

Realizability and Compositional Compiler Correctness for a Polymorphic Language

Nick Benton
Microsoft Research
Cambridge

Chung-Kil Hur
PPS, Université Paris Diderot
Paris

April 2010

Technical Report
MSR-TR-2010-62

We construct operationally-based realizability relations between phrases in a language with both universal and existential types and programs for a variant SECD machine. The relations, defined using parametricity, biorthogonality and step-indexing, give extensional and compositional specifications of when low-level code and values realize typed source-level terms. We prove full functional correctness of a compiler in terms of these relations and show how they also justify both source-level transformations and the linking of compiled code with hand-optimized code fragments that exploit non-parametric and non-functional low-level operations whilst being extensionally well-behaved. The definitions and results have been fully formalized in Coq.

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
<http://www.research.microsoft.com>

1 Introduction

As the title suggests, this paper will describe a (mechanized) proof that a particular compiler is correct. The real subject, however, is how to define a good specification of when a low-level code fragment should be said to ‘correspond to’ a phrase in a high-level language.

A straightforward compiler correctness theorem says that for every closed, ground type source program P , the result $C(P)$ of running compiler C on P is a target program whose observable behaviour (termination, final result, IO behaviour) ‘matches’ that of C . Proving such a theorem involves a strengthened induction hypothesis, relating *open* source phrases of higher types to target code (and values). This richer relation, typically some kind of (bi)simulation, is often essentially the simplest extension of the function C itself that suffices to establish the ‘big’ theorem about complete programs, which are, after all, the only ones we can run according to the semantics of our source language. But the ‘closed’ systems we really run are not the result of a single compilation: they are composed by linking code from many places, including libraries, the operating system, the runtime system and foreign functions, which may be compiled with different compilers and written in many languages, including ‘cleverly’ handcrafted machine code. To reason modularly about all these components, we need a clean specification of the interface between compiled code and its environment, a job for which a naive induction hypothesis is inappropriate. The kind of specification we want should constrain only the observable behaviour of code, rather than intensional details of just how it executes, and make no reference to details of a particular compiler beyond those aspects of data representation and calling conventions that *have* to be agreed upon for interoperability. Similarly, the correctness relation should not be tweaked to admit individual source-level optimizations, but should rather be closed under a rich set of high-level equations (possibly even source contextual equivalence) by construction.

In a previous paper, we addressed the question of when low-level code correctly realizes a source term by defining relations between the denotational semantics of a simply-typed source language and the operational behaviour of a lower-level virtual machine [3]. Here we take another step forward, treating a source language with impredicative universal and existential types and defining relations that express how parametricity and data abstraction principles from the typed source translate to an untyped target. This is a non-trivial technical extension, and, as one of the examples will demonstrate, our realizability relation captures the requirement on low-level code realizing quantified types to behave parametrically from an extensional perspective, whilst allowing sufficient freedom for it to implement that behaviour in a decidedly non-parametric manner. We compositionally prove full functional correctness for a compiler for the polymorphic language, which also performs tail-call optimizations. Another major difference from our previous work is that we work with an operational, rather than a denotational, semantics for the high-level language. Rather than fix a high-level notion of equivalence that should be respected, we parameterize

Values:	$\frac{\Theta; \Gamma, f : \tau \rightarrow \tau', x : \tau \vdash M : \tau'}{\Theta; \Gamma \vdash \text{rec } f(x : \tau) : \tau' = M : \tau \rightarrow \tau'} \quad \frac{\Theta \vdash \Gamma \quad \Theta, X \vdash \tau \quad \Theta, X; \Gamma \vdash V : \tau}{\Theta; \Gamma \vdash \Lambda X.V : \forall X.\tau}$ $\frac{\Theta \vdash \Gamma \quad \Theta, X \vdash \tau \quad \Theta \vdash \tau' \quad \Theta; \Gamma \vdash V : \tau[\tau'/X]}{\Theta; \Gamma \vdash \text{pack } \tau', V \text{ to } \exists X.\tau : \exists X.\tau}$
Expressions:	$\frac{\Theta; \Gamma \vdash V : \tau \quad \Theta; \Gamma \vdash M : \tau \quad \Theta; \Gamma, x : \tau \vdash N : \tau'}{\Theta; \Gamma \vdash [V] : \tau \quad \Theta; \Gamma \vdash \text{let } x = M \text{ in } N : \tau'}$ $\frac{\Theta \vdash \Gamma \quad \Theta, X \vdash \tau \quad \Theta; \Gamma \vdash V : \forall X.\tau \quad \Theta \vdash \tau'}{\Theta; \Gamma \vdash V \tau' : \tau[\tau'/X]}$ $\frac{\Theta \vdash \Gamma \quad \Theta \vdash \tau \quad \Theta, X \vdash \tau' \quad \Theta; \Gamma \vdash V : \exists X.\tau' \quad \Theta, X; \Gamma, x : \tau' \vdash M : \tau}{\Theta; \Gamma \vdash \text{unpack } V \text{ as } X, x \text{ in } M : \tau}$
Transition semantics:	$\text{let } x = [V] \text{ in } N \mapsto N[V/x] \quad (\Lambda X.V)\tau \mapsto V[\tau/X]$ $(\text{rec } f(x : \tau) : \tau' = M) V \mapsto M[(\text{rec } f(x : \tau) : \tau' = M)/f, V/x]$ $\text{unpack } (\text{pack } \tau, V \text{ to } \exists X.\tau') \text{ as } X, x \text{ in } M \mapsto M[\tau/X][V/x]$

Figure 1: Selected typing and transition rules for F_v

our definitions and results by a novel form of adequate precongruence relation on the source, incorporating an abstract notion of chain to capture analogues of domain-theoretic admissibility.

All the metatheory and examples have been formally verified in the Coq proof assistant, and the proof script is available from the authors' web pages.

2 High- and Low-Level Languages

High-Level Language. F_v is a conventional call-by-value functional language with recursion and impredicative universal and existential types. The types are:

$$\tau ::= X \mid \text{Int} \mid \tau \rightarrow \tau' \mid 1 \mid \tau \times \tau' \mid \tau + \tau' \mid \forall X.\tau \mid \exists X.\tau$$

We separate values, V , from expressions, M , and restrict the syntax to ANF, with explicit sequencing by **let** and inclusion of values into expressions by $[\cdot]$. Selected typing rules for F_v are shown in Figure 1, where Θ and Γ are contexts for type and term variables respectively. Write $\text{Value } \Theta \Gamma \tau$ for the set of values of type τ in contexts $\Theta; \Gamma$ (where $\Theta; \Gamma \vdash \tau$), $\text{CValue } \tau$ for the closed values of (closed) type τ , and similarly for expressions. If $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ then the set of environments $\text{EValue } \Gamma$ is $\prod_{i=1 \dots n} \text{CValue } \tau_i$.

Selected transitions from the standard CBV semantics of F_v are also shown in Figure 1. There is an equivalent big-step semantics and we write $M \mapsto$ (resp. $M \mapsto V$) when the closed expression M converges (to the closed value V).

Our realizability relations will be parameterized by a notion of observational approximation \sqsubseteq on F_v . To cope abstractly with recursion, we take as basic $\hat{\sqsubseteq}_{\Theta \Gamma \tau} \subseteq (\text{Value } \Theta \Gamma \tau) \times (\text{Value } \Theta \Gamma \tau)^\omega$, a slightly unusual (type- & context-indexed, but we usually omit indices) relation between values and *sequences* of values, and derive the associated order on values via constant sequences:

$V \sqsubseteq V'$ iff $V \hat{\sqsubseteq} (\lambda i \in \omega. V')$. There are homonymous orders on expressions and environments. A sequence $\langle V_i \rangle_i$ is a *chain* if $V_i \sqsubseteq V_j$ for all $i \leq j$. One should think of $V \hat{\sqsubseteq} \langle W_i \rangle_i$ meaning $V \sqsubseteq \sqcup_i W_i$ in a domain-theoretic sense, but without requiring lubs to exist.

The conditions on $\hat{\sqsubseteq}$ and \sqsubseteq , eliding types and with the same conditions applied, *mutatis mutandis*, to the order on expressions, are: *[Chain]*: if $V \hat{\sqsubseteq} \langle W_i \rangle_i$ then $\langle W_i \rangle_i$ is a chain; *[Elem]*: if $\langle W_i \rangle_i$ is a chain, $W_j \hat{\sqsubseteq} \langle W_i \rangle_i$ for all j ; *[Refl]*: $V \sqsubseteq V$ for all V of the appropriate type; *[Trans]*: if $U \hat{\sqsubseteq} \langle V_i \rangle_i$ and $V_j \hat{\sqsubseteq} \langle W_i \rangle_i$ for all j , then $U \hat{\sqsubseteq} \langle W_i \rangle_i$; *[Subst]*: if $V \hat{\sqsubseteq} \langle W_i \rangle_i$ then $V[U/x] \hat{\sqsubseteq} \langle W_i[U/x] \rangle_i$ and similarly for type substitutions; *[Compat]*: all constructs of F_v preserve $\hat{\sqsubseteq}$, e.g.

$$\frac{\Theta; \Gamma \vdash V \hat{\sqsubseteq} \langle W_i \rangle_i : \tau[\tau'/X]}{\Theta; \Gamma \vdash \text{pack } \tau', V \text{ to } \exists X. \tau \hat{\sqsubseteq} \langle \text{pack } \tau', W_i \text{ to } \exists X. \tau \rangle_i : \exists X. \tau}$$

[Beta]: $\text{let } x = [V] \text{ in } N \sqsubseteq N[V/x]$ and vice versa; *[Adeq]*: If $M \hat{\sqsubseteq} \langle N_i \rangle_i : \text{Int}$ and $M \Vdash n$ for some n , then there exists j such that $N_j \Vdash n$, and similarly for the unit type; *[Unfold]*: $\text{rec } f x. M \hat{\sqsubseteq} \langle \text{recn}_i f x. M \rangle_i$ where $\text{recn}_0 f x. M = \text{rec } f x. f x$ and $\text{recn}_{i+1} f x. M = \lambda x. M[\text{recn}_i f x. M/f]$. These conditions imply that \sqsubseteq is an adequate precongruence satisfying an ‘unwinding theorem’ [12]. An important example is given by defining $V \hat{\sqsubseteq} \langle W_i \rangle_i$ iff $\forall C[\cdot], C[V] \Vdash \implies \exists j, \forall i \geq j, C[W_i] \Vdash$; another is generated by a (non fully-abstract) step-indexed logical relation like that of Ahmed [1]. The extension of $\hat{\sqsubseteq}$ to environments, which are typed lists of closed values, is pointwise.

Low-Level Machine. Our target is a variant SECD machine [9]. Although the SECD machine was designed as a target for compiling functional languages, the kind of relations we construct will work for lower-level targets too. The substantial independence of our definitions from the fine detail of exactly what compiled code looks like is part of the point of the compositional, extensional approach we are espousing, and we have added new, non-functional, operations to the original SECD machine to express more interesting and realistic low-level optimizations.

A configuration is a quadruple $\langle c, e, s, d \rangle \in \text{CESD}$, as defined in Figure 2. An *MVal* is either a natural, a closure, a recursive closure, or a pair of values. The deterministic transition relation \mapsto between configurations is defined in Figure 3. The non-standard ‘no dump’ forms of application and selection do not save a continuation on the dump and are used for tail-call optimizations. **PushE** and **PopE** allow the environment to be modified. **IsNum** tests for numberhood, and **Eq** for intensional equality. Our previous paper [3] explains how non-functional ‘reflective’ operations, such as **Eq**, in the target break a straightforward realizability interpretation of types in the presence of term-level recursion, requiring step-indexing (or similar) in defining low-level interpretations of high-level terms.

Configurations with no successor are *terminated*. Write $\text{cesd} \mapsto^k$ if cesd takes at least k steps without having terminated, and say it diverges, written $\text{cesd} \mapsto^\omega$, if it can always take a step. We say cesd terminates, and write $\text{cesd} \mapsto^* \not\downarrow$, if it does not diverge.

$CESD \stackrel{\text{def}}{=} Code \times MEnv \times Stack \times Dump$ $c \in Code \stackrel{\text{def}}{=} list\ Instruction$ $e \in MEnv \stackrel{\text{def}}{=} list\ MVal$ $s \in Stack \stackrel{\text{def}}{=} list\ MVal$ $d \in Dump \stackrel{\text{def}}{=} list\ (Code \times MEnv \times Stack)$ $Instruction \ni inst := Pop \mid Push\ i \mid PushE \mid PopE \mid PushN\ n \mid Op\ \star \mid PushC\ c$ $\mid PushRC\ c \mid App \mid AppNoDump \mid Ret \mid Sel\ (c_1, c_2) \mid SelNoDump\ (c_1, c_2) \mid Join$ $\mid MkPair \mid Fst \mid Snd \mid Eq \mid IsNum$ $n, i \in \mathbb{N} \quad \star \in \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ $MVal \ni v := \underline{n} \mid CL\ (e, c) \mid RCL\ (e, c) \mid PR\ (v_1, v_2)$

Figure 2: Extended SECD machine

Compiling F_v to SECD. The compiler is shown in Figure 4 and comprises mutually-recursive functions, both written $\langle \cdot \rangle$, mapping typed F_v values and expressions into *Code*. The compilation of expressions is parameterized by a boolean flag *ret* that identifies expressions that are in ‘tail position’ and hence expect to be immediately followed by a return instruction. Applications in tail position are compiled with the `AppNoDump` instruction, which does *not* push the calling context to the dump, so allowing the called function to return directly to the caller’s caller. Similarly, conditionals normally push a common continuation to the dump, and each branch ends with a `Join`; when the conditional is in tail position, however, the pushing of the context is elided and each branch compiled in tail position.

3 Logical Relations

We will now define two logical relations between components of the SECD machine and terms of F_v . \preceq specifies when a low-level component approximates (in the observational sense of ‘diverging in more contexts than’) a source term at a particular type, whilst \succeq expresses the converse. These two relations can be seen as corresponding to the traditional soundness and adequacy theorems used to show correspondence between an operational and a denotational semantics.

We start by giving a broad overview of the constructions used in defining the relations. Firstly, both relations are parameterized by $\hat{\sqsubseteq}$, an approximation relation on the source satisfying the conditions we gave in Section 2. \sqsubseteq can be taken to be the contextual preorder for F_v , but may be something weaker, such as the order of some non-fully abstract denotational model. The factorization separates concerns and provides some ‘tuneability’ in the degree to which the realizability relation is required to preserve source-level equivalence. Secondly, the \preceq relation, which is intuitively about specifying that low-level code should diverge in certain contexts, involves step-indexing [2] on the low-level side. Di-

$\langle \text{Pop} :: c, e, v :: s, d \rangle$	\mapsto	$\langle c, e, s, d \rangle$
$\langle \text{Push } i :: c, [v_1, \dots, v_k], s, d \rangle$	\mapsto	$\langle c, [v_1, \dots, v_k], v_i :: s, d \rangle$
$\langle \text{PushE} :: c, e, v :: s, d \rangle$	\mapsto	$\langle c, v :: e, s, d \rangle$
$\langle \text{PopE} :: c, v :: e, s, d \rangle$	\mapsto	$\langle c, e, v :: s, d \rangle$
$\langle \text{PushN } n :: c, e, s, d \rangle$	\mapsto	$\langle c, e, \underline{n} :: s, d \rangle$
$\langle \text{PushC } bod :: c, e, s, d \rangle$	\mapsto	$\langle c, e, \text{CL}(e, bod) :: s, d \rangle$
$\langle \text{PushRC } bod :: c, e, s, d \rangle$	\mapsto	$\langle c, e, \text{RCL}(e, bod) :: s, d \rangle$
$\langle \text{App} :: c, e, v :: \text{CL}(e', bod) :: s, d \rangle$	\mapsto	$\langle bod, v :: e', [], (c, e, s) :: d \rangle$
$\langle \text{App} :: c, e, v :: \text{RCL}(e', bod) :: s, d \rangle$	\mapsto	$\langle bod, v :: \text{RCL}(e', bod) :: e', [], (c, e, s) :: d \rangle$
$\langle \text{AppNoDump} :: c, e, v :: \text{CL}(e', bod) :: s, d \rangle$	\mapsto	$\langle bod, v :: e', [], d \rangle$
$\langle \text{AppNoDump} :: c, e, v :: \text{RCL}(e', bod) :: s, d \rangle$	\mapsto	$\langle bod, v :: \text{RCL}(e', bod) :: e', [], d \rangle$
$\langle \text{Op } \star :: c, e, \underline{n_2} :: \underline{n_1} :: s, d \rangle$	\mapsto	$\langle c, e, \underline{n_1} \star \underline{n_2} :: s, d \rangle$
$\langle \text{Ret} :: c, e, v :: s, (c', e', s') :: d \rangle$	\mapsto	$\langle c', e', v :: s', d \rangle$
$\langle \text{Sel}(c_1, c_2) :: c, e, v :: s, d \rangle$	\mapsto	$\langle c_1, e, s, (c, [], []) :: d \rangle$ (if $v \neq \underline{0}$)
$\langle \text{Sel}(c_1, c_2) :: c, e, \underline{0} :: s, d \rangle$	\mapsto	$\langle c_2, e, s, (c, [], []) :: d \rangle$
$\langle \text{SelNoDump}(c_1, c_2) :: c, e, v :: s, d \rangle$	\mapsto	$\langle c_1, e, s, d \rangle$ (if $v \neq \underline{0}$)
$\langle \text{SelNoDump}(c_1, c_2) :: c, e, \underline{0} :: s, d \rangle$	\mapsto	$\langle c_2, e, s, d \rangle$
$\langle \text{Join} :: c, e, s, (c', e', s') :: d \rangle$	\mapsto	$\langle c' ++ c, e, s, d \rangle$
$\langle \text{MkPair} :: c, e, v_1 :: v_2 :: s, d \rangle$	\mapsto	$\langle c, e, \text{PR}(v_2, v_1) :: s, d \rangle$
$\langle \text{Fst} :: c, e, \text{PR}(v_1, v_2) :: s, d \rangle$	\mapsto	$\langle c, e, v_1 :: s, d \rangle$
$\langle \text{Snd} :: c, e, \text{PR}(v_1, v_2) :: s, d \rangle$	\mapsto	$\langle c, e, v_2 :: s, d \rangle$
$\langle \text{Eq} :: c, e, v_1 :: v_2 :: s, d \rangle$	\mapsto	$\langle c, e, \underline{1} :: s, d \rangle$ (if $v_1 = v_2$)
$\langle \text{Eq} :: c, e, v_1 :: v_2 :: s, d \rangle$	\mapsto	$\langle c, e, \underline{0} :: s, d \rangle$ (if $v_1 \neq v_2$)
$\langle \text{IsNum} :: c, e, \underline{n} :: s, d \rangle$	\mapsto	$\langle c, e, \underline{1} :: s, d \rangle$
$\langle \text{IsNum} :: c, e, v :: s, d \rangle$	\mapsto	$\langle c, e, \underline{0} :: s, d \rangle$ (if $\nexists n, v = \underline{n}$.)

Figure 3: Operational Semantics of Extended SECD Machine

vergence arises as the limit as k increases of ‘takes at least k steps without terminating’.

A third important construction is the use of biorthogonality to ‘extensionalize’ the sets of low-level values that are related to particular F_v terms. We are trying to define compositional specifications for *components* of SECD configurations, in particular instruction sequences $c \in \text{Code}$, but we want those specifications to ultimately depend only on the observable behaviour of complete, runnable configurations. This is achieved, building on ideas of Pitts and of Krivine, by making our specifications ‘ $\top\top$ -closed’ [12, 13]. The rough idea here is that one starts with an over-intensional set of computations, constructs the set of all contexts that yield some particular observation when linked with any element of the initial set, and then constructs the set – larger than that with which one started – of those computations that yield the observation when combined with any of those contexts. In the case of \preceq , the observation will be divergence (actually, stepping for at least some number of steps), whilst for \succeq the observation will be termination.

Values:	$\langle \Theta; \vec{x}_j : \vec{\tau}_j \vdash x_i : \tau_i \rangle$	=	[Push i]
	$\langle () \rangle$	=	[PushN 0]
	$\langle n \rangle$	=	[PushN n]
	$\langle (V_1, V_2) \rangle$	=	$\langle V_1 \rangle ++ \langle V_2 \rangle ++ [\text{MkPair}]$
	$\langle \text{inl } V \rangle$	=	[PushN 1] ++ $\langle V \rangle ++ [\text{MkPair}]$
	$\langle \text{inr } V \rangle$	=	[PushN 0] ++ $\langle V \rangle ++ [\text{MkPair}]$
	$\langle \lambda x. M \rangle$	=	[PushC ($\langle M \rangle_{\text{true}}$)]
	$\langle \text{rec } f x = M \rangle$	=	[PushRC ($\langle M \rangle_{\text{true}}$)]
	$\langle \Lambda X. V \rangle$	=	$\langle V \rangle$
	$\langle \text{pack } \tau', V \text{ to } \exists X. \tau \rangle$	=	$\langle V \rangle$
Expressions:	$\langle \text{ret} \rangle$	=	if $\text{ret} = \text{true}$ then [Ret] else []
	$\langle [V] \rangle_{\text{ret}}$	=	$\langle V \rangle ++ \langle \text{ret} \rangle$
	$\langle V_1 \star V_2 \rangle_{\text{ret}}$	=	$\langle V_1 \rangle ++ \langle V_2 \rangle ++ [\text{Op } \star] ++ \langle \text{ret} \rangle$
	$\langle V_1 > V_2 \rangle_{\text{ret}}$	=	$\langle V_1 \rangle ++ \langle V_2 \rangle ++ [\text{Op } (\lambda(n_1, n_2). n_1 > n_2 \supset 1 \mid 0), \text{PushN 0, MkPair}] ++ \langle \text{ret} \rangle$
	$\langle \pi_1(V) \rangle_{\text{ret}}$	=	$\langle V \rangle ++ [\text{Fst}] ++ \langle \text{ret} \rangle$
	$\langle \pi_2(V) \rangle_{\text{ret}}$	=	$\langle V \rangle ++ [\text{Snd}] ++ \langle \text{ret} \rangle$
	$\langle \text{case } V \text{ of inl } x. M_1 \mid \text{inr } y. M_2 \rangle_{\text{true}}$	=	$\langle V \rangle ++ [\text{Dup, Snd, PushE, Fst, SelNoDump } (\langle M_1 \rangle_{\text{true}}, \langle M_2 \rangle_{\text{true}}), \text{PopE}]$
	$\langle \text{case } V \text{ of inl } x. M_1 \mid \text{inr } y. M_2 \rangle_{\text{false}}$	=	$\langle V \rangle ++ [\text{Dup, Snd, PushE, Fst, Sel } (\langle M_1 \rangle_{\text{false}} ++ [\text{Join}], \langle M_2 \rangle_{\text{false}} ++ [\text{Join}]), \text{PopE}]$
	$\langle \text{let } x = M \text{ in } N \rangle_{\text{ret}}$	=	$\langle M \rangle_{\text{false}} ++ [\text{PushE}] ++ \langle N \rangle_{\text{ret}} ++ [\text{PopE}]$
	$\langle V_1 V_2 : \rangle_{\text{true}}$	=	$\langle V_1 \rangle ++ \langle V_2 \rangle ++ [\text{AppNoDump}]$
	$\langle V_1 V_2 : \rangle_{\text{false}}$	=	$\langle V_1 \rangle ++ \langle V_2 \rangle ++ [\text{App}]$
	$\langle V \tau' \rangle_{\text{ret}}$	=	$\langle V \rangle ++ \langle \text{ret} \rangle$
	$\langle \text{unpack } V \text{ as } X, x \text{ in } M \rangle_{\text{ret}}$	=	$\langle V \rangle ++ [\text{PushE}] ++ \langle M \rangle_{\text{ret}} ++ [\text{PopE}]$

Figure 4: Compiler for F_v

We will clearly want to relate source terms both to constructed machine values and to instruction sequences, but it is convenient to work with pairs $(c, s) \in MComp \stackrel{\text{def}}{=} Code \times Stack$ instead of isolated bits of code. If $c \in Code$, write $\hat{c} \in MComp$ for $(c, [])$, and if $v \in MVal$, write $\hat{v} \in MComp$ for $([], [v])$. We will use concatenation to link elements of $MComp$ with contexts, which are themselves elements of $CESD$.

Approximating High-level By Low-level. The \preceq relation works with step-indexed entities. We write $iMValue$ for $\mathbb{N} \times MVal$, $iMComp$ for $\mathbb{N} \times MComp$ and $iCESD$ for $\mathbb{N} \times CESD$. The relation is ‘logical’, with each type constructor having an associated relational action. Biorthogonality is used in defining the action of the lifting monad that, whilst not reflected explicitly in F_v types, is morally there in the distinction between expressions and values. Given a relation R between F_v values and machine values, we’ll want a relation between F_v expressions and machine computations. A direct approach would say something about the expression and the computation evaluating to R -related values, but

that would overspecify just *how* the low-level code must compute, requiring a non-observable intermediate configuration of a particular shape. We instead exploit two maps, each of which is half of a contravariant Galois connection, between the lattice of subsets of $iCESD$ and those of subsets of $iMValue$ and of $iMComp$, respectively. If $pe \in MEnv$ and $P \subseteq iMValue$, define

$$\begin{aligned} \downarrow^{pe}(P) &\stackrel{\text{def}}{=} \{(j, \langle c, e, s, d \rangle) \mid \forall (i, v) \in P, \langle c, pe++e, v :: s, d \rangle \mapsto^{\min(i,j)}\} \\ &\subseteq iCESD. \end{aligned}$$

So an indexed context is in $\downarrow^{pe}(P)$ if whenever we link it with an indexed value from P and the machine environment pe , the resulting configuration takes a number of steps that is at least the minimum of the two indices. Coming back the other way, if $pe \in MEnv$ and $Q \subseteq iCESD$, define

$$\begin{aligned} \uparrow^{pe}(Q) &\stackrel{\text{def}}{=} \{(i, \langle c', s' \rangle) \mid \forall (j, \langle c, e, s, d \rangle) \in Q, \langle c'++c, pe++e, s'++s, d \rangle \mapsto^{\min(i,j)}\} \\ &\subseteq iMComp. \end{aligned}$$

The maps are indexed by an (an extension of) the environment as a way of sharing the environment between the computation and the context.

If $e \in MEnv$, τ is a closed type and $RT_i \subseteq MVal \times (CValue \tau)$ is a \mathbb{N} -indexed relation between machine values and closed F_v values of type τ , then define the indexed relation $\langle RT_{\perp}^e \rangle_i \subseteq MComp \times (CExp \tau)$ by¹

$$\begin{aligned} \langle RT_{\perp}^e \rangle_i &= \{(comp, M) \mid \\ &\quad (i, comp) \in \uparrow^e(\downarrow^e(\{(j, v) \mid \exists V : \tau, M \mapsto V \wedge (v, V) \in RT_j\}))\} \end{aligned}$$

If τ is a closed source type, then let $iRel_{\tau}$ be the set of \mathbb{N} -indexed relations $R_i \subseteq MVal \times (CValue \tau)$. We say such a relation is *decreasing* when $R_0 \supseteq R_1 \supseteq \dots$. The intuition of step-indexing is that a relation R is approximated by relations of the form ‘not detectably un-related within i -steps’, so the relations should get finer as more steps are available for testing. If $\Theta = X_1, \dots, X_n$ is a type variable environment, then a relation environment for Θ is a vector \vec{R} of pairs $\tau_1 R_1, \dots, \tau_n R_n$ where each τ_k is a closed type and $R_k \in iRel_{\tau_k}$.

Now for each Θ , matching relation environment \vec{R} and type σ such that $\Theta \vdash \sigma$, the indexed relation

$$\leq_i^{\vec{R}, \sigma} \subseteq MVal \times (CValue (\sigma[\tau_k/X_k]))$$

is defined by induction on σ , as shown in Figure 5. The relational interpretation of a type variable is looked up in the relation environment, machine integers approximate the corresponding high-level value, any machine value approximates the unit value, and machine pairs approximate F_v pairs pointwise. High-level sum values are approximated by tagged pairs on the machine. The case for functions follows the usual pattern of monadic Kripke logical relations: at all future worlds (smaller indices) take related arguments to results related by the monadic lifting of the relational interpretation of the result type. This is where the low-level interface for function types, the calling convention, is specified: a machine value approximates a high-level function when putting it on the stack

¹The lifting is a form of possibility modality, in the sense of Evaluation Logic [11], whence the notation $\langle R \rangle$.

$$\begin{aligned}
\trianglelefteq_i^{\tau\vec{R}, X_k} &= (R_k)_i \\
\trianglelefteq_i^{\tau\vec{R}, \text{Int}} &= \{(n, n) \mid n \in \mathbb{N}\} \\
\trianglelefteq_i^{\tau\vec{R}, 1} &= \{(v, ()) \mid v \in MVal\} \\
\trianglelefteq_i^{\tau\vec{R}, \sigma_1 \times \sigma_2} &= \{(\text{PR}(v_1, v_2), \langle V_1, V_2 \rangle) \mid (v_1, V_1) \in \trianglelefteq_i^{\tau\vec{R}, \sigma_1} \wedge (v_2, V_2) \in \trianglelefteq_i^{\tau\vec{R}, \sigma_2}\} \\
\trianglelefteq_i^{\tau\vec{R}, \sigma_1 + \sigma_2} &= \{(\text{PR}(n + \underline{1}, v), \text{inl } V) \mid (v, V) \in \trianglelefteq_i^{\tau\vec{R}, \sigma_1}, n \in \mathbb{N}\} \\
&\quad \cup \{(\text{PR}(\underline{0}, v), \text{inr } V) \mid (v, V) \in \trianglelefteq_i^{\tau\vec{R}, \sigma_2}\} \\
\trianglelefteq_i^{\tau\vec{R}, \sigma_1 \rightarrow \sigma_2} &= \{(f, F) \mid \forall k \leq i, \forall (v, V) \in \trianglelefteq_k^{\tau\vec{R}, \sigma_1}, \\
&\quad (([\text{App}], [v, f]), (F V)) \in \langle \langle \trianglelefteq_k^{\tau\vec{R}, \sigma_2} \rangle \rangle_k \} \\
\trianglelefteq_i^{\tau\vec{R}, \forall X. \sigma} &= \{(v, V) \mid \forall \tau', \forall R' \in iRel_{\tau'}, \\
&\quad \text{decreasing } R' \rightarrow (\hat{v}, (V \tau')) \in \langle \langle \trianglelefteq_i^{\tau\vec{R}, \tau' R', \sigma} \rangle \rangle_i \} \\
\trianglelefteq_i^{\tau\vec{R}, \exists X. \sigma} &= \{(v, \text{pack } \tau', V \text{ to } \exists X. \sigma) \mid \exists R' \in iRel_{\tau'}, \\
&\quad \text{decreasing } R' \wedge (v, V) \in \trianglelefteq_i^{\tau\vec{R}, \tau' R', \sigma} \}
\end{aligned}$$

Figure 5: The relation $\trianglelefteq_i^{\tau\vec{R}, \sigma}$

with a value approximating a high-level argument and executing an **App** instruction yields behaviour that approximates that of the high-level application. Universally quantified types are interpreted using relational parametricity, where we quantify over all decreasing indexed relations between machine values and values of some F_v type. Similarly, a machine value v is related to a high-level existential package if there is *some* decreasing relation between low-level values and values of the witnessing F_v type such that v is related to the packed value.

We lift \trianglelefteq to environments pointwise. If $\Theta \vdash \Gamma$ where Γ is $x_1 : \sigma_1, \dots, x_m : \sigma_m$, and $\tau\vec{R}$ is a relation environment for Θ , then the indexed relation $\trianglelefteq_i^{\tau\vec{R}, \Gamma} \subseteq MEnv \times (EValue(\Gamma[\tau_k/X_k]))$ is the set of pairs $([v_1, \dots, v_m], [V_1, \dots, V_m])$ such that $v_j \trianglelefteq_i^{\tau\vec{R}, \sigma_j} V_j$ for all $1 \leq j \leq m$. Then for expressions in context:

$$\begin{aligned}
\trianglelefteq_i^{\tau\vec{R}, \Gamma \vdash \sigma} &\subseteq MComp \times (Exp [](\Gamma[\tau_k/X_k])(\sigma[\tau_k/X_k])) \\
&= \{(comp, M) \mid \forall i' \leq i, \forall (e, \vec{V}) \in \trianglelefteq_{i'}^{\tau\vec{R}, \Gamma}, (comp, M[V_j/x_j]) \in \langle \langle \trianglelefteq_{i'}^{\tau\vec{R}, \sigma} \rangle \rangle_{i'} \}
\end{aligned}$$

which is again ‘logical’, taking related environments to related computations for all smaller indices. That was for closed types.

For open expressions of $\Theta; \Gamma \vdash \sigma$, we define

$$\begin{aligned}
\trianglelefteq_i^{\Theta; \Gamma \vdash \sigma} &\subseteq MComp \times Exp \Theta \Gamma \sigma \\
&= \{(comp, M) \mid \forall \tau\vec{R} : \Theta, \text{decreasing } \vec{R} \rightarrow comp \trianglelefteq_i^{\tau\vec{R}, \Gamma \vdash \sigma} (M[\tau_k/X_k])\}
\end{aligned}$$

by quantifying over all decreasing relation environments (i.e. those for which every R_k is decreasing) for Θ . Finally, we can define the true \preceq relations between

the machine and F_v by quantifying over all step indices and combining with the \sqsubseteq relation:

$$comp \preceq^{\Theta; \Gamma \vdash \sigma} M \in Exp \Theta \Gamma \sigma \stackrel{\text{def}}{\iff} \exists M', \Theta; \Gamma \vdash M' \sqsubseteq M : \sigma \wedge \forall i, comp \preceq_i^{\Theta; \Gamma \vdash \sigma} M'.$$

Approximating Low-Level By High-Level. The \succeq relation expresses when a machine computation is approximated by an F_v term. We again use biorthogonality on the machine side, but this time with respect to the observation of termination. The property of being less than some fixed high-level term is a safety property of machine computations, for which step indexing is appropriate, whereas being greater is a liveness property: we will be specifying that certain machine configurations terminate. We again start with maps between sets of contexts and sets of values and of computations, defining a notion of orthogonality. For $pe \in MEnv$, $P \subseteq MVal$, $Q \subseteq CESD$, define

$$\begin{aligned} \downarrow^{pe}(P) &= \{\langle c, e, s, d \rangle \mid \forall v \in P, \langle c, pe++e, v :: s, d \rangle \mapsto^* \downarrow\} \subseteq CESD \\ \uparrow^{pe}(Q) &= \{\langle c', s' \rangle \mid \forall \langle c, e, s, d \rangle \in Q, \langle c'++c, pe++e, s'++s, d \rangle \mapsto^* \downarrow\} \subseteq MComp \end{aligned}$$

and now, for a closed type τ and a relation $RT \subseteq MVal \times (CValue \tau)$, define a relational lifting modality by

$$\begin{aligned} [RT_{\perp}^e] &\subseteq MComp \times (CExp \tau) \\ &= \{(comp, M) \mid \forall V, M \mapsto V \rightarrow comp \in \overline{\uparrow}^e(\downarrow^e(\{v \mid v RT V\}))\} \end{aligned}$$

That is, given a low-high relation RT on values, the lifted relation holds between a low computation to a high one if whenever the high computation yields a value V , the low computation terminates in all contexts that terminate whenever they are fed values v that are related to V by RT .

For closed τ , let Rel_{τ} be $\mathbb{P}(MVal \times (CValue \tau))$. If $\Theta = X_1, \dots, X_m$, a relation environment for Θ is now a vector $\vec{\tau R}$ of pairs of closed types and relations, where $R_k \in Rel_{\tau_k}$. For each $\vec{\tau R} : \Theta$ and $\Theta \vdash \sigma$, we now define the relation

$$\succeq^{\vec{\tau R}, \sigma} \subseteq MVal \times (CValue (\sigma[\tau_k/X_k]))$$

by induction on σ , as shown in Figure 6. This relation also lifts pointwise to environments. For $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$ and $\Theta \vdash \Gamma$ and $\vec{\tau R} : \Theta$, define $\succeq^{\vec{\tau R}, \Gamma} \subseteq MEnv \times (EValue (\Gamma[\tau_k/X_k]))$ to be the set of pairs $([v_1, \dots, v_n], (V_1, \dots, V_n))$ such that $(v_j, V_j) \in \succeq^{\vec{\tau R}, \sigma_j}$ for all j . Then for computations in context,

$$\begin{aligned} \succeq^{\vec{\tau R}, \Gamma \vdash \sigma} &\subseteq MComp \times (Exp [] (\Gamma[\tau_k/X_k]) (\sigma[\tau_k/X_k])) \\ &= \{(comp, M) \mid \forall (e, \vec{V}) \in \succeq^{\vec{\tau R}, \Gamma}, (comp, M[V_j/x_j]) \in [(\succeq^{\vec{\tau R}, \sigma})_{\perp}^e]\} \end{aligned}$$

For open types, we define $\succeq^{\Theta; \Gamma \vdash \sigma}$ by universally quantifying over all relation environments $\vec{\tau R} : \Theta$, just as we did for the $\preceq_i^{\Theta; \Gamma \vdash \sigma}$. The true \succeq relations are then given by combining with $\hat{\sqsubseteq}$ as follows:

$$comp \succeq^{\Theta; \Gamma \vdash \sigma} M \in Exp \Theta \Gamma \sigma \stackrel{\text{def}}{\iff} \exists \langle N_i \rangle_i, \Theta; \Gamma \vdash M \hat{\sqsubseteq} \langle N_i \rangle_i : \sigma \wedge \forall j, comp \succeq^{\Theta; \Gamma \vdash \sigma} N_j.$$

Realizability. The conjunction of the two approximation relations is our notion of when a machine computation realizes an F_v term:

$$comp \models^{\Theta; \Gamma \vdash \sigma} M \stackrel{\text{def}}{\iff} comp \preceq^{\Theta; \Gamma \vdash \sigma} M \wedge comp \succeq^{\Theta; \Gamma \vdash \sigma} M$$

$$\begin{array}{l}
\sqsubseteq^{\tau\bar{R}, X_k} = R_k \\
\sqsubseteq^{\tau\bar{R}, \text{Int}} = \{(\underline{n}, n) \mid n \in \mathbb{N}\} \\
\sqsubseteq^{\tau\bar{R}, 1} = \{(v, ()) \mid v \in MVal\} \\
\sqsubseteq^{\tau\bar{R}, \tau_1 \times \tau_2} = \{(\text{PR}(v_1, v_2), (V_1, V_2)) \mid (v_1, V_1) \in \sqsubseteq^{\tau\bar{R}, \tau_1} \wedge (v_2, V_2) \in \sqsubseteq^{\tau\bar{R}, \tau_2}\} \\
\sqsubseteq^{\tau\bar{R}, \sigma_1 + \sigma_2} = \{(\text{PR}(\underline{n+1}, v), \text{inl } V) \mid (v, V) \in \sqsubseteq^{\tau\bar{R}, \sigma_1}, n \in \mathbb{N}\} \\
\quad \cup \{(\text{PR}(\underline{0}, v), \text{inl } V) \mid (v, V) \in \sqsubseteq^{\tau\bar{R}, \sigma_2}\} \\
\sqsubseteq^{\tau\bar{R}, \sigma_1 \rightarrow \sigma_2} = \{(f, F) \mid \forall (v, V) \in \sqsubseteq^{\tau\bar{R}, \sigma_1}, (([\text{App}], [v, f]), (F V)) \in [(\sqsubseteq^{\tau\bar{R}, \sigma_2})_{\perp}^{\square}]\} \\
\sqsubseteq^{\tau\bar{R}, \forall X.\sigma} = \{(v, V) \mid \forall \tau', \forall R' \in \text{Rel}_{\tau'}, (\hat{v}, (V \tau')) \in [(\sqsubseteq^{\tau\bar{R}, \tau' R', \sigma})_{\perp}^{\square}]\} \\
\sqsubseteq^{\tau\bar{R}, \exists X.\sigma} = \{(v, \text{pack } \tau', V \text{ to } \exists X.\sigma) \mid \exists R' \in \text{Rel}_{\tau'}, (v, V) \in \sqsubseteq^{\tau\bar{R}, \tau' R', \sigma}\}
\end{array}$$

Figure 6: The relation $\sqsubseteq^{\tau\bar{R}, \sigma}$

It is immediate from the definitions that the \models relations are closed on the right under $=$, the symmetric closure of \sqsubseteq . We note again that these definitions make surprisingly little reference to actual low-level code and values. We specify the encoding for base types, pairs and sums and otherwise the only real piece of code that shows up is the application instruction in the definition of the relations at function types. That calling convention is the interface across which linked components communicate.

If $(c', s') \in MComp$, we say (c', s') diverges unconditionally if for any c, e, s, d , $\langle c' ++ c, e, s' ++ s, d \rangle \mapsto^{\omega}$.

Lemma 3.1 (Ground divergence adequacy) *For any $comp \in MComp$, type $\tau = \text{Int}$ or 1 , and $M \in CExp \tau$ if $comp \models^{\square; \square^{\tau}} M$ and the F_v term M diverges, then $comp$ diverges unconditionally.*

Say a computation (c', s') converges to a natural n if plugging it into an arbitrary context equiterminates with plugging n into that context:

$$\begin{aligned}
\forall c, e, s, d, \langle c, e, \underline{n} :: s, d \rangle \mapsto^* \frac{1}{2} &\implies \langle c' ++ c, e, s' ++ s, d \rangle \mapsto^* \frac{1}{2} \\
\wedge \langle c, e, \underline{n} :: s, d \rangle \mapsto^{\omega} &\implies \langle c' ++ c, e, s' ++ s, d \rangle \mapsto^{\omega}.
\end{aligned}$$

If a computation realizes a closed F_v term that evaluates to an integer n , then it converges to n :

Lemma 3.2 (Ground convergence adequacy) *For any $comp \in MComp$, $M \in CExp \text{Int}$ and $n \in \mathbb{N}$, if $comp \models^{\square; \square^{\text{Int}}} M$ and $M \mapsto n$ then $comp$ converges to n .*

Convergence adequacy also holds for observation at the unit type, with a definition of convergence that quantifies over test contexts $cesd$ whose termination is independent of the value on the top of the stack (since any value realizes $()$).

What allows the realizability semantics to be used for modular reasoning about linking code obtained from different places, and proved by different means, is that it is compositional. An important case is that of function application:

Lemma 3.3 (Compositionality for application) *For any $cf, cx \in \text{Code}$ and $V_f \in \text{Value} \Theta \Gamma (\tau \rightarrow \tau')$, $V_x \in \text{Value} \Theta \Gamma \tau$,*

$$(cf, []) \models^{\Theta; \Gamma \vdash \tau \rightarrow \tau'} [V_f] \wedge (cx, []) \models^{\Theta; \Gamma \vdash \tau} [V_x] \implies (cf ++ cx ++ [\mathbf{App}], []) \models^{\Theta; \Gamma \vdash \tau'} V_f V_x.$$

4 Examples

Compiler Correctness. Whilst the realizability relation is defined without reference to our compiler, we do of course intend that the compiler is correct, in the sense that compiled code always realizes the original source term:

Theorem 4.1 (Compositional Compiler Correctness) *1. For all Θ, Γ, V, τ , if $\Theta; \Gamma \vdash V : \tau$ then $\langle V \rangle \models^{\Theta; \Gamma \vdash \tau} [V]$.*

2. For all Θ, Γ, M, τ , if $\Theta; \Gamma \vdash M : \tau$ then $\langle M \rangle_{\text{false}} \models^{\Theta; \Gamma \vdash \tau} M$.

Both directions of the above are proved by simultaneous structural induction, with a strengthened induction hypothesis to account for compilation of expressions with the *ret* flag set to true. In the \preceq direction, the case for recursive functions involves a nested induction over step indices, whilst in the \succeq direction, we appeal to the unwinding property of the \sqsubseteq parameter relation.

A consequence is that compiled code for whole programs has the correct operational behaviour according to the operational semantics of the programs:

Corollary 4.2 (Correctness for whole programs) *For any $M \in \text{CExp Int}$, if M diverges, then $\langle M \rangle_{\text{false}}$ diverges unconditionally and if $M \mapsto n$, then $\langle M \rangle_{\text{false}}$ converges to n .*

The corollary is normally thought of as compiler correctness, but it is Theorem 4.1 that lets us reason about combining compiled code with code from elsewhere.

Hand-written fixed-point combinator. One can use the **rec** construct to write a CBV fixed-point combinator in F_v :

$$\text{FixC} = \Lambda X. \Lambda Y. \lambda F : (X \rightarrow Y) \rightarrow (X \rightarrow Y). \mathbf{rec} \ f \ x. F \ f \ x$$

which compiles to code using SECD's recursive closures:

`[PushC [PushRC [Push 2, Push 1, App, PushE, Push 0, Push 1, AppNoDump, PopE], Ret]]`

Alternatively, we can hand-encode $\lambda F. \lambda x. (\lambda y. F(\lambda z. y y z)) (\lambda y. F(\lambda z. y y z)) x$, which is an *untyped* CBV fixpoint combinator, as the following SECD code:

```
YCombinator =
[PushC [PushC [
  PushC [Push 2, PushC [Push 1, Push 1, App, Push 0, App, Ret], App, Ret],
  PushC [Push 2, PushC [Push 1, Push 1, App, Push 0, App, Ret], App, Ret],
  App, Push 0, App, Ret], Ret]]
```

Direct reasoning about the operational semantics of low-level code and of F_v , *not* involving anything to do with the compiler, shows the following:

Lemma 4.3 YCombinator , $\square \models^{\vdash \forall X. \forall Y. ((X \rightarrow Y) \rightarrow X \rightarrow Y) \rightarrow X \rightarrow Y} \text{FixC}$.

Together with Theorem 4.1 and the compositionality of the realizability relation, the above implies that linking the handcrafted code with code produced by the compiler for any term with an appropriately typed free variable will be observationally indistinguishable from linking in the compiled code for any source-level term that is \equiv to the explicitly recursive FixC .

Context manipulation. For this machine, most interesting examples of the difference between intensional and extensional specifications involve higher-order functions and cunning use of the intensional equality test. But one simple first-order example is that for any $n \in \mathbb{N}$,

$$([\text{PushN } n, \text{PushN } 1, \text{Sel } ([\text{Join}, \text{PushN } 0, \text{Pop}], []), []]) \models^{\vdash \text{Int}} [n]$$

This code would not be in a simple direct-style realizability relation because as well as pushing n onto the stack, it modifies its context (non-observably), appending push and pop instructions to the end of the current continuation.

Hand-optimized polymorphic list module.. We now give an example that involves relational parametricity, for both universal (polymorphic) and existential (abstract) types. Consider a signature for a polymorphic list module:

$\text{SigPolList} = \forall X. \exists LX. LX \times (X \times LX \rightarrow LX) \times (LX \rightarrow \text{Option}(X \times LX))$
 where $\text{Option } \tau = 1 + \tau$ and we write *none* and *some* for the constructors as usual. The signature says that an implementation of polymorphic lists should have some private representing type LX , equipped with three standard list manipulation operations: constructors $\text{nil} : LX$, $\text{cons} : X \times LX \rightarrow LX$ and a destructor $\text{split} : LX \rightarrow \text{Option}(X \times LX)$.

Now one concrete implementation of the signature is given by the well-known Church-encoding of lists:

$$\text{PList} : \text{SigPolList} = \Lambda X. \text{packList } X, (\text{nil}_X, \text{cons}_X, \text{split}_X) \text{ to } \dots$$

where $\text{List } \tau = \forall Y. Y \times (\tau \times Y \rightarrow Y) \rightarrow Y$ for any type τ and a fresh type variable Y , and where

$$\begin{aligned} \text{nil}_\tau : \text{List } \tau &= \Lambda Y. \lambda(n, c). n \\ \text{cons}_\tau (hd : \tau) (tl : \text{List } \tau) : \text{List } \tau &= \Lambda Y. \lambda(n, c). c(hd, tl Y (n, c)) \\ \text{split}_\tau (l : \text{List } \tau) : \text{Option}(\tau \times \text{List } \tau) &= \\ & l(\text{Option}(\tau \times \text{List } \tau)) (\text{none}_{\tau \times \text{List } \tau}, \\ & \lambda(hd, htl). \text{case } htl \text{ of inl } (). \text{some}_{\tau \times \text{List } \tau} (hd, \text{nil}_\tau) \\ & \quad | \text{inr } (hd', tl). \text{some}_{\tau \times \text{List } \tau} (hd, \text{cons}_\tau hd' tl)) \end{aligned}$$

The Church encoding is elegant, but inefficient: the list splitting operation is $\mathcal{O}(n)$, rather than $\mathcal{O}(1)$, for example. But without recursive types, we can do no better in the F_v source language. However, one can treat the inefficient encoding of lists as a specification, and write some cunning, hand-optimized SECD machine code that provably realizes, i.e. behaves exactly the same as from the point of view of well-behaved clients, PList . There are two tricks

in our implementation. First, we represent lists as nested tuples of elements, which would not be typeable in F_v . We then further optimize by playing a highly non-parametric low-level representation trick: when the representation of list elements is (dynamically!) observed to be a natural number, we further compress the representation via a (potentially iterated) Gödel numbering of sequences of natural numbers. The optimized implementation of the list module in SECD code is parameterized by pairing and projection functions on natural numbers:

$$\text{npair} : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}, \quad \text{nfst} : \mathbb{N} \rightarrow \mathbb{N}, \quad \text{nsnd} : \mathbb{N} \rightarrow \mathbb{N}$$

such that,

$$\forall n, m \in \mathbb{N}, \text{npair } n \ m > 0, \text{nfst } (\text{npair } n \ m) = n, \text{and } \text{nsnd } (\text{npair } n \ m) = m.$$

One could take $\text{npair } n \ m$ to be $2^n \times 3^m$, with matching projection functions, for example. We assume low-level implementations $\text{NPair}, \text{NSplit} \in \text{Code}$ of these pairing operations, such that

$$\begin{aligned} \forall c \ e \ s \ d \ n \ m \ \exists k \ \langle \text{NPair} ++ c, e, \underline{n} :: \underline{s}, d \rangle &\mapsto^k \langle c, e, \text{npair } n \ m :: s, d \rangle \\ \forall c \ e \ s \ d \ n \ \exists k \ \langle \text{NSplit} ++ c, e, \underline{n} :: s, d \rangle &\mapsto^k \langle c, e, \text{PR } (\text{nfst } \underline{n}, \text{nsnd } \underline{n}) :: s, d \rangle \end{aligned}$$

Exploiting the instruction `IsNum` that checks if a given machine value is a number, one can compactly represent a list of machine values as follows:

$$\text{encodeM } (vs : \text{list } MVal) : MVal = \begin{cases} \underline{0} & \text{if } vs = [] \\ \text{npair } n \ m & \text{if } vs = \underline{n} :: tl \wedge \text{encodeM } tl = \underline{m} \\ \text{PR } (hd, tl) & \text{otherwise } (vs = hd :: tl) \end{cases}$$

The following are implementations of the three functions of `SigPolList` in SECD according to the above representation.

$$\begin{aligned} \text{NilM} &= [\text{PushN } 0] \\ \text{ConsM} &= [\text{PushC } [\text{Push } 0, \text{Fst}, \text{IsNum}, \\ &\quad \text{Sel } ([\text{Push } 0, \text{Snd}, \text{IsNum}, \text{Join}], [\text{Push } 0, \text{Snd}, \text{PushN } 0, \text{Join}]), \\ &\quad \text{Sel } (\text{NPair} ++ [\text{Join}], [\text{MkPair}, \text{Join}]), \\ &\quad \text{Ret}]] \\ \text{SplitM} &= [\text{PushC } [\text{Push } 0, \\ &\quad \text{Sel } ([\text{PushN } 0, \text{Push } 0, \text{IsNum}, \text{Sel } (\text{NSplit} ++ [\text{Join}], [\text{Join}]), \\ &\quad \quad \text{MkPair}, \text{Join}], \\ &\quad [\text{PushN } 1, \text{PushN } 1, \text{MkPair}, \text{Join}]), \\ &\quad \text{Ret}]]. \end{aligned}$$

To aid understanding, we give pseudo- F_v terms interpreting the above code:

$$\begin{aligned} \text{NilM} &\approx \underline{0} \\ \text{ConsM} &\approx \lambda(hd, tl). \text{if } hd = \underline{n} \wedge tl = \underline{m} \text{ then } \text{npair } n \ m \text{ else } \text{PR } (hd, tl) \\ \text{SplitM} &\approx \lambda l. \text{if } l \neq \underline{0} \text{ then if } l = \underline{n} \text{ then } \text{some } (\text{nfst } \underline{n}, \text{nsnd } \underline{n}) \text{ else } \text{some } l \\ &\quad \text{else } \text{none} \end{aligned}$$

One can then verify that the above optimized low-level implementation realizes, and is therefore substitutable for (in any related context) the code compiled from, the Church-encoded F_v module `PList`:

Lemma 4.4 $(\text{NilM} ++ \text{ConsM} ++ \text{SplitM} ++ [\text{MkPair}, \text{MkPair}], []) \models^{\text{SigPolList}} \text{PList}.$

Whilst frivolous, this shows a non-trivial extensionally parametric specification being realized by an implementation that is intensionally unequivocally non-parametric.

5 Discussion

We have presented a compositional and extensional operational realizability relation between a parametrically polymorphic source language and a low-level untyped abstract machine. The relation was used to establish both full functional correctness for an optimizing compiler and to justify the linking of code from any compiler meeting the specification with hand-optimized fragments of low-level code. The relation involves relational parametricity, biorthogonality and step-indexing, as well as being parameterized by a novel form of adequate precongruence, amounting to a kind of ideal closure operation, on the source language.

The results and examples are all formalized in the Coq proof assistant, using a strongly-typed representation of polymorphic terms and substitutions that we describe in a companion paper [4]. Using the strongly-typed representation involves some occasionally-vexing manipulation of casts and heterogeneous (‘John Major’) equality; this led the second author to develop a Coq library for rewriting with heterogeneous equality [8], which is exploited in the latest version of the proofs for the present paper. The full formalization of the source language, machine, logical relations, compiler correctness proof and examples is around 6000 lines, which seems very reasonable, though this is now at least our third completely fresh version of a mechanized compiler correctness theorem. Our initial formalization for this language was over twice as long; the improvement is partly in the formulation of the logical relations and partly in the details of the mechanization. The reasons for mechanizing at all are twofold. Firstly, the reasoning is already sufficiently complex that we really would have low confidence in (and would find it hard to manage) paper proofs, especially for non-trivial examples such as the polymorphic list module. Secondly, mechanization would be an integral part of any realistic infrastructure based on certified code, so we simultaneously want to establish the feasibility of (and extend (or encourage others to extend) the state of the art in) doing mechanized proofs in this area.

The closest related work is our own paper on relating the denotational semantics of a simply-typed language to a similar machine [3]. There have been many other compiler correctness proofs done in the last 35 years or so, amongst which we particularly mention the classic work on the VLISP verified Scheme compiler [7], which relates a denotational semantics to, ultimately, real machine code; the Coq formalization of compiler correctness for a total functional language by Chlipala [6]; and the work of Leroy [10] on mechanically verifying a realistic compiler for a C-like language. A distinguishing feature of our work is the focus on compositional specifications that are independent of any particular compiler and can be used to independently verify foreign code. We have also looked at low-level semantic type soundness in a similar style [5] and the

present work arose naturally from an attempt to understand the way in which that semantics failed to be sufficiently abstract.

Step-indexing has been widely used to tame recursive phenomena in operational semantics in the last decade. During the course of this research, however, we have found that it does inhibit some aggressive low-level program transformations that we believe *should* be legal. Intuitively, non-functional operations such as comparing function pointers allow one to look ‘too far’ into the future, so clever optimizations make misbehaving computations misbehave ‘too soon’ to respect current forms of step-indexed relations. Finding a way around this problem will be interesting and challenging. Other avenues for further work include looking at recursive types and references, and transferring our results to a lower level target machine. We would also like to make our specifications even more independent of the source language, by expressing them in a logic that talks only about the low-level machine.

References

- [1] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *15th European Symposium on Programming (ESOP)*, volume 3924 of *Lecture Notes in Computer Science*, 2006.
- [2] A. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 23(5), 2001.
- [3] N. Benton and C.-K. Hur. Biorthogonality, step-indexing and compiler correctness. In *ACM International Conference on Functional Programming*, 2009.
- [4] N. Benton, C.-K. Hur, A. Kennedy, and C. McBride. Strongly typed term representations in Coq. Submitted, November 2009.
- [5] N. Benton and N. Tabareau. Compiling functional types to relational specifications for low level imperative code. In *4th ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI)*, 2009.
- [6] A. Chlipala. A certified type-preserving compiler from lambda calculus to assembly language. In *ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation (PLDI)*, 2007.
- [7] J. Guttman, J. Ramsdell, and M. Wand. VLISP: A verified implementation of scheme. *Lisp and Symbolic Computation*, 8(1/2), 1995.
- [8] C. K. Hur. Heq: A Coq library for heterogeneous equality. <http://www.pps.jussieu.fr/~gil/Heq/>, May 2010.
- [9] P. Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4), 1964.

- [10] X. Leroy. Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In *33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2006.
- [11] A. M. Pitts. Evaluation logic. In *IVth Higher Order Workshop, Banff 1990*, Workshops in Computing, 1991.
- [12] A. M. Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7. 2005.
- [13] J. Vouillon and P.-A. Melliès. Semantic types: A fresh look at the ideal model for types. In *31st ACM Symposium on Principles of Programming Languages (POPL)*, 2004.