essentially take the form of simple memory relations, *i.e.*, memory *invariants*. As we have seen in the “awkward” example, however, memory invariants are not necessarily enough; we need additionally the ability to describe assumptions about state that may change in a controlled and monotone way. (It is thus not a surprise that Pitts and Stark put forth the “awkward” example as an example for which their method was inadequate.)

To address this limitation, Ahmed *et al.* [2] proposed generalizing possible worlds to include the ability for a memory relation to *evolve*. Dreyer *et al.* [14] later streamlined and extended Ahmed *et al.*’s approach in various ways, and cast Ahmed *et al.*’s possible worlds as collections of *state transition systems (STS’s)*. In the case of the “awkward” example, one can understand the restrictions placed on e ’s local variable x according to the following STS:



When this STS is first added to the initial world W , it starts out in the $x \mapsto 0$ state, because after x is first allocated, it points to 0. However, under Dreyer *et al.*’s model, future worlds of W may not only place additional restrictions on fresh pieces of memory but also update the state of existing STS’s in W . Thus, in some future world W' , the above STS may be switched to the $x \mapsto 1$ state, and memories satisfying W' would have to map x to 1. Furthermore, any future world of W' would have to remain in the $x \mapsto 1$ state as there is no transition out of it. This corresponds to the intuitive reasoning about the example that we described in Section 3.1.

3.4 State Transition Systems for the Motivating Example

Using Kripke logical relations based on state transition systems, we can now roughly sketch the proof of equivalence of p and e .

We will prove that p and e are logically related in some initial world W_0 that includes some basic assumptions (in the form of STS’s) about registers, the stack, etc. (See Section 7 for details.)

First, since we can assume p has just been loaded into memory, we can think of its code as a freshly allocated piece of memory, and we are therefore given the opportunity to extend W_0 with an STS governing p ’s code. For most programs, we would extend W_0

at this point with a one-state STS, representing the simple invariant that the code of the program never changes. For our motivating example, we instead extend W_0 with an STS of the form:



When the STS is in the left state, the associated memory relation will require that p ’s code be in its initial, encrypted form, and when the STS is in the right state, the memory relation will require that p ’s code be in decrypted form. Once decrypted, always decrypted.

When p and e are executed, the former allocates a fresh closure on the heap, setting its constituent code pointer to $\text{bg} + 5$, and the latter allocates the local ref cell x , setting its contents to 0. Since both the closure and x are freshly allocated, we may at that point also extend the world with a new STS governing both of them:



We have joined the assumptions about them into one STS because they change in lockstep: when the functions returned by p and e are applied for the first time, x gets updated to 1 and the closure’s code pointer gets updated to $\text{bg} + 13$ simultaneously.

As noted in Section 3.2, it should never be the case that the code is still encrypted while the closure returned by p points to $\text{bg} + 13$. (The closure would behave in an unspecified manner if it were called in such a state.) When adding the second STS, it is therefore important that we outlaw this possibility up front so that we will not have to consider it later. Fortunately, Dreyer *et al.*’s model, on which our logical relation is based, allows us to define the second STS in such a way that we can only be in its right state if the first STS is also in its right state.

Let W be the world resulting from extending W_0 with the above two STS’s. What remains to be shown is that the functions returned by p and e are related in world W . So, suppose that W' is a future world of W , and that we begin executing the high- and low-level functions in some corresponding high- and low-level memories that are related by W' . There are three cases to consider, depending on the states of the two STS’s of interest in W' (the fourth case was outlawed, as described above):

Case 1: Code is encrypted, $x \mapsto 0$, closure $\mapsto \text{bg} + 5$. In this case, we first decrypt the code, and then set x to 1 and the closure’s code pointer to $\text{bg} + 13$.

Case 2: Code is decrypted, $x \mapsto 0$, closure $\mapsto \text{bg} + 5$. In this case, we set x to 1 and the closure’s code pointer to $\text{bg} + 13$.

Case 3: Code is decrypted, $x \mapsto 1$, closure $\mapsto \text{bg} + 13$. In this case, we set x to 1 and don’t touch the closure’s code pointer.

In all three cases, we end up transitioning to a future world W'' in which both STS’s of interest are in the right state, if they weren’t already there in W' . After making the state transition, both the high- and low-level functions invoke their callback arguments. Assuming they return, they will do so with memories that satisfy some future world of W'' , but in the STS’s of interest, there is nowhere to transition to. So we know that x must still point to 1. The high- and low-level functions must therefore both return the same result, namely 1, along with memories that satisfy a future world of the starting world W' .

3.5 Well-Bracketed State Changes and Private Transitions

For simplicity, we have glossed over many details in the above proof sketch. One important detail is how we reason about the stack. When the functions invoke their unknown callback arguments, we need to know that the callback will return the stack—*i.e.*, the contents of the stack segment up to the stack pointer sp —as it found it. One approach would be to bake this condition into

our logical relation for functions. However, as we explained in the introduction, we have set out to define our relation in a largely language-generic fashion, and it would not make sense to bake a low-level property about the stack into a language-generic relation.

Instead, we wish to build this condition into the initial world under which we relate high- and low-level programs, but the property we desire of the stack is not expressible in terms of STS’s as we have described them so far. To account for stack-like behavior, we employ another aspect of Dreyer *et al.*’s possible worlds, namely the idea of *private* vs. *public* transitions [14]. Private transitions were introduced in order to reason about so-called “well-bracketed” state changes, of which the behavior of the stack is a perfect example.

The basic idea is to label local state transitions as either public or private. Functions may make either private or public transitions internally, but viewed extensionally (*i.e.*, end-to-end), they must appear to make a public transition. In the STS that we use to reason about the stack, the states of the STS correspond to the possible states of the stack; every state is accessible from every other state by a private transition, but the only public transitions are self-transitions. This grants logically-related functions plenty of flexibility in how they manipulate the stack, but requires that, when they return, they leave the stack exactly how they found it.

We also use private transitions to reason about callee-save registers, which logically-related functions are expected to return as they found them, even if they modify them internally.

3.6 Reasoning in the Presence of Garbage Collection

Another important detail we have glossed over is how the presence of a garbage collector affects our proof of equivalence of p and e . Let us assume that we are using a standard mark-and-sweep or copying collector. The main effect such a collector has on our reasoning is that, whenever we pass control to the allocator (or to an unknown function that may call the allocator), we need to make sure that (1) all data that we care about being able to access in the future is in the reachable portion of the heap, and (2) there are no dangling pointers in the reachable portion of the heap.

Typically, guaranteeing these two conditions is straightforward. In our example, there is one call to the allocator (at $bg + 2$) and one call to an unknown function (at $bg + 17$). In the case of the former, there is nothing interesting to show. But in the case of the latter, it is important that when we increment the stack pointer before invoking the callback, we set the contents at the top of the stack to a value from which no dangling pointer may be reached. Here, we store in that stack slot the return address $[wk_0]$, which trivially satisfies this requirement.

So at an abstract level, reasoning in the presence of garbage collection is no big deal. What is more interesting is the technical question of how we actually implement this reasoning in the context of our logical relation. The central difficulty is that Kripke logical relations traditionally enjoy a *monotonicity* property, meaning that when two values are related in a world W , they are related in any future world of W . Monotonicity is essential when reasoning about unknown functions, such as the callback argument in our example. There, we were given the assumption that the callback argument was logically related in the world W' , but we did not actually invoke it until we had transitioned to the future world W'' . This step requires a use of monotonicity to show that the callback argument is still related in W'' .

Unfortunately, garbage collection seems superficially to throw a wrench into monotonicity. For instance, in our logical relation we want to be able to relate a **HIGH** pair value $\langle v_1, v_2 \rangle$ with a memory location pointing to a representation of that pair value on the heap. How can we expect those two value representations to be related in all future worlds if, at some point in the future, the memory location

may get deallocated (or its contents moved by a moving collector) and later reused for storing something else?

Our solution is to employ *logical memories*, which form a layer of abstraction over physical memories. Pointers in a logical memory are never moved or deallocated. Their connection to reality is established by a component of the logical memory called the *lookup table*, which specifies for any given logical pointer whether it is live and, if so, what physical pointer it corresponds to. We say that a logical memory \mathbf{M} *represents* a physical memory Φ only if \mathbf{M} ’s lookup table describes a bijection between the reachable portions of \mathbf{M} and Φ .

We maintain a global invariant on the logical memory, which must hold before and after calls to the allocator, requiring that all reachable data be live (and thus not dangling) according to the lookup table. Together with the definition of what it means for a logical memory to represent a physical one, this invariant guarantees the allocator’s precondition—(2) above—that there are no reachable dangling pointers. After the allocator returns, the lookup table of the logical memory may have completely changed—*e.g.*, due to a semi-space collection—but the only other change to the logical memory will be its extension with a freshly allocated logical pointer. Thus, if any data we care about was reachable prior to allocation—condition (1) above—it will still be reachable post-allocation, and by the global invariant it will also still be live.

With logical memories in hand, we can adapt our logical relations accordingly so that they relate values in **HIGH** with *logical values* (*i.e.*, logical pointers or non-pointer data) in **LOW**. We also define our possible worlds to impose invariants on *logical* memories, not physical ones. In this way, we regain monotonicity, as well as a clean, abstract account of memory locations that gives the garbage collector significant flexibility in how it implements them and allows us to essentially ignore how it does so.

4. A Language-Generic Kripke Logical Relation

Figure 3 defines a Kripke logical relation between two languages that is parameterized by abstract specifications for those languages ($\in \text{LangSpec}$), as well as by a specification for the possible worlds relating the memories of those languages ($\in \text{WorldSpec}$).

4.1 Language Specifications

A language specification (upper left of Figure 3) must provide sets of values (Val), computations (Com), continuations (Cont), memories (Mem), and configurations (Conf), together with a number of operations on these sets. (Note that we are assuming here a standard stratification on sets, and that Val , Com , etc. are “smaller” than LangSpec .) Most of these operations take elements of some of these sets and return a predicate on another of the sets. In most cases, this is because there may be a number of different representations of the same thing—*e.g.*, in **LOW**, a pair value $\langle v_1, v_2 \rangle$ is represented by a pointer that satisfies some conditions, but many pointers may satisfy those conditions.

“plugv” forms a configuration by plugging a value into a continuation under a given memory. “plugc” does the same, but plugging a computation instead of a value. “step” executes a configuration for one step of computation, resulting in either a new configuration, termination (*halt*), or failure (*fail*). “mdom” returns the domain of a memory, and “mdisj” takes two memories and returns a memory that contains their disjoint union, if it exists.

“oftype(τ)” determines whether a value is considered syntactically to have type τ under a certain memory. In our specification for **HIGH**, we include heap typings Σ in memories in order to define this predicate at ref type. For **LOW**, there is no notion of syntactic typing, but we find oftype convenient for expressing assumptions about the “syntactic” structure of closures (see Section 5).

$CType \stackrel{\text{def}}{=} \{ \tau \mid \text{ftv}(\tau) = \emptyset \}$
 $LangSpec \stackrel{\text{def}}{=} \{ (\text{Val}, \text{Com}, \text{Cont}, \text{Mem}, \text{Conf}, \text{plugv}, \text{plugc}, \text{step}, \text{mdom}, \text{mdisj}, \text{oftype}, \text{base}_b, \text{pair}, \text{app}, \text{appty}, \text{pack}, \text{roll}, \text{ref}, \text{asgn}) \mid$
 $\text{Val}, \text{Com}, \text{Cont}, \text{Mem}, \text{Conf} \in \text{Set} \wedge$
 $\text{plugv} \in \text{Val} \times \text{Cont} \times \text{Mem} \rightarrow \mathbb{P}(\text{Conf}) \wedge$
 $\text{plugc} \in \text{Com} \times \text{Cont} \times \text{Mem} \rightarrow \mathbb{P}(\text{Conf}) \wedge$
 $\text{step} \in \text{Conf} \rightarrow \text{Conf} \uplus \{ \text{fail}, \text{halt} \} \wedge$
 $\text{mdom} \in \text{Mem} \rightarrow \mathbb{P}(\text{Val}) \wedge$
 $\text{mdisj} \in \text{Mem} \times \text{Mem} \rightarrow \mathbb{P}(\text{Mem}) \wedge$
 $\text{oftype} \in CType \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{base}_b \in \llbracket b \rrbracket \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{pair} \in \text{Val} \times \text{Val} \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{app} \in \text{Val} \times \text{Val} \rightarrow \mathbb{P}(\text{Com}) \wedge$
 $\text{appty} \in \text{Val} \times CType \rightarrow \mathbb{P}(\text{Com}) \wedge$
 $\text{pack} \in CType \times \text{Val} \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{roll} \in \text{Val} \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{ref} \in \text{Val} \rightarrow \mathbb{P}(\text{Val} \times \text{Mem}) \wedge$
 $\text{asgn} \in \text{Mem} \times \text{Val} \times \text{Val} \rightarrow \text{Mem} \wedge$
 $\forall M_1, M_2. \forall M \in \text{mdisj}(M_1, M_2).$
 $\text{mdom}(M) \supseteq \text{mdom}(M_1) \uplus \text{mdom}(M_2) \}$

For $\mathcal{L}_1, \mathcal{L}_2 \in LangSpec,$
 $WorldSpec \stackrel{\text{def}}{=} \{ (\text{World}, \text{lev}, \mathcal{M}, \mathcal{B}, \mathcal{O}, \triangleright, \sqsupseteq, \sqsupseteq_{\text{pub}}) \mid$
 $\text{World} \in \text{Set} \wedge$
 $\text{lev} \in \text{World} \rightarrow \mathbb{N} \wedge$
 $\mathcal{M} \in \text{World} \rightarrow \mathbb{P}(\mathcal{L}_1.\text{Mem} \times \mathcal{L}_2.\text{Mem}) \wedge$
 $\mathcal{B} \in \text{World} \rightarrow \mathbb{P}(\mathcal{L}_1.\text{Val} \times \mathcal{L}_2.\text{Val}) \wedge$
 $\mathcal{O} \in \text{World} \rightarrow \mathbb{P}(\mathcal{L}_1.\text{Conf} \times \mathcal{L}_2.\text{Conf}) \wedge$
 $\triangleright \in \text{World} \rightarrow \text{World} \wedge$
 $\sqsupseteq \in \mathbb{P}(\text{World} \times \text{World}) \wedge$
 $\sqsupseteq_{\text{pub}} \in \mathbb{P}(\text{World} \times \text{World}) \wedge$
 $\sqsupseteq, \sqsupseteq_{\text{pub}} \text{ are preorders } \wedge \sqsupseteq_{\text{pub}} \subseteq \sqsupseteq \wedge$
 $\forall W' \sqsupseteq W. \triangleright W' \sqsupseteq \triangleright W \wedge$
 $\forall W' \sqsupseteq_{\text{pub}} W. \triangleright W' \sqsupseteq_{\text{pub}} \triangleright W \wedge$
 $\forall W. \triangleright W \sqsupseteq_{\text{pub}} W \wedge$
 $\forall W' \sqsupseteq W. \text{lev}(W') \leq \text{lev}(W) \wedge$
 $\forall W. \text{lev}(W) > 0 \implies \text{lev}(\triangleright W) = \text{lev}(W) - 1 \}$

$\triangleright \stackrel{\text{def}}{=} \{ (W', W) \mid \text{lev}(W) > 0 \wedge W' \sqsupseteq \triangleright W \}$
 $WVRel \stackrel{\text{def}}{=} \{ R \in \mathbb{P}(\text{World} \times \mathcal{L}_1.\text{Val} \times \mathcal{L}_2.\text{Val}) \}$
 $R(W) \stackrel{\text{def}}{=} \{ (v_1, v_2) \mid (W, v_1, v_2) \in R \}$
 $\triangleright R \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \mid \text{lev}(W) > 0 \implies (\triangleright W, v_1, v_2) \in R \}$
 $\square R \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \mid \forall W' \sqsupseteq W. (W', v_1, v_2) \in R \}$
 $R_{\triangleright W} \stackrel{\text{def}}{=} \{ (W', v_1, v_2) \mid W' \triangleright W \wedge (W', v_1, v_2) \in R \}$
 $(R_1, R_2) \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \mid \forall (M_1, M_2) \in \mathcal{M}(W). (v_1, M_1) \in R_1 \wedge (v_2, M_2) \in R_2 \}$
 for $R_1 \in \mathbb{P}(\mathcal{L}_1.\text{Val} \times \mathcal{L}_1.\text{Mem}), R_2 \in \mathbb{P}(\mathcal{L}_2.\text{Val} \times \mathcal{L}_2.\text{Mem})$
 $TyValRel \stackrel{\text{def}}{=} \{ (\tau_1, \tau_2, R) \mid \tau_1, \tau_2 \in CType \wedge R \in WVRel \}$
 $\rho \in \text{TypeVar} \rightarrow TyValRel$
 $\rho.1(\tau) \stackrel{\text{def}}{=} \tau[\rho(\alpha).\tau_1/\alpha] \quad \rho.2(\tau) \stackrel{\text{def}}{=} \tau[\rho(\alpha).\tau_2/\alpha]$
 $\text{oftype}(\tau, \rho) \stackrel{\text{def}}{=} \square(\mathcal{L}_1.\text{oftype}(\rho.1(\tau)), \mathcal{L}_2.\text{oftype}(\rho.2(\tau)))$
 $\mathcal{V}[\llbracket \alpha \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\alpha, \rho) \mid (W, v_1, v_2) \in \square \rho(\alpha).R \}$
 $\mathcal{V}[\llbracket b \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(b, \rho) \mid \exists x \in \llbracket b \rrbracket.$
 $(W, v_1, v_2) \in \square(\mathcal{L}_1.\text{base}_b(x), \mathcal{L}_2.\text{base}_b(x)) \}$
 $\mathcal{V}[\llbracket \tau \times \tau' \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\tau \times \tau', \rho) \mid$
 $\exists (u_1, u_2) \in \triangleright \mathcal{V}[\llbracket \tau \rrbracket \rho](W). \exists (u'_1, u'_2) \in \triangleright \mathcal{V}[\llbracket \tau' \rrbracket \rho](W).$
 $(W, v_1, v_2) \in \square(\mathcal{L}_1.\text{pair}(u_1, u'_1), \mathcal{L}_2.\text{pair}(u_2, u'_2)) \}$
 $\mathcal{V}[\llbracket \tau' \rightarrow \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\tau' \rightarrow \tau, \rho) \mid \forall W' \sqsupseteq W. \forall (u_1, u_2) \in \mathcal{V}[\llbracket \tau' \rrbracket \rho](W').$
 $\forall e_1 \in \mathcal{L}_1.\text{app}(v_1, u_1). \forall e_2 \in \mathcal{L}_2.\text{app}(v_2, u_2). (W', e_1, e_2) \in \mathcal{E}[\llbracket \tau \rrbracket \rho] \}$
 $\mathcal{V}[\llbracket \forall \alpha. \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\forall \alpha. \tau, \rho) \mid \forall W' \sqsupseteq W. \forall (\tau_1, \tau_2, R) \in TyValRel.$
 $\forall e_1 \in \mathcal{L}_1.\text{appty}(v_1, \tau_1). \forall e_2 \in \mathcal{L}_2.\text{appty}(v_2, \tau_2).$
 $(W', e_1, e_2) \in \mathcal{E}[\llbracket \tau \rrbracket \rho][\alpha \mapsto (\tau_1, \tau_2, R)] \}$
 $\mathcal{V}[\llbracket \exists \alpha. \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\exists \alpha. \tau, \rho) \mid$
 $\exists (\tau_1, \tau_2, R) \in TyValRel. \exists (u_1, u_2) \in \mathcal{V}[\llbracket \tau \rrbracket \rho][\alpha \mapsto (\tau_1, \tau_2, R)](W).$
 $(W, v_1, v_2) \in \square(\mathcal{L}_1.\text{pack}(\tau_1, u_1), \mathcal{L}_2.\text{pack}(\tau_2, u_2)) \}$
 $\mathcal{V}[\llbracket \text{ref } \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, v_1, v_2) \in \text{oftype}(\text{ref } \tau, \rho) \mid \forall W' \sqsupseteq W. \forall (M_1, M_2) \in \mathcal{M}(W').$
 $(v_1, v_2) \in \mathcal{B}(W') \wedge$
 $(\exists (u_1, u_2) \in \triangleright \mathcal{V}[\llbracket \tau \rrbracket \rho](W')). (v_1, M_1) \in \mathcal{L}_1.\text{ref}(u_1) \wedge (v_2, M_2) \in \mathcal{L}_2.\text{ref}(u_2) \} \wedge$
 $(\forall (u_1, u_2) \in \triangleright \mathcal{V}[\llbracket \tau \rrbracket \rho](W')). (\mathcal{L}_1.\text{asgn}(M_1, v_1, u_1), \mathcal{L}_2.\text{asgn}(M_2, v_2, u_2)) \in \mathcal{M}(W') \}$
 $\mathcal{V}[\llbracket \mu \alpha. \tau \rrbracket \rho] \stackrel{\text{def}}{=} \mu(F_{\alpha, \tau, \rho})$
 $F_{\alpha, \tau, \rho} \stackrel{\text{def}}{=} \lambda R. \{ (W, v_1, v_2) \in \text{oftype}(\mu \alpha. \tau, \rho) \mid$
 $\exists (u_1, u_2) \in \mathcal{V}[\llbracket \tau \rrbracket \rho][\alpha \mapsto (\rho.1(\mu \alpha. \tau), \rho.2(\mu \alpha. \tau), R)](W).$
 $(W, v_1, v_2) \in \square(\mathcal{L}_1.\text{roll}(u_1), \mathcal{L}_2.\text{roll}(u_2)) \}$
 $\mu(F)(W) \stackrel{\text{def}}{=} F(\mu(F)_{\triangleright W})(W)$
 $\mathcal{K}[\llbracket \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, K_1, K_2) \in \text{World} \times \mathcal{L}_1.\text{Cont} \times \mathcal{L}_2.\text{Cont} \mid \forall W' \sqsupseteq_{\text{pub}} W.$
 $\forall (v_1, v_2) \in \mathcal{V}[\llbracket \tau \rrbracket \rho](W'). \forall (M_1, M_2) \in \mathcal{M}(W').$
 $\forall C_1 \in \mathcal{L}_1.\text{plugv}(v_1, K_1, M_1). \forall C_2 \in \mathcal{L}_2.\text{plugv}(v_2, K_2, M_2).$
 $(C_1, C_2) \in \mathcal{O}(W') \}$
 $\mathcal{E}[\llbracket \tau \rrbracket \rho] \stackrel{\text{def}}{=} \{ (W, e_1, e_2) \in \text{World} \times \mathcal{L}_1.\text{Com} \times \mathcal{L}_2.\text{Com} \mid$
 $\forall (K_1, K_2) \in \mathcal{K}[\llbracket \tau \rrbracket \rho](W). \forall (M_1, M_2) \in \mathcal{M}(W).$
 $\forall C_1 \in \mathcal{L}_1.\text{plugc}(e_1, K_1, M_1). \forall C_2 \in \mathcal{L}_2.\text{plugc}(e_2, K_2, M_2).$
 $(C_1, C_2) \in \mathcal{O}(W) \}$

Figure 3. Language Specifications, World Specifications, and a Language-Generic Kripke Logical Relation

The remaining operations define the various syntactic language forms that are referenced in the definition of the logical relation. One point of note: determining whether a value represents a particular canonical form may require one to consider an accompanying memory. It is easy to see why—*e.g.*, one cannot determine if a pointer to a pair of heap cells represents $\langle v_1, v_2 \rangle$ without inspecting the heap. Determining whether a computation represents a function application also may require inspecting the memory, but we have found it technically more convenient to assume that any memory inspection is somehow built into the notion of computation (see Section 5 to see how this works in the **LOW** specification).

4.2 World Specifications

A world specification (lower left of Figure 3) defines a set of possible worlds relating the memories of two languages specified by \mathcal{L}_1 and \mathcal{L}_2 , together with a variety of operations on and relations indexed by those worlds. In order to understand some of these, it is necessary to first say a word about our use of *step-indexing*.

Appel and McAllester [3] introduced step-indexing in order to model recursive types in foundational proof-carrying code. The basic idea is to use a natural number index (“step level”) to stratify what would otherwise be a circular construction. Step-indexing has also proven useful in modeling other semantically circular notions like higher-order state, which is how we use them here. Due to space considerations, since our step-indexed construction follows closely that of [2] and [14], we refer the interested reader to those previous works for the relevant background.

In the world specification, the important step-indexing-related bits are the lev function, which returns the step level of a given world, and the \triangleright (pronounced “later”) operator, which returns an approximated version of the given world at one lower step level. We use \triangleright to ensure well-foundedness of the logical relation.

“ $\mathcal{M}(W)$ ” is the memory relation associated with W , which specifies when two memories (from \mathcal{L}_1 and \mathcal{L}_2) satisfy the constraints of the STS’s in W . “ $\mathcal{B}(W)$ ” is a bijection on values representing memory locations in \mathcal{L}_1 and \mathcal{L}_2 . This is used in defining the logical relation for ref type. “ $\mathcal{O}(W)$ ” is an observation relation

on configurations. For the possible worlds we employ in this paper, $\mathcal{O}(W)$ actually only depends on $\text{lev}(W)$, and it is defined to relate configurations that either both terminate (without failure) or that both run for at least $\text{lev}(W)$ steps of computation. When we prove that two programs are logically related, we will prove it in starting worlds of an arbitrary step level, thus ensuring that the programs are observably equivalent for arbitrarily many computation steps.

Finally, “ \sqsupset ” defines the general “future world” relation between worlds, and “ \sqsupset_{pub} ” defines a restricted “public” version of that relation: if $W' \sqsupset_{\text{pub}} W$, then for any STS in W , the new state of that STS in W' must be accessible from its old state in W only by public transitions (see Section 3.5). Both \sqsupset and \sqsupset_{pub} are preorders.

4.3 Kripke Logical Relation

The right side of Figure 3 displays our Kripke logical relation, whose definition is parametric w.r.t. \mathcal{L}_1 , \mathcal{L}_2 , and an instance of WorldSpec thereon. In the definition, we adopt the convention that the entities (values, continuations, etc.) from \mathcal{L}_1 appear in boldface (\mathbf{v} , \mathbf{K} , etc.) and the entities from \mathcal{L}_2 appear in italics (v , K , etc.). The coincidence of the notation for \mathcal{L}_2 entities with the notation for the corresponding entities from **HIGH** is deliberate, for in the next section we will instantiate \mathcal{L}_1 and \mathcal{L}_2 with our specifications for **LOW** and **HIGH**, respectively. We abuse notation in this way in order to avoid the proliferation of more than two fonts.

Our logical relation is based very closely on Dreyer *et al.*'s [14], with the principal difference being that the relevant linguistic forms have been abstracted away in the language specifications \mathcal{L}_1 and \mathcal{L}_2 . For instance, in the logical relation for arrow types, we do not construct the applications $\mathbf{v}_1 \mathbf{u}_1$ and $v_2 u_2$ directly, since \mathcal{L}_1 and \mathcal{L}_2 may not include an explicit application construct. Rather, we quantify over arbitrary computations e_1 and e_2 drawn from $\mathcal{L}_1.\text{app}(\mathbf{v}_1, \mathbf{u}_1)$ and $\mathcal{L}_2.\text{app}(v_2, u_2)$, respectively.

The logical relation consists of a relation for values ($\mathcal{V}[\tau]\rho$), one for continuations ($\mathcal{K}[\tau]\rho$), and one for computations ($\mathcal{E}[\tau]\rho$). Here, we assume that ρ is a relational interpretation of the free variables of τ , mapping them to arbitrary world-indexed value relations. For $\tau = \alpha$, $\mathcal{V}[\tau]\rho$ is defined as the restriction of $\rho(\alpha)$ to triples (W, \mathbf{v}_1, v_2) where \mathbf{v}_1 and v_2 are “well-typed” (according to $\mathcal{L}_i.\text{oftype}$) and continue to be related in all future worlds of W . This last part, which is specified using the $\square R$ operator defined at the top right of the figure, is key to ensuring monotonicity of the value relation. The pair and existential cases of the value relation also use the $\square R$ operator in order to ensure monotonicity of data representations—*e.g.*, that if \mathbf{v}_1 represents a pair of \mathbf{u}_1 and \mathbf{u}'_1 , it will continue to do so in all future worlds.

As in Appel *et al.* [4], the interpretation of recursive types is defined by induction on the “strictly future world” relation \sqsupset_{p} . This relation is well-founded because $W' \sqsupset_{\text{p}} W$ implies that W' has a lower step level than W . By defining the recursive type case this way, we can relate **HIGH** programs, where roll and unroll are explicit coercions, to **LOW** programs where they have been erased.

The ref type case relates two memory locations if dereferencing, assigning and testing them for pointer equality will always produce related results. The condition on pointer equality testing is guaranteed by the requirement that the locations be in the bijection of the world in which they are related.

The value relation is lifted to a relation on computations by the technique of *biorthogonality* (aka $\top\top$ -closure) [22, 17]. The idea is to define two computations to be related if they behave in an observably equivalent manner when plugged into related continuations. Two continuations are in turn related if they behave in an observably equivalent manner when plugged with related values. By quantifying only over public future worlds in the definition of $\mathcal{K}[\tau]\rho$, we ensure that computations may only make public transitions when viewed end-to-end, as per the discussion in Section 3.5.

This mind-bending technique is well-suited to languages where the evaluation of a computation is *context-sensitive* in the sense that it cannot be performed in ignorance of its continuation. Such is the case with **LOW**, where a computation always ends with a jump to its return address. As we shall see, the fact that the return address really is a valid return address, which is part of the contract between computation and continuation, will be encoded in the **LOW** implementation of plug that we define in Section 5.

5. Implementing the Specifications

In order to instantiate our logical relations to relate **LOW** and **HIGH** entities, we must first show how to implement the abstract LangSpec interface for both languages.

For **HIGH**, the implementation of the interface is almost entirely straightforward, as all the required entities (values, computations) have direct correspondents in the **HIGH** language. The only slightly unusual bit is that we define **HIGH** memories to be pairs of heaps *and* heap typings Σ . The inclusion of the heap typing is necessary for defining oftype. For the remaining details, please see the companion technical appendix [15].

Low-Level Entities The implementation of LangSpec for **LOW** (Figure 4) is much more interesting. As described in Section 3.6, we employ a notion of *logical values* \mathbf{v} , which are either non-pointer words \underline{w} or logical pointers \hat{l} . Logical Lvalues are similar to physical PLvalues except that logical Lvalues include logical heap locations $(l : o)_h$ in place of physical ones $\langle a \rangle_h$. We include the offset o because the logical heap is (for convenience) modeled two-dimensionally as a list of blocks.

LOW computations e are 4-tuples $(\text{cpc}, \text{kpc}, \text{vloc}, \text{data})$, where: cpc is the code address where the computation begins, kpc is the return address, vloc is the Lvalue where the return value will be stored, and data is a memory predicate that must be satisfied in order for the computation to be correctly executed. **LOW** continuations \mathbf{K} are pairs $(\text{kpc}, \text{vloc})$, where: kpc is the code address where the continuation begins, and vloc is the Lvalue where the input to the continuation should be placed. It is worth noting that, unlike $e.\text{cpc}$ and $e.\text{kpc}$, the kpc in \mathbf{K} must be a code address, not an Rvalue, because when we plug an e into a \mathbf{K} , we need to know that the value of $\mathbf{K}.\text{kpc}$ will be the same before and after the execution of e . Were $\mathbf{K}.\text{kpc}$ an Rvalue, we would not know that.

Logical memories \mathbf{M} are 6-tuples $(\text{code}, \text{reg}, \text{stk}, \text{hp}, \text{tab}, \text{shp})$, where: code is the code segment, reg is the register file (minus sp since it is determined by the size of the stack), stk is the stack, hp is the heap, tab is the lookup table (described in Section 3.6), and shp is the *system heap*, which is a separate portion of the heap controlled by the runtime system. The lookup table maps each logical pointer to a physical pointer and the size of the memory block starting at that address. The pointer is live iff the size is > 0 . Note that reg , stk , and hp are all maps from various Lvalues to *logical* values, whereas the tab and shp are maps to physical values. In the proofs, we end up treating tab and shp as essentially black boxes, since they can be changed at whim by the allocator, whereas the allocator should not mess around with the logical portion of the memory represented by reg , stk , and hp .

Lastly, note that reg , stk , tab , and shp may all be undefined (undef). This is a useful technical device for defining the disjoint union of several partial memories: it enables us to specify that only one of those partial memories contains information about, say, the stack. (See the definition of “ mdisj ” below.)

Low-Level Representations of High-Level Constructs The bottom left of Figure 4 defines **LOW** representations of various high-level constructs, as required by the LangSpec interface. A pair of \mathbf{v}_1 and \mathbf{v}_2 is represented as a pointer to a pair of cells containing \mathbf{v}_1

Loc	$\stackrel{\text{def}}{=} \{l \in \mathbb{N}\}$	$ v \stackrel{\text{def}}{=} \begin{cases} w & \text{if } v = \underline{w} \\ l & \text{if } v = \widehat{l} \end{cases}$
Word	$\stackrel{\text{def}}{=} \{w \in \mathbb{N}\}$	$M(v) \stackrel{\text{def}}{=} v \quad M([r]) \stackrel{\text{def}}{=} M.\text{reg}(r)$
$v \in \text{Val}$	$::= \underline{w} \mid \widehat{l}$	$M(\langle a \rangle_s) \stackrel{\text{def}}{=} M.\text{stk}(a) \quad M(\langle r - o \rangle_s) \stackrel{\text{def}}{=} M.\text{stk}(M.\text{reg}(r) - o)$
$lv \in \text{Lvalue}$	$::= [r] \mid \langle a \rangle_s \mid \langle r - o \rangle_s \mid \langle l : o \rangle_h \mid \langle r + o \rangle_h$	$M(\langle l : o \rangle_h) \stackrel{\text{def}}{=} M.\text{hp}(l)(o) \quad M(\langle r + o \rangle_h) \stackrel{\text{def}}{=} M.\text{hp}(M.\text{reg}(r))(o)$
$rv \in \text{Rvalue}$	$::= lv \mid v$	$M[[T, S]] \stackrel{\text{def}}{=} (M.\text{code}, M.\text{reg}, M.\text{stk}, M.\text{hp}, T, S)$
Com	$\stackrel{\text{def}}{=} \{e = (\text{cpc}, \text{kpc}, \text{vloc}, \text{data}) \in \text{Rvalue} \times \text{Rvalue} \times \text{Lvalue} \times \mathbb{P}(\text{Mem})\}$	$M[l : o \mapsto v]_{\text{hp}} \stackrel{\text{def}}{=} (M.\text{code}, M.\text{reg}, M.\text{stk}, M.\text{hp}[l \mapsto (v_0, \dots, v_{o-1}, v, v_{o+1}, \dots, v_{n-1})], M.\text{tab}, M.\text{shp})$
Cont	$\stackrel{\text{def}}{=} \{K = (\text{kpc}, \text{vloc}) \in \text{PAddr} \times \text{Lvalue}\}$	if $M.\text{hp}(l) = (v_0, \dots, v_{n-1}) \wedge o < n$
CodeFrag	$\stackrel{\text{def}}{=} \text{PAddr} \rightarrow_{\text{fin}} \text{Instruction}$	
RegFile	$\stackrel{\text{def}}{=} (\text{Register} \setminus \{\text{sp}\} \rightarrow \text{Val}) \uplus \{\text{undef}\}$	
List X	$\stackrel{\text{def}}{=} \{(x_0, \dots, x_{n-1}) \mid n \in \mathbb{N} \wedge x_0, \dots, x_{n-1} \in X\}$	
Stack	$\stackrel{\text{def}}{=} \text{List Val} \uplus \{\text{undef}\}$	
Heap	$\stackrel{\text{def}}{=} \text{Loc} \rightarrow_{\text{fin}} \text{List Val}$	
Table	$\stackrel{\text{def}}{=} (\text{Loc} \rightarrow_{\text{fin}} \mathbb{N} \times \text{PAddr}) \uplus \{\text{undef}\}$	
SysHeap	$\stackrel{\text{def}}{=} (\text{PAddr} \rightarrow \text{Word}) \uplus \{\text{undef}\}$	
Mem	$\stackrel{\text{def}}{=} \{M = (\text{code}, \text{reg}, \text{stk}, \text{hp}, \text{tab}, \text{shp}) \in \text{CodeFrag} \times \text{RegFile} \times \text{Stack} \times \text{Heap} \times \text{Table} \times \text{SysHeap}\}$	
Conf	$\stackrel{\text{def}}{=} \text{PConf}$	
$\text{oftype}(\tau)$	$\stackrel{\text{def}}{=} \{(v, M) \in \text{Val} \times \text{Mem} \mid \forall \tau_1, \tau_2. \tau = \tau_1 \rightarrow \tau_2 \implies \exists l, w. v = \widehat{l} \wedge M.\text{hp}(l)(0) = \underline{w} \wedge \forall \alpha, \tau'. \tau = \forall \alpha. \tau' \implies \exists l, w. v = \widehat{l} \wedge M.\text{hp}(l)(0) = \underline{w}\}$	
$\text{base}_b(x)$	$\stackrel{\text{def}}{=} \{(v, M) \in \text{Val} \times \text{Mem} \mid v \text{ is a representation of } x\}$	
$\text{pair}(v_1, v_2)$	$\stackrel{\text{def}}{=} \{(v, M) \in \text{Val} \times \text{Mem} \mid \exists l. v = \widehat{l} \wedge M.\text{hp}(l)(0) = v_1 \wedge M.\text{hp}(l)(1) = v_2\}$	
$\text{app}(v_1, v_2)$	$\stackrel{\text{def}}{=} \{e \in \text{Com} \mid \exists l. v_1 = \widehat{l} \wedge e.\text{cpc} = \langle l : 0 \rangle_h \wedge e.\text{kpc} = [wk_0] \wedge e.\text{vloc} = [wk_5] \wedge e.\text{data} = \{M \in \text{Mem} \mid M.\text{reg}(wk_1) = v_1 \wedge M.\text{reg}(wk_2) = v_2\}\}$	
$\text{appty}(v, \tau)$	$\stackrel{\text{def}}{=} \{e \in \text{Com} \mid \exists l. v = \widehat{l} \wedge e.\text{cpc} = \langle l : 0 \rangle_h \wedge e.\text{kpc} = [wk_0] \wedge e.\text{vloc} = [wk_5] \wedge e.\text{data} = \{M \in \text{Mem} \mid M.\text{reg}(wk_1) = v\}\}$	
$\text{pack}(\tau, v)$	$\stackrel{\text{def}}{=} \{(v', M) \in \text{Val} \times \text{Mem} \mid v' = v\}$	
$\text{roll}(v)$	$\stackrel{\text{def}}{=} \{(v', M) \in \text{Val} \times \text{Mem} \mid v' = v\}$	
$\text{ref}(v)$	$\stackrel{\text{def}}{=} \{(v', M) \in \text{Val} \times \text{Mem} \mid \exists l. v' = \widehat{l} \wedge M.\text{hp}(l)(0) = v\}$	
$\text{asgn}(M, v_1, v_2)$	$\stackrel{\text{def}}{=} \begin{cases} M[l : 0 \mapsto v_2]_{\text{hp}} & \text{if } v_1 = \widehat{l} \wedge M.\text{hp}(l) > 0 \\ \text{undef} & \text{otherwise} \end{cases}$	
$\text{plugv}(v, K, M)$	$\stackrel{\text{def}}{=} \{(\Phi, \text{pc}) \in \text{Conf} \mid M \text{ repr } \Phi \wedge \text{pc} = K.\text{kpc} \wedge M(K.\text{vloc}) = v\}$	
$\text{plugc}(e, K, M)$	$\stackrel{\text{def}}{=} \{(\Phi, \text{pc}) \in \text{Conf} \mid M \text{ repr } \Phi \wedge M \in e.\text{data} \wedge \text{pc} = M(e.\text{cpc}) \wedge M(e.\text{kpc}) = K.\text{kpc} \wedge e.\text{vloc} = K.\text{vloc}\}$	
$\text{step}(\Phi, \text{pc})$	$\stackrel{\text{def}}{=} R \quad \text{with } (\Phi, \text{pc}) \leftrightarrow R$	
$\text{mdom}(M)$	$\stackrel{\text{def}}{=} \{\widehat{l} \in \text{Val} \mid l \in \text{dom}(M.\text{hp})\}$	
$\text{mdisj}(M_1, M_2)$	$\stackrel{\text{def}}{=} \{M \in \text{Mem} \mid M.\text{code} \supseteq M_1.\text{code} \uplus M_2.\text{code} \wedge M.\text{hp} \supseteq M_1.\text{hp} \uplus M_2.\text{hp} \wedge \text{nosh}(M.\text{reg}, M_1.\text{reg}, M_2.\text{reg}) \wedge \text{nosh}(M.\text{stk}, M_1.\text{stk}, M_2.\text{stk}) \wedge \text{nosh}(M.\text{tab}, M_1.\text{tab}, M_2.\text{tab}) \wedge \text{nosh}(M.\text{shp}, M_1.\text{shp}, M_2.\text{shp})\}$	
$\text{nosh}(X, X_1, X_2)$	$\stackrel{\text{def}}{=} \begin{cases} X_1 \neq \text{undef} \implies X_2 = \text{undef} \wedge X = X_1 \\ X_2 \neq \text{undef} \implies X_1 = \text{undef} \wedge X = X_2 \end{cases}$	

Figure 4. The Implementation of LangSpec for LOW

and v_2 . $\text{pack}(\tau, v)$ and $\text{roll}(v)$ are represented the same as v . References are represented directly as pointers. We also use $\text{oftype}(\tau)$ to enforce that values of arrow and universal type are represented by pointers to closures whose first cell is *not* a logical pointer. (This ensures that we can jump to the code address directly.)

The most interesting bit is the representation of $\text{app}(v_1, v_2)$ (and $\text{appty}(v, \tau)$, which is similar). In order for a computation e to represent this application, v_1 is assumed to be a pointer \widehat{l} to a closure. The starting address of the function application ($e.\text{cpc}$) is thus taken to be the code address stored in the first cell of the closure, *i.e.*, $\langle l : 0 \rangle_h$. Our calling convention is that, when the function is called, the return address should be stored in wk_0 , and when the function returns, the return value should be stored in wk_5 , so $e.\text{kpc}$ and $e.\text{vloc}$ reflect this convention. Finally, the memory predicate $e.\text{data}$ requires that when control is passed to $e.\text{cpc}$, the function v_1 and argument v_2 are stored in wk_1 and wk_2 .

Connecting Logical and Physical Memories The right side of Figure 4 defines the remaining elements of LangSpec, along with a number of auxiliary operations. The operations at the top right give shorthand for various lookup and update operations on logical memories. “ $\text{phyv}(M)(v)$ ” returns the physical interpretation of v according to M ’s lookup table, if one exists, and “ $\text{phyh}(M)$ ” returns the live portion of the physical heap according to M ’s lookup table. Note that the definition of phyh demands that the physical representations of logical memory blocks with distinct

head pointers be disjoint, thus ensuring a proper bijection between the reachable parts of the physical and logical heaps.

“ $M \text{ repr } \Phi$ ” says that M is a valid logical abstraction of Φ . The definition is fairly straightforward, making use of phyv and phyh as one would expect. The fourth line of the definition guarantees that the reachable heap is disjoint from the system heap, and the fifth condition just checks that the block sizes specified for live data in $M.\text{tab}$ are correct. Note that Φ may contain arbitrary other junk (in the code, stack, and heap segments) that is not described by M .

Plugging Continuations Using $M \text{ repr } \Phi$, it is easy to specify how to plug continuations with values and computations. A configuration (Φ, pc) belongs to $\text{plugv}(v, K, M)$ if (1) M is a valid abstraction of Φ , (2) the program counter pc is set to the starting address of the continuation ($K.\text{kpc}$), and (3) the value v is stored in the location where the continuation is expecting it ($K.\text{vloc}$).

A configuration (Φ, pc) belongs to $\text{plugc}(e, K, M)$ if (1) M is a valid abstraction of Φ , (2) M satisfies the memory constraints demanded by e , (3) the program counter pc is set to the starting address of the computation ($e.\text{cpc}$), (4) the starting address of the continuation ($K.\text{kpc}$) is stored in the place where the computation is expecting to find its return address ($e.\text{kpc}$), and (5) the place where e will store its return value ($e.\text{vloc}$) is the same place where K is expecting to find its input value ($K.\text{vloc}$).

The remaining definitions (of step , mdom , and mdisj) are fairly self-explanatory. As mentioned earlier, the definition of

$$\begin{array}{l}
[\text{bg} \Rightarrow \text{instrs}] \stackrel{\text{def}}{=} [\text{bg} \mapsto \text{instrs}(0), \dots, \text{instrs}(|\text{instrs}| - 1)] \\
\{C\}_{\text{code}} \stackrel{\text{def}}{=} (C, \text{undef}, \text{undef}, \emptyset, \text{undef}, \text{undef}) \in \text{Mem} \\
\{H\}_{\text{heap}} \stackrel{\text{def}}{=} (\emptyset, \text{undef}, \text{undef}, H, \text{undef}, \text{undef}) \in \text{Mem} \\
\mathbf{v} \text{ live in } \mathbf{M} \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } \mathbf{v} = \underline{w} \\ \exists n, a. \mathbf{M}.\text{tab}(\mathbf{l}) = (n, a) \wedge n > 0 & \text{if } \mathbf{v} = \widehat{\mathbf{l}} \end{cases} \\
\text{reach}_0(\mathbf{M}) \stackrel{\text{def}}{=} \{ \mathbf{l} \mid \exists r \in \text{Register}. \widehat{\mathbf{l}} = \mathbf{M}.\text{reg}(r) \} \cup \\
\{ \mathbf{l} \mid \exists j < |\mathbf{M}.\text{stk}|. \widehat{\mathbf{l}} = \mathbf{M}.\text{stk}(j) \} \\
\text{reach}_{i+1}(\mathbf{M}) \stackrel{\text{def}}{=} \text{reach}_i(\mathbf{M}) \cup \\
\{ \mathbf{l} \mid \exists \mathbf{l}' \in \text{reach}_i(\mathbf{M}). \exists j. \widehat{\mathbf{l}} = \mathbf{M}.\text{hp}(\mathbf{l}')(j) \} \\
\text{reach}(\mathbf{M}) \stackrel{\text{def}}{=} \bigcup_{i \in \mathbb{N}} \text{reach}_i(\mathbf{M}) \\
\text{AllocSpec} \stackrel{\text{def}}{=} \\
\{ \mathcal{A} \in \text{PAddr} \rightarrow \{ (\text{init}, \text{alloc}, \text{instrs}, I) \\
\quad \in \text{PAddr} \times \text{PAddr} \times \text{List Instruction} \times \\
\quad \mathbb{P}(\text{Table} \times \text{SysHeap}) \} \mid \\
\forall \text{gcbg}, \Phi, \text{pc}. \\
\Phi.\text{code} \supseteq [\text{gcbg} \Rightarrow \mathcal{A}(\text{gcbg}).\text{instrs}] \wedge \Phi.\text{reg}(\text{wk}_4) = \underline{\text{pc}} \implies \\
\exists \mathbf{M}', \Phi'. \\
(\Phi, \mathcal{A}(\text{gcbg}).\text{init}) \xrightarrow{*} (\Phi', \text{pc}) \wedge \\
\Phi'.\text{code} = \Phi.\text{code} \wedge \mathbf{M}'.\text{code} = [\text{gcbg} \Rightarrow \mathcal{A}(\text{gcbg}).\text{instrs}] \wedge \\
\mathbf{M}' \text{ repr } \Phi' \wedge \mathbf{M}' \in \mathcal{A}.\text{GR}(\text{gcbg}) \wedge \mathbf{M}' \in \mathcal{A}.\text{MR}(\text{gcbg}) \wedge \\
\forall \text{gcbg}, \mathbf{M}, \Phi, \text{pc}, n. \\
\mathbf{M} \text{ repr } \Phi \wedge \mathbf{M} \in \mathcal{A}.\text{GR}(\text{gcbg}) \wedge \mathbf{M} \in \mathcal{A}.\text{MR}(\text{gcbg}) \wedge \\
\mathbf{M}.\text{reg}(\text{wk}_4) = \underline{\text{pc}} \wedge \mathbf{M}.\text{reg}(\text{wk}_5) = \underline{n} \implies \\
\exists \Phi', \mathbf{M}', T, S, w, \mathbf{l}, w_0, \dots, w_{n-1}. \\
(\Phi, \mathcal{A}(\text{gcbg}).\text{alloc}) \xrightarrow{*} (\Phi', \text{pc}) \wedge \\
\mathbf{M}' \text{ repr } \Phi' \wedge \mathbf{M}' \in \mathcal{A}.\text{GR}(\text{gcbg}) \wedge \mathbf{M}' \in \mathcal{A}.\text{MR}(\text{gcbg}) \wedge \\
\mathbf{M}' = \mathbf{M}[[T, S]][[\text{wk}_4 \mapsto \underline{w}]_{\text{reg}}][[\text{wk}_5 \mapsto \widehat{\mathbf{l}}]_{\text{reg}}] \uplus \\
\{ \mathbf{l} \mapsto (\underline{w}_0, \dots, \underline{w}_{n-1}) \}_{\text{heap}} \} \\
\mathcal{A}.\text{GR}(\text{gcbg}) \stackrel{\text{def}}{=} \{ \mathbf{M} \in \text{Mem} \mid \forall \mathbf{l} \in \text{reach}(\mathbf{M}). \widehat{\mathbf{l}} \text{ live in } \mathbf{M} \} \\
\mathcal{A}.\text{MR}(\text{gcbg}) \stackrel{\text{def}}{=} \{ \mathbf{M} \in \text{Mem} \mid (\mathbf{M}.\text{tab}, \mathbf{M}.\text{shp}) \in \mathcal{A}(\text{gcbg}).I \wedge \\
\mathbf{M}.\text{code} \supseteq [\text{gcbg} \Rightarrow \mathcal{A}(\text{gcbg}).\text{instrs}] \}
\end{array}$$

Figure 5. Abstract Specification of the Memory Allocator

mdisj uses the `undef` option for `reg`, `stk`, `tab`, and `shp` to ensure that if a memory is split into disjoint pieces, each of these indivisible components can only appear in one of the pieces.

Possible Worlds The model of possible worlds that we use to implement `WorldSpec` is based very closely on Dreyer *et al.* [14]. For space reasons, we will therefore not present the details of the model in this paper and instead refer the reader to the appendix [15]. (See Section 8 for a more detailed discussion of how our model relates to Dreyer *et al.*'s.)

That said, there are a few salient aspects of the model that are needed for understanding the definition of compositional program equivalence that we will give in Section 7. In the model, worlds W are 3-tuples (k, ω, GR) , where: k is W 's step level, ω is a finite set of state transition systems (STS's) of the sort described in Section 3, and GR is a global invariant governing the entire memory. We call the STS's *islands* because they govern disjoint pieces of memory. Two memories are related by W if (1) they satisfy its global invariant GR , and (2) they can be split into disjoint memories (one for each island) such that the n -th pair of memories satisfies the “local” memory relation determined by the current state of the n -th island in ω . In our definition of program equivalence, we will use GR to enforce the property that all reachable data is live. We need to use a global invariant since reachability cannot be determined by looking at a local subheap. We use islands to express all other assumptions about memory.

6. Assumptions About the Memory Allocator

Figure 5 shows the assumptions we make about memory allocation and garbage collection, in the form of the specification `AllocSpec`.

Given as input a starting physical address `gcbg`, the runtime system represented by \mathcal{A} will return a 4-tuple $(\text{init}, \text{alloc}, \text{instrs}, I)$, where: `init` is the starting address of the initialization routine that sets up the runtime system, `alloc` is the starting address of the allocator, `instrs` is the list of instructions defining the runtime system, which are assumed to be loaded at address `gcbg`, and I is a private invariant of the runtime system, which describes when a logical memory's lookup table is in sync with its system heap. The assumption about `instrs` is joined together with the private invariant I to form the memory predicate MR defined at the bottom of the figure.

Assuming `init` is invoked with the runtime system code in the right place, and with a return address placed in `wk4`, its specification says it will return control in a memory that satisfies $MR(\text{gcbg})$, along with the global invariant $GR(\text{gcbg})$ that all reachable data is live.

Assuming that `alloc` is invoked in a physical memory Φ represented abstractly by the logical memory \mathbf{M} , that \mathbf{M} satisfies the MR and GR properties, that the number of cells to be allocated (n) is stored in `wk5`, and that the return address is stored in `wk4`, the specification of `alloc` says that it will return a pointer to a fresh n -word block in `wk5`, and that the memory it returns (\mathbf{M}') will continue to satisfy all the aforementioned invariants. Moreover, while the lookup table and system heap of \mathbf{M}' may be completely different from those of \mathbf{M} , the contents of \mathbf{M} must remain otherwise unchanged. This does not of course prevent the allocator from having performed a GC: any logical pointer that was not reachable in \mathbf{M} before the call to `alloc` may very well be marked as dead in the lookup table of the post-allocation \mathbf{M}' , but any pointer that was reachable in \mathbf{M} will still be reachable in \mathbf{M}' and thus, by the definition of GR , still be live.

Our specification of the runtime system provides considerable flexibility—for example, it should be satisfied by either a mark-and-sweep or a copying collector because the specification says nothing about the private invariant of the runtime system. However, it *does* assume that the collector places no restrictions (such as read or write barriers) on what the mutator does to live data. We believe it should be possible to adapt our approach to a wider range of collectors, but we leave that to future work.

7. Compositional Program Equivalence

The logical relation $\mathcal{E}[[\tau]]\rho$ characterizes what it means for two *computations* to be logically equivalent, but ultimately what we really care about is whether a pair of **HIGH** and **LOW** *programs* are logically equivalent. What, one may wonder, is the difference between computations and programs? In short, a program is what you *write*, and a computation is what you *run*. That is, a program is a piece of relocatable code that must be linked with other programs and loaded into memory before it can be executed, whereas a computation describes the “next” thing to be executed in a running machine configuration. For the **HIGH** language, the distinction between computations and programs can be easily glossed over because the operational semantics of **HIGH** is defined directly on **HIGH** programs. For the **LOW** language, however, especially given the ability to write self-modifying code, it is important to distinguish the two notions. In this section, we explain what a **LOW** program is and how to define logical equivalence between **HIGH** and **LOW** programs, and we then present our key technical results.

7.1 Equivalence of HIGH and LOW Programs

As can already be seen from our motivating example in Section 3.1, we define a **LOW** program p to be a function from two code pointers to a list of instructions. The first of the code pointer inputs is assumed to be the address of the memory allocation routine, and the second is assumed to be the address where the list of

$\mathcal{H}.\text{Prog} \stackrel{\text{def}}{=} \{ e \mid \text{floc}(e) = \emptyset \}$ $\mathcal{L}.\text{Prog} \stackrel{\text{def}}{=} \{ p \in \text{PAddr} \times \text{PAddr} \rightarrow \text{List Instruction} \}$ $\mathcal{D}[\cdot] \stackrel{\text{def}}{=} \emptyset$ $\mathcal{D}[\Delta, \alpha] \stackrel{\text{def}}{=} \{ (\rho, \alpha \mapsto R) \mid \rho \in \mathcal{D}[\Delta] \wedge R \in \text{TyValRel} \}$ $\mathcal{G}[\cdot] \rho \stackrel{\text{def}}{=} \{ (W, \mathbf{v}, \emptyset) \mid W \in \text{World} \wedge \mathbf{v} \in \mathcal{L}.\text{Val} \}$ $\mathcal{G}[\Gamma, x : \tau] \rho \stackrel{\text{def}}{=} \{ (W, \mathbf{v}, (\gamma, x \mapsto v)) \mid \exists \mathbf{v}_1, \mathbf{v}_2. \\ (W, \mathbf{v}, \langle \rangle) \in \square(\mathcal{L}.\text{pair}(\mathbf{v}_1, \mathbf{v}_2), \mathcal{H}.\text{Val} \times \mathcal{H}.\text{Mem}) \wedge \\ (W, \mathbf{v}_1, v) \in \mathcal{V}[\tau] \rho \wedge (W, \mathbf{v}_2, \gamma) \in \mathcal{G}[\Gamma] \rho \}$ $W_k^\circ(\mathcal{A}, \text{gcbg}) \stackrel{\text{def}}{=} (k, [\iota^{\text{regstk}}, \iota^{\text{htyping}}, \iota^{\text{gc}}(\mathcal{A}, \text{gcbg})], GR^\circ(\mathcal{A}, \text{gcbg}))$ $\Delta; \Gamma \vdash \text{bg} \approx_W e : \tau \stackrel{\text{def}}{=} \\ \forall W' \sqsupseteq W. \forall \rho \in \mathcal{D}[\Delta]. \forall (\mathbf{v}, \gamma) \in \mathcal{G}[\Gamma] \rho(W'). \\ ((\text{bg}, [\text{wk}_0], [\text{wk}_5], \{ \mathbf{M} \mid \mathbf{M}.\text{reg}(\text{sv}_0) = \mathbf{v} \}), \gamma \rho e) \in \mathcal{E}[\tau] \rho(W')$ <p>where $\gamma \rho e ::= e[\rho(\alpha).\tau_2/\alpha][\gamma/x]$</p> $\Delta; \Gamma \vdash p \approx e : \tau \stackrel{\text{def}}{=} \\ \emptyset; \Delta; \Gamma \vdash e : \tau \wedge \\ \forall \mathcal{A}, \text{gcbg}, \text{bg}. \forall k, W \sqsupseteq W_k^\circ(\mathcal{A}, \text{gcbg}). \forall (\mathbf{M}, M) \in \mathcal{M}(W). \\ \forall \mathbf{M}'. \mathbf{M}' = \mathbf{M} \uplus \{ [\text{bg} \mapsto p(\mathcal{A}(\text{gcbg}).\text{alloc}, \text{bg})] \}_{\text{code}} \implies \\ \exists W' \sqsupseteq W. \text{lev}(W') = \text{lev}(W) \wedge (\mathbf{M}', M) \in \mathcal{M}(W') \wedge \\ \Delta; \Gamma \vdash \text{bg} \approx_{W'} e : \tau$
--

Figure 6. Program Equivalence

instructions returned by the program will be loaded into memory. When listing the code of a program (e.g., in Figure 2), we write line numbers on selected lines of code (e.g., $\text{bg} + 3$) to indicate the physical addresses where we expect the code to be loaded, but note that (1) these addresses are always relative to the second parameter of the program (typically named bg), so that the code is always relocatable, and (2) the notation is merely suggestive—the line numbers are not part of the actual program.

Now, concerning equivalence of **HIGH** and **LOW** programs: it is possible to define a notion of logical equivalence strictly between *closed* programs, but we will find it useful when reasoning about compiler correctness to generalize this relation to one on *open* programs. On the **HIGH** side, an open program is simply an expression e with free variables (and no free locations), but what is an open program on the **LOW** side? In order to answer this, we need to pick an environment-passing convention, specifying what is an official low-level representation of a high-level environment and where the low-level program expects to get its environment data. In Figure 6, we give a logical relation $\mathcal{G}[\Gamma] \rho$ between low-level values and high-level environments. This relation specifies that a high-level environment γ should be represented as a linked list of low-level values that are componentwise related to the high-level values in the range of γ . We assume that the environment data is passed in the register sv_0 —this assumption will be codified in the definition of the program equivalence relation \approx (see below).

Before we define program equivalence, we must also specify the invariants on memories that are required for executing programs. These conditions are represented by the *initial* worlds W_k° in Figure 6. (The k in W_k° simply determines the step level but does not otherwise affect the definition.) The initial worlds consist of three islands and one global invariant: ι^{regstk} owns the register file and the stack in the **LOW** memory and requires that callee-save registers and stack be preserved before and after function calls (this is accomplished using a combination of private and public transitions as discussed in Section 3.5); ι^{htyping} requires that the heap typing in the **HIGH** memory should only grow in future worlds; ι^{gc} owns the lookup table and the system heap and enforces the private $\mathcal{A}.\text{MR}(\text{gcbg})$ invariant of the runtime system (Figure 5); and the global invariant $GR^\circ(\mathcal{A}, \text{gcbg})$ specifies that all reachable blocks in the **LOW** memory are live ($\mathcal{A}.\text{GR}(\text{gcbg})$ in

Figure 5) and that the **HIGH** memory satisfies its heap typing. All these components of W_k° are formally defined in the appendix [15].

We are now ready to define the program equivalence relation \approx . The judgment $\Delta; \Gamma \vdash p \approx e : \tau$ says that a **LOW** program p and a **HIGH** program e are equivalent if the following is true. First, let k be any starting step level, and let W be any future world of W_k° that does not already impose any invariants on the code segment of p . (This latter condition is guaranteed by the assumption that we are given initial memories \mathbf{M} and M related by W , but where \mathbf{M} does not contain the code segment of p . We use the notation $\mathbf{M}_1 \uplus \mathbf{M}_2$ here to denote the “smallest” memory in $\text{mdisj}(\mathbf{M}_1, \mathbf{M}_2)$, in a sense defined formally in the appendix.)

Under these assumptions, we must be able to:

1. Extend W to a future world W' (of the same step level) such that W' relates \mathbf{M}' and M , where \mathbf{M}' is \mathbf{M} extended with the code of p . Intuitively, this step affords p the opportunity to “own” its own code segment by extending W with an island governing it. Typically, this island will take the form of an invariant stating that p ’s code must never be modified, although in the case of self-modifying code we would instead define the island to be a state transition system (as described in Section 3.4).
2. Show that the **HIGH** computation e is related (under the world W' constructed in the previous step) to the **LOW** computation e that starts at the beginning (bg) of p ’s code segment. This subgoal is encapsulated in the *open computation equivalence* judgment $\Delta; \Gamma \vdash \text{bg} \approx_{W'} e : \tau$ (Figure 6). This judgment says that, for any future world W'' of W' , for any relational interpretation ρ of the type variables in Δ , and for any environments \mathbf{v} and γ related under W'' by $\mathcal{G}[\Gamma] \rho$, the **HIGH** computation $\gamma \rho e$ is logically related (by $\mathcal{E}[\tau] \rho$, under W'') to the **LOW** computation starting at address bg . The other three components of the **LOW** computation stipulate, respectively, that (1) the computation expects to find its return address stored in wk_0 , (2) the computation will store its resulting value in wk_5 , and (3) the computation expects to find its environment stored in sv_0 . One can think of this last assumption about sv_0 as having the effect of “closing” p , in much the same way that γ and ρ close e .

7.2 Compiler Correctness and Other Technical Results

Our first result is an *adequacy* theorem for our program equivalence relation \approx . The statement of the theorem refers to the following simple loader for the **LOW** language: $\text{load}(\mathcal{A}, p)$ first runs the initialization routine of the memory allocator \mathcal{A} , then executes p , and finally halts as soon as it gets control back from p :

```

load( $\mathcal{A}, p$ ) ::= let (init, alloc, gcinstrs, _) :=  $\mathcal{A}(105)$ ,
               instrs $p$  :=  $p(\text{alloc}, 105 + |\text{gcinstrs}|)$ ,
               loadinstrs $p$  := [
(* 100 *)   move   [wk4]           102
              jmp    init
(* 102 *)   move   [wk0]           104
              jmp    105 + |gcinstrs|
(* 104 *)   halt
] ++ gcinstrs ++ instrs $p$  in
{ ( $\Phi, 100$ )  $\in$  PConf  $\mid \Phi.\text{code} \sqsupseteq [100 \mapsto \text{loadinstrs}^p]$  }

```

The adequacy theorem states that closed **HIGH** and **LOW** programs that are equivalent according to \approx must equi-terminate when loaded by the above loader.

Theorem 1 (Adequacy). For all $\emptyset; \emptyset \vdash p \approx e : \tau$,
 $\forall \mathcal{A} \in \text{AllocSpec}. \forall (\Phi, \text{pc}) \in \text{load}(\mathcal{A}, p). \forall h.$
both (Φ, pc) and (h, e) diverge or both halt without fail.

One might think that this adequacy result is weak because the loader does not inspect the result returned by the program p . How-

Papp(p_1, p_2) ::= λ alloc, bg.				
let instrs ₁ := p_1 (alloc, bg + 4), c_1 := instrs ₁ ,				
instrs ₂ := p_2 (alloc, bg + c_1 + 6), c_2 := instrs ₂ in [
bg	plus	[sp]	[sp]	<u>2</u>
	move	\langle sp - 2 \rangle_s	[wk ₀]	
	move	\langle sp - 1 \rangle_s	0	
	move	[wk ₀]	<u>bg + c_1 + 4</u>	
	instrs ₁			
bg + c_1 + 4	move	\langle sp - 1 \rangle_s	[wk ₅]	
	move	[wk ₀]	<u>bg + c_1 + c_2 + 6</u>	
	instrs ₂			
bg + c_1 + c_2 + 6	move	[wk ₀]	\langle sp - 2 \rangle_s	
	move	[wk ₁]	\langle sp - 1 \rangle_s	
	move	[wk ₂]	[wk ₅]	
	minus	[sp]	[sp]	<u>2</u>
	jump	\langle wk ₁ + 0 \rangle_h		
]				

Figure 7. Compilation of Function Application

ever, together with the compositionality result below, one can link the program p with arbitrary well-behaved “test” programs and the linked programs are guaranteed to behave the same.

In order to show that our logical relations are sufficiently populated—an important desideratum, as we explained in the introduction—we have written a very naïve compiler from **HIGH** to **LOW** and proved that every source program is related to its compiled low-level program by our program equivalence. Specifically, we have implemented a low-level construct corresponding to each high-level construct and then defined the compiler purely inductively on the structure of source programs using these low-level constructs. For each construct, we have shown a corresponding “compatibility” lemma (following Pitts’ terminology [21]), meaning that program equivalence is preserved under said construct. This implies that equivalent programs behave the same under arbitrary well-behaved contexts.

One example of a high-level construct is function application. Figure 7 shows its low-level realization as a simple linking program $\text{Papp}(p_1, p_2)$. We assume here that p_1 is some program computing a value of function type $\tau' \rightarrow \tau$, and p_2 is some program computing a value of the argument type τ' .

The program begins by bumping up the stack pointer twice, pushing the return address (stored in wk_0) onto the first new stack slot, and clearing the second one with 0. The clearing instruction (at $\text{bg} + 2$) is needed because the stack slot at $\langle \text{sp} - 1 \rangle_s$ might otherwise contain a dangling pointer; thus, in order to maintain the global invariant that all reachable data is live, we must clear $\langle \text{sp} - 1 \rangle_s$ before passing control to p_1 . When p_1 returns, we assume it returns a pointer to a function closure in wk_5 , we store that pointer in the cleared stack slot, and we proceed to execute p_2 . After p_2 returns, we move p_1 ’s closure pointer into wk_1 , the argument value (returned by p_2) into wk_2 , and the original return address of $\text{Papp}(p_1, p_2)$ into wk_0 . We then pop off two stack slots and make a tail call to the code pointer stored in wk_1 ’s closure.

The compatibility result for Papp can be seen as a compositionality result for our program equivalence relation.

Lemma 1 (Compatibility: App).

$$\Delta; \Gamma \vdash p_1 \approx e_1 : \tau' \rightarrow \tau \wedge \Delta; \Gamma \vdash p_2 \approx e_2 : \tau' \implies \Delta; \Gamma \vdash \text{Papp}(p_1, p_2) \approx e_1 e_2 : \tau$$

Another related construct of course is λ -abstraction. Figure 8 shows its low-level realization as the program $\text{Pabs}(p)$. We assume here that p is a program implementing the body of a λ -abstraction, under the assumption that the argument of that abstraction is the first element in the linked list environment stored at sv_0 .

Pabs(p) ::= λ alloc, bg.				
let code := p (alloc, bg + 16), c := code in [
bg	move	[wk ₄]	<u>bg + 3</u>	
	move	[wk ₅]	<u>2</u>	
	jump	alloc		
bg + 3	move	\langle wk ₅ + 0 \rangle_h	<u>bg + 6</u>	
	move	\langle wk ₅ + 1 \rangle_h	[sv ₀]	
	jump	[wk ₀]		
	bg + 6			
	plus	[sp]	[sp]	<u>2</u>
	move	\langle sp - 2 \rangle_s	[wk ₀]	
	move	\langle sp - 1 \rangle_s	[sv ₀]	
	move	[wk ₄]	<u>bg + 12</u>	
	move	[wk ₅]	<u>2</u>	
	jump	alloc		
bg + 12	move	\langle wk ₅ + 0 \rangle_h	[wk ₂]	
	move	\langle wk ₅ + 1 \rangle_h	\langle wk ₁ + 1 \rangle_h	
	move	[sv ₀]	[wk ₅]	
	move	[wk ₀]	<u>bg + c + 16</u>	
	code			
bg + c + 16	move	[sv ₀]	\langle sp - 1 \rangle_s	
	minus	[sp]	[sp]	<u>2</u>
	jump	\langle sp - 0 \rangle_s		
]				

Figure 8. Compilation of λ -Abstraction

The actual computation of the program is quite simple, because it merely allocates a closure representing the λ -abstraction and returns it (just like the example in Section 3.1). The closure consists of a code pointer (to $\text{bg} + 6$) and an environment pointer, which is of course just whatever environment was passed to $\text{Pabs}(p)$ in sv_0 . Whenever the closure is invoked (by jumping to $\text{bg} + 6$), it first saves the return address (stored in wk_0), as well as the callee-save register sv_0 , by pushing them on the stack. It then pushes the argument value $[\text{wk}_2]$ onto the front of the environment linked-list headed by $\langle \text{wk}_1 + 1 \rangle_h$ (which requires allocating two new memory cells), and stores a pointer to this extended environment in sv_0 before executing p . Finally, when p returns, it pops the stack twice, and restores the original contents of the callee-save sv_0 , before jumping to the return address that was stored on the stack.

The compatibility result for Pabs is as follows:

Lemma 2 (Compatibility: Abs).

$$\Delta; \Gamma, x : \tau' \vdash p \approx e : \tau \implies \Delta; \Gamma \vdash \text{Pabs}(p) \approx \lambda x : \tau'. e : \tau' \rightarrow \tau$$

Using Papp , Pabs , and all the other **LOW** program constructors, we can define a compiler $(\Gamma \vdash e)$ in a simple, syntax-directed fashion, e.g.,

$$\begin{aligned} (\Gamma \vdash e_1 e_2) &::= \text{Papp}((\Gamma \vdash e_1), (\Gamma \vdash e_2)) \\ (\Gamma \vdash \lambda x : \tau. e) &::= \text{Pabs}((\Gamma, x : \tau \vdash e)) \end{aligned}$$

and then establish the following compiler correctness result:

Theorem 2 (Compiler Correctness). For $\emptyset; \Delta; \Gamma \vdash e : \tau$,

$$\Delta; \Gamma \vdash (\Gamma \vdash e) \approx e : \tau$$

The theorem is easily provable by induction on e , using the appropriate compatibility lemma in each case.

Finally, we prove the self-modifying awkward program p_{awk} is equivalent to the high-level awkward program e_{awk} from Figure 2.

Theorem 3. $\emptyset; \emptyset \vdash p_{\text{awk}} \approx e_{\text{awk}} : (\text{unit} \rightarrow \text{unit}) \rightarrow \text{int}$

As a corollary, we can see that for any $\emptyset; \emptyset; \emptyset \vdash e : \text{unit} \rightarrow \text{unit}$, both $e_{\text{awk}} e$ and $\text{load}(\mathcal{A}, \text{Papp}(p_{\text{awk}}, (\emptyset \vdash e)))$ equi-terminate. Similarly, for any $\emptyset; \emptyset; \emptyset \vdash e : ((\text{unit} \rightarrow \text{unit}) \rightarrow \text{int}) \rightarrow \tau$, both $e e_{\text{awk}}$ and $\text{load}(\mathcal{A}, \text{Papp}((\emptyset \vdash e), p_{\text{awk}}))$ equi-terminate.

Detailed proofs of all these results appear in the companion technical appendix [15].

8. Related and Future Work

There is a huge body of work on compiler correctness and semantics for low-level code. We focus on the most closely related work.

Compositional Compiler Correctness As explained in the introduction, the overall motivation of our work is very similar to that of Benton and Hur [5], and our use of logical relations to build an extensional, compositional notion of equivalence between high- and low-level languages is inspired directly by their work. However, there are significant differences between our work and theirs.

First, they define a relation between a purely functional PCF-like language and an SECD machine, whereas we relate a more expressive, impure, ML-like high-level language to an assembly language that is significantly more low-level and realistic than SECD. Reasoning about compositional equivalence in our setting is significantly more complex, not least because we must deal with reasoning about the heap and the presence of a garbage collector. We make essential use of Kripke logical relations for this purpose.

That said, there is a sense in which our setting makes the problem easier. One of Benton and Hur’s goals was to develop a model of low-level programs that would admit program equivalences (such as commutativity of addition) whose validity depend on the purely functional nature of the source language. Toward this end, they related low-level programs to *denotations* of high-level programs, so that one could use domain-theoretic reasoning to establish the purely functional equivalences of interest. These half-operational, half-denotational relations were of necessity asymmetric. In particular, they employ biorthogonality—but only on the low-level side of the relation—as well as step-indexing—but only in defining one direction of approximation (in the other direction, they use an admissible closure operation). As a result, proving in their setting that a high- and low-level program are *equivalent* really involves doing two proofs (one for each direction of approximation) using very different technical machinery.

In our work, we sidestep this problem because the ML-like lack of effect encapsulation in our high-level language causes it to have a relatively weak equational theory that simply does not admit the kinds of purely functional equivalences that Benton and Hur were interested in. Nevertheless, as our motivating example illustrates, there are still plenty of interesting equivalences in our setting, particularly involving uses of local state. Moreover, our logical relations, being entirely operational and defined in language-generic fashion, are inherently symmetric, making them easier to use.

More recently, Benton and Hur [6] have generalized their technique to a compiler for a polymorphic (yet still purely functional) language, but their logical relations are still asymmetric.

Chlipala [11] proposes a syntactic approach to proving compositional compiler correctness. His idea is to establish a set of criteria for high-low compilation relations, such that the results of different compilers can be correctly linked so long as their compilation relations satisfy these criteria. However, the criteria are still syntactic, and thus he cannot reason for instance about our motivating example, wherein the high-low equivalence depends on semantic reasoning about local state. Moreover, Chlipala only considers a fairly high-level target language that is a CPS version of the source.

Jaber and Tabareau [16] propose an alternative approach to compositional compiler correctness based on type preservation. Instead of proving compiler correctness directly, they prove that the compiler is type-preserving, but their source language has such a rich type system (with dependent refinement types) that this effectively implies correctness. Like Benton and Hur, they compile a

purely functional language to an SECD machine, but their correctness result only applies to terminating programs.

None of the aforementioned work considers a target language that supports garbage collection.

Kripke Logical Relations Our logical relation is based closely on Dreyer, Neis and Birkedal’s (hereafter, DNB) [14], which was in turn a refinement and generalization of Ahmed, Dreyer and Rossberg’s [2]. (Of course, these are but the latest in a long line of work on Kripke logical relations spanning decades—see *loc. cit.*, as well as Pitts and Stark [22], for further pointers to the literature.)

DNB’s main goal was to show how a Kripke model based on state transition systems could be extended in orthogonal ways to exploit the *absence* of certain features, namely higher-order state and/or control effects. For the **HIGH** language considered here, in which there are no control operators, DNB showed how to extend their baseline model with *private transitions* (Section 3.5) and demonstrated the utility of private transitions in reasoning about a variety of challenging contextual equivalences involving local state. For our purposes, private transitions have proved useful in formalizing the “well-bracketing” assumptions about the stack and callee-save registers, assumptions which indeed rely on the absence of control operators. Extending the **HIGH** language with control operators would thus necessitate a significant change to our Kripke model, precisely because the compilation strategy for **HIGH** would need to change as well. We leave this problem to future work.

Despite the close connection, our logical relation diverges from DNB’s in several ways. First and foremost, whereas DNB’s relation was only designed to reason about high-level programs, the whole point of our model is to allow us to relate high- and low-level programs. As a result, we have no “fundamental theorem of logical relations” because one cannot even state such a theorem for a relation between two languages. Instead, we prove a *compiler correctness* result, whose proof mirrors that of the usual fundamental theorem. Our consideration of low-level programs has also led us to make a clear distinction between *programs* and *computations*, and between the notions of equivalence thereon. This seems to us an interesting and important distinction that is worth investigating further.

Dealing with low-level programs introduces significant technical complexity. In order to isolate this complexity, we factor the presentation of our relation generically w.r.t. a language specification “interface” *LangSpec*. This helps to clarify the structure of our Kripke logical relation, bringing into relief its essentially symmetric high-level structure, as well as the components of the model that are language-dependent (namely the two implementations of *LangSpec*, where most of the complexity lies). Although we have in this paper only instantiated our model so as to relate **HIGH** and **LOW** programs, one may also instantiate the model to relate **HIGH** and **HIGH** programs, or **LOW** and **LOW** programs. The former model would be largely similar to Dreyer *et al.*’s; the latter would enable one to reason about equivalence of low-level programs directly, which may prove useful in reasoning about correctness of low-level optimizations, although this remains to be explored.

As far as possible worlds are concerned, although ours are largely similar to DNB’s, we have extended their worlds in one relatively straightforward way: in addition to *local* invariants (expressed in our possible worlds by *islands*), our worlds also permit one to express a *global* invariant. We exploit this added functionality to encode the allocator’s invariant that all reachable data are live (Sections 3.6 and 6).

Lastly, as far as the high-level structure of the logical relation is concerned, ours is also quite similar to DNB’s, except that our interpretation of reference types is somewhat different. Various approaches to interpreting reference types have been proposed in the literature; our present interpretation is in a more “extensional” style than either DNB’s or Ahmed *et al.*’s [2], in the sense that

it avoids dependence on too many details of how possible worlds are structured. This enables us to present it, as we have, in a more “world-generic” fashion. Among existing accounts, our present formulation is fairly close to the denotational one given by Birkedal, Støvring and Thamsborg [7], but the jury is still out on which of these formulations is most felicitous.

Although the primary benefit of presenting our logical relation language- and world-generically is to clarify its (admittedly complex) structure, it also enables us to prove a few “structural” lemmas generically as well, most notably *monotonicity*. (That we can prove monotonicity generically should not be surprising, given that it is essentially baked into the logical relation via the quantification over future worlds in certain cases and the use of the \square operator in others.) Most of the interesting theorems, however, cannot be stated, let alone proven, without talking about the details of the LangSpec’s for **HIGH** and **LOW**.

Garbage Collection Torp-Smith *et al.* [24] prove the correctness of a Cheney copying collector using separation logic. They specify the behavior of a correct garbage collector in terms of an isomorphism between the reachable portions of the initial and final heaps. Our specification is broadly similar in that we express the reachable portion of the heap in our logical memories. We construct our entire LangSpec for **LOW** around these logical memories.

Building on the work of Torp-Smith *et al.*, McCreight *et al.* [20] develop a garbage collector interface that is general enough to characterize a variety of different collectors, including incremental copying collectors with read and write barriers. They prove in Coq that various collectors implement this interface, and that various mutator programs respect it. In more recent work, McCreight *et al.* [19] extend Leroy’s Compcert compiler [18] with support for garbage collection, by building the mutator-collector interface into the design of a new intermediate language, GCminor. They prove semantics preservation (mostly) for a compiler from a purely functional, typed language—Dminor—to GCminor, but (as in Compcert) not compositional correctness.

Self-Modifying Code Cai *et al.* [9] provide one of the only formal accounts of how to verify self-modifying code. They use Hoare-style separation-based reasoning to enable local reasoning about modifications to code, in much the same way that standard separation logic enables local reasoning about the heap. We adopt a similar approach, using possible worlds to impose local invariants on—and, more generally, to establish local state transition systems governing—both the heap and the code segment. Ours is the first relational model for reasoning about (compositional) equivalence of self-modifying programs.

Future Work Given our ability to reason about self-modifying code, one direction for future work is to adapt this functionality to reason about more practical applications, such as a just-in-time compiler or a dynamic linker/loader.

We have shown compositionality of our high-low relation by showing that it is closed under the linking constructs expressible at the level of our **HIGH** language, such as function application. In future work, we would like to prove that our relation is also compositional w.r.t. a more realistic linking language.

The concrete logical relation we have presented here assumes a uniform data representation. It is possible in principle to define the meaning of language forms like $\text{pair}(v_1, v_2)$ in a non-uniform way—*e.g.*, to enable flattening—but it is not currently possible in our language-generic framework to define such language forms in a *type*-specialized manner. We leave a serious examination of this issue to future work.

Lastly, it is unclear how to scale our techniques to reason about compositional correctness of a *multi-phase* compiler because the step-indexed logical relations we define are not obviously transi-

tive. This is a well-known problem with step-indexed logical relations [1], and it seems a fresh idea is needed to circumvent it.

Acknowledgments

We would like to thank Jacob Thamsborg for pointing out a very subtle technical oversight in an earlier draft of this paper.

References

- [1] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *ESOP*, 2006.
- [2] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, 2009.
- [3] A. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *TOPLAS*, 23(5):657–683, 2001.
- [4] A. Appel, P.-A. Melliès, C. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *POPL*, 2007.
- [5] N. Benton and C.-K. Hur. Biorthogonality, step-indexing and compiler correctness. In *ICFP*, 2009.
- [6] N. Benton and C.-K. Hur. Realizability and compositional compiler correctness for a polymorphic language. Technical Report MSR-TR-2010-62, Microsoft Research, Apr. 2010.
- [7] L. Birkedal, K. Støvring, and J. Thamsborg. A relational realizability model for higher-order stateful ADTs. Submitted for publication, 2010.
- [8] N. Bohr. *Advances in Reasoning Principles for Contextual Equivalence and Termination*. PhD thesis, IT University of Copenhagen, 2007.
- [9] H. Cai, Z. Shao, and A. Vaynberg. Certified self-modifying code. In *PLDI*, 2007.
- [10] A. Chlipala. A certified type-preserving compiler from lambda calculus to assembly language. In *PLDI*, 2007.
- [11] A. Chlipala. Syntactic proofs of compositional compiler correctness, 2009. Submitted for publication.
- [12] A. Chlipala. A verified compiler for an impure functional language. In *POPL*, 2010.
- [13] Z. Dargaye. *Vérification formelle d’un compilateur pour langages fonctionnels*. PhD thesis, Université Paris 7 Denis Diderot, July 2009.
- [14] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. In *ICFP*, 2010.
- [15] C.-K. Hur and D. Dreyer. Technical appendix for this paper, 2010. URL: <http://www.mpi-sws.org/~dreyer/papers/lrmlasm/>.
- [16] G. Jaber and N. Tabareau. Krivine realizability for compiler correctness. In *LOLA*, 2010.
- [17] J.-L. Krivine. Classical logic, storage operators and second-order lambda-calculus. *Annals of Pure and Applied Logic*, 68:53–78, 1994.
- [18] X. Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [19] A. McCreight, T. Chevalier, and A. Tolmach. A certified framework for compiling and executing garbage-collected languages. In *ICFP*, 2010.
- [20] A. McCreight, Z. Shao, C. Lin, and L. Li. A general framework for certifying garbage collectors and their mutators. In *PLDI*, 2007.
- [21] A. Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7. MIT Press, 2005.
- [22] A. Pitts and I. Stark. Operational reasoning for functions with local state. In *HOOTS*, 1998.
- [23] E. Sumii. A complete characterization of observational equivalence in polymorphic lambda-calculus with general references. In *CSL*, 2009.
- [24] N. Torp-Smith, L. Birkedal, and J. C. Reynolds. Local reasoning about a copying garbage collector. *TOPLAS*, 30(4), 2008.