

Verifying an Open Compiler Using Multi-language Semantics

James T. Perconti and Amal Ahmed

Northeastern University

Abstract. Existing verified compilers are proved correct under a closed-world assumption, i.e., that the compiler will only be used to compile *whole* programs. We present a new methodology for verifying correct compilation of program *components*, while formally allowing linking with target code of arbitrary provenance. To demonstrate our methodology, we present a two-pass type-preserving open compiler and prove that compilation preserves semantics. The central novelty of our approach is that we define a combined language that embeds the source, intermediate, and target languages and formalizes a semantics of interoperability between them, using boundaries in the style of Matthews and Findler. Compiler correctness is stated as contextual equivalence in the combined language.

Note to reader: We use **blue**, **red**, and **purple** to typeset terms in various languages. This paper will be difficult to follow unless read/printed in color.

1 Introduction

There has been remarkable progress on formally verified compilers over the last few years, with researchers proving the correctness of increasingly sophisticated compilers for increasingly realistic languages. The most well known instance of this is the CompCert compiler [1,2] which uses the Coq proof assistant to both implement and verify a multi-pass optimizing compiler from C to PowerPC, ARM, and x86 assembly, proving that the compiler preserves semantics of source programs. Several other compiler-verification efforts have successfully followed CompCert's lead and basic methodology, for instance, focusing on multithreaded Java [3], just-in-time compilation [4], and C with relaxed memory concurrency [5].

Unfortunately, these projects prove compiler correctness under a *closed-world* assumption, that is, assuming that the verified compiler will always compile *whole* programs. Despite the immense effort put into verification, the compiler correctness theorem provides no guarantees about correct compilation of *components*. This whole-program assumption is completely unrealistic since most software systems today are comprised of many components written in different languages compiled to a common target, as well as runtime-library routines that may be handwritten in the target language. We need compiler correctness theorems applicable to the way we actually use these compilers.

Formally verifying that components are compiled correctly—often referred to as *compositional compiler correctness*—is a challenging problem. A key difficulty is that, in the setting of compiling components, it is not clear how to even state the compiler correctness theorem. CompCert's compiler correctness theorem is easy to state thanks to

the whole program assumption: informally, it says that if a source program P_S compiles to a target program P_T , then running P_S and P_T results in the same trace of observable events. The same sort of theorem does not make sense when we compile a component e_S to a component e_T : we cannot “run” a component since it is not a complete program.

Intuitively, we want the compiler correctness theorem to say that if a component e_S compiles to e_T , then some desired relationship $e_S \simeq e_T$ holds between e_S and e_T . The central question is: how do we *formally specify* $e_S \simeq e_T$? To answer this question, we must consider how the compiled component is actually used: it needs to be linked with some e'_T , creating a whole program that can be run. Informally, the compiler correctness theorem should guarantee that if we link e_T with e'_T , then the resulting target-level program should correspond to the source component e_S linked with e'_T . But, formally speaking, how can one link a source component with a target component and what are the rules for running the resulting source-target hybrid? These questions demand a *semantics of interoperability* between the source and target languages. We give our semantics of interoperability as a multi-language operational model. We then define $e_S \simeq e_T$ as a contextual equivalence in that model.

There are two other important issues to consider when evaluating a compositional compiler correctness theorem and its supporting formalism. The first is the degree of *horizontal compositionality* that the model allows, that is, which target components e'_T may formally be linked with a compiled component. At the lower end of the horizontal compositionality spectrum are *fully abstract* compilers. Full abstraction states that the compiler both preserves and reflects contextual equivalence. Hence, a fully abstract compiler preserves all of the source language’s abstractions, and compiled components are only allowed to link with components that can be expressed in the source language.

But real systems often link together components from multiple languages with different guarantees and different expressive power. We are particularly interested in supporting interoperability between parametric typed languages like ML and low-level languages like C. Thus, full abstraction is often too restrictive. To support the whole programs that we actually run, the compiler correctness theorem should formally support linking with as large a class of programs as possible, and in particular, should not require an e'_T to have been compiled from the same source language as e_T .

Abandoning full abstraction in favor of greater horizontal compositionality does not require giving up all the guarantees of the source language. The compiler and its verification framework can be designed to preserve the source-level equivalences that are critically needed without forbidding all foreign behavior. To show that different levels of abstraction preservation are possible, we will deliberately pick a target language that is more expressive than the source and design our compiler so that it is *not* fully abstract. Our focus in this paper is on how to preserve the representation independence and information hiding guarantees provided by type abstraction in our source language.

The second important issue for a compiler correctness framework is that we want to be able to verify multi-pass compilers. For example, if we have a two-pass compiler that compiles a source component e_S to an intermediate-language component e_I to a target component e_T , we should be able to verify each pass separately, showing $e_S \simeq e_I$ and $e_I \simeq e_T$, and then compose these results to get a correctness theorem for the whole compiler saying $e_S \simeq e_T$. This is typically referred to as *vertical compositionality*.

We will show that our approach of using a multi-language operational model succeeds at both horizontal and vertical compositionality. In particular, we validate our methodology by applying it to a two-pass type-preserving compiler. The compiler deals with three languages: our source language F (System F with existential and recursive types), an intermediate language C (the target of a typed closure conversion pass), and our target language A (the target of a heap allocation pass).¹ The target language A allows tuples and closures to live only on the heap and supports both mutable and immutable references. Our closure conversion pass translates F components of type τ to C components of type τ^C , where τ^C denotes the type translation of τ . The subsequent allocation pass translates C components of type τ to A components of type τ^A , where τ^A is the type translation of τ .

To define the semantics of interoperability between these languages, we embed them all into one language, FCA, and add syntactic boundary forms between each pair of adjacent languages, in the style of Matthews and Findler [7] and of Ahmed and Blume [8]. For instance, the term $\mathcal{CF}^\tau(e_F)$ allows an F component e_F of type τ to be used as a C component of type τ^C , while ${}^\tau\mathcal{FC}(e_C)$ allows a C component e_C of translation type τ^C to be used as an F component of type τ . Similarly, we have boundary forms \mathcal{AC} and \mathcal{CA} for the next language pair. Non-adjacent languages can interact by stacking up boundaries: for example, $\mathcal{FC}(\mathcal{CA} e_A)$ (abbreviated $\mathcal{FCA}(e_A)$) allows an A component e_A to be embedded in an F term.

FCA Design Principles. Our goal is for the FCA interoperability semantics to give us a useful specification of when a component in one of the underlying languages should be considered equivalent to a component in another language. We realize that goal by following three principles.

First, we define the operational semantics of FCA so that the original languages are *embedded* into FCA unchanged: running an FCA program that's written solely in one of the embedded languages is identical to running it in that language alone. For instance, execution of the A program e_A proceeds in exactly the same way whether we use the operational semantics of A or the augmented semantics for FCA.

Next, we ensure that the typing rules are similarly embedded: a component that contains syntax from only one underlying language should typecheck under that language's individual type system if and only if it typechecks under FCA's type system.

The final property we need is *boundary cancellation*, which says that wrapping two opposite language boundaries around a component yields the same behavior as the underlying component with no boundaries. For example, any $e_F : \tau$ must be contextually equivalent to ${}^\tau\mathcal{FC}(\mathcal{CF}^\tau e_F)$, and any $e_C : \tau^C$ must be equivalent to $\mathcal{CF}^\tau({}^\tau\mathcal{FC} e_C)$.

Compiler Correctness. We state the correctness criterion for our compiler as a contextual equivalence. For each pass of the compiler from a source S to a target T , where S and T interoperate via boundaries \mathcal{ST} and ${}^\tau\mathcal{S}$, define our source-target relationship by

$$e_S \simeq e_T \stackrel{\text{def}}{=} e_S \simeq_{\text{FCA}}^{ctr} {}^\tau\mathcal{ST}(e_T) : \tau.$$

We prove that if $e_S : \tau$ compiles to e_T , then $e_S \simeq e_T$. Since contextual equivalence is transitive, our framework achieves vertical compositionality immediately: it is easy

¹ We have extended our F to A compiler with a code-generation pass to an assembly language, much like Morrisett *et al.*'s stack-based TAL [6]. We will report on that work in a future paper.

to combine the two correctness proofs for the individual compiler passes, giving the overall correctness result that if e_F compiles to e_A , then $e_F \simeq e_A$, or

$$e_F \approx_{FCA}^{ctx} \tau \mathcal{FCA}(e_A) : \tau.$$

Reasoning About Linking. Our approach enjoys a strong horizontal compositionality property: we can link with any target component e'_A that has an appropriate type, with no requirement that e'_A was produced by any particular means or from any particular source language. Specifically, if e_F expects to be linked with a component of type τ' and compiles to e_A , then e_A will expect to be linked with a component of type $((\tau')^c)^A$. If e'_A has this type, then using our compiler correctness theorem, we can conclude that

$$(e_F \tau' \mathcal{FCA}(e'_A)) \approx^{ctx} \mathcal{FCA}(e_A e'_A),$$

or equivalently,

$$\mathcal{ACF}(e_F \tau' \mathcal{FCA}(e'_A)) \approx^{ctx} e_A e'_A.$$

The right-hand side of this equality is exactly the A program we ultimately want to run, and the left-hand side is an FCA program that models that program.

Contributions. Our main contributions are our methodology and that we have proven correctness for an open multi-pass compiler. We have designed a multi-language semantics that lets us state a strong compiler-correctness theorem, and to prove the theorems, we have developed a logical relation for proving contextual equivalences between FCA components. The most significant technical challenges were related to interoperability between languages with type abstraction, specifically, in designing the multi-language semantics so it preserves type abstraction between languages (§5), and in designing the parts of the logical relation that model the handling of type abstraction in a multi-language setting (§9).

Due to space constraints, we elide various technical details and omit proofs. All definitions, lemmas, and proofs are spelled out in full detail in the accompanying technical report [9], available at: <http://ccs.neu.edu/home/amal/voc/>

2 Related Work: Benton-Hur Approach

Before beginning our technical development, we compare our methodology to the only prominent existing approach to compositional compiler correctness.

To eliminate the closed-world assumption, Benton and Hur [10] advocate setting up a logical relation between the source and target languages, specifying when a source term semantically approximates target code and vice versa. We will refer to a logical relation that relates terms from two *different* languages as a *cross-language* logical relation. The relation is defined by induction on source-language types. Benton and Hur verified a compiler from the simply-typed λ -calculus with recursion [10]—and later, from System F with recursion [11]—to an SECD machine, proving that if source component e_S compiles to target code e_T , then e_S and e_T are logically related. Later, Hur and Dreyer [12] used essentially the same approach to prove correctness of a compiler from an idealized ML to assembly.

However, the Benton-Hur (henceforth, BH) approach suffers from serious drawbacks in both vertical and horizontal compositionality. First, the cross-language framework does not scale to a multi-pass compiler. Both Benton-Hur and Hur-Dreyer handle only

a single pass. To achieve vertical compositionality in the BH style, one would have to define separate cross-language logical relations relating the source and target of each compiler pass, and then prove that the logical relations compose transitively in order to establish that the correctness of each pass implies correctness for the entire compiler. But this kind of transitive composition of cross-language logical relations has been an open problem for some time. (We'll discuss recent work towards addressing this problem in §11.)

The second drawback to the BH approach is its limited horizontal compositionality. Consider the situation where a verified compiler from language S to language T is used to compile a source component e_S to some target code e_T . The BH compiler correctness theorem tells us that e_S and e_T are logically related. We wish to link the compiled code e_T with some other target code e'_T and verify the resulting program. To do this using the BH framework, we must now *come up with a source-level component* e'_S and show that it is logically related to e'_T . This is an onerous requirement: while it may be reasonable to come up with e'_S when the given e'_T is very simple, it seems almost impossible when e'_T consists of hundreds of lines of assembly! Further, if e'_T is compiled from some other source language R , it may not even be possible to write down an e'_S in language S that is related to e'_T .

Technically speaking, the BH approach does support linking with any target code that can be proved logically related to a source component. But it cannot support linking with any components that are not expressible in the source language. And we contend that even for the theoretically-allowed cases, in practice the approach is limited to allowing linking between only very simple components or components that were all compiled from the same source language.

Overcoming BH Limitations. By reasoning about components in the FCA setting, we can overcome both limitations of the BH framework. We have already pointed out that our framework admits vertical compositionality thanks to the transitivity of contextual equivalence.

For the second limitation of the BH approach, consider a target component e'_A . While the BH approach would need to find a related source component to fit e'_A into their framework, we only need to find an FCA component that looks like a source component. Specifically, we can use e'_A itself in a source context by wrapping it in appropriate boundaries: $\mathcal{FCA}(e'_A)$.

3 The Languages

We begin our technical development with a few notes on typesetting and notational conventions. We typeset the terms, types, and contexts of our various languages as follows:

- F (System \underline{F}) in a **blue sans-serif font**;
- C (\underline{C} losure conversion) in a **red bold font with serifs**;
- A (\underline{A} llocation) in a **purple sans-serif bold font**.

For each of our languages, we will use the metavariable e for *components* and t for *terms*. In the first two languages, F and C, terms and components coincide, but the distinction will be meaningful in language A. Similarly, all languages use τ for types, v for values, E for evaluation contexts, and C for general contexts. We write $\text{fv}(e)$

$\tau ::= \alpha \mid \text{unit} \mid \text{int} \mid \forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau \mid \langle \bar{\tau} \rangle \mid \exists \alpha. \tau \mid \mu \alpha. \tau$
 $e ::= t$
 $t ::= x \mid () \mid n \mid \text{t pt} \mid \text{if0 t t t} \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t \mid t[\bar{\tau}]\bar{t} \mid \langle \bar{t} \rangle \mid \pi_i(t) \mid \text{pack}(\tau, t) \text{ as } \exists \alpha. \tau$
 $\quad \mid \text{unpack} \langle \alpha, x \rangle = t \text{ in } t \mid \text{fold}_{\mu \alpha. \tau} t \mid \text{unfold } t$
 $p ::= + \mid - \mid *$
 $v ::= () \mid n \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t \mid \langle \bar{v} \rangle \mid \text{pack}(\tau, v) \text{ as } \exists \alpha. \tau \mid \text{fold}_{\mu \alpha. \tau} v$
 $E ::= [\cdot] \mid E \text{ pt} \mid v \text{ p } E \mid \text{if0 } E \text{ t t} \mid E[\bar{\tau}]\bar{t} \mid v[\bar{\tau}]\bar{v} E \bar{t} \mid \dots$

$e \mapsto e' \quad E[\lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t[\bar{\tau}]\bar{v}] \mapsto E[t[\bar{\tau}'/\bar{\alpha}][\bar{v}/\bar{x}]] \quad \dots$

$\Delta; \Gamma \vdash e: \tau$ where $\Delta ::= \cdot \mid \Delta, \alpha$ and $\Gamma ::= \cdot \mid \Gamma, x: \tau$

$\tau ::= \alpha \mid \text{unit} \mid \text{int} \mid \forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau \mid \langle \bar{\tau} \rangle \mid \exists \alpha. \tau \mid \mu \alpha. \tau$
 $e ::= t$
 $t ::= x \mid () \mid n \mid \text{t pt} \mid \text{if0 t t t} \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t \mid t[\bar{\tau}]\bar{t} \mid \langle \bar{t} \rangle \mid \pi_i(t)$
 $\quad \mid \text{pack}(\tau, t) \text{ as } \exists \alpha. \tau \mid \text{unpack} \langle \alpha, x \rangle = t \text{ in } t \mid \text{fold}_{\mu \alpha. \tau} t \mid \text{unfold } t$
 $p ::= + \mid - \mid *$
 $v ::= () \mid n \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t \mid \langle \bar{v} \rangle \mid \text{pack}(\tau, v) \text{ as } \exists \alpha. \tau \mid \text{fold}_{\mu \alpha. \tau} v \mid v[\bar{\tau}]$
 $E ::= [\cdot] \mid \dots \mid E[\bar{\tau}]\bar{t} \mid v[\bar{\tau}]\bar{v} E \bar{t} \mid E[\bar{\tau}] \mid \dots$

$e \mapsto e' \quad E[\lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t[\bar{\tau}]\bar{v}] \mapsto E[t[\bar{\tau}'/\bar{\alpha}][\bar{v}/\bar{x}]] \quad \dots$

$\Delta; \Gamma \vdash e: \tau$ where $\Delta ::= \cdot \mid \Delta, \alpha$ and $\Gamma ::= \cdot \mid \Gamma, x: \tau$

$$\frac{\frac{\bar{\alpha}; \bar{x}:\bar{\tau} \vdash t: \tau'}{\Delta; \Gamma \vdash \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t: \forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'} \quad \frac{\Delta; \Gamma \vdash t: \forall[\cdot].(\bar{\tau}) \rightarrow \tau' \quad \Delta; \Gamma \vdash \bar{t}: \bar{\tau}}{\Delta; \Gamma \vdash t[\bar{\tau}]\bar{t}: \tau'}}{\Delta; \Gamma \vdash t: \forall[\beta, \bar{\alpha}].(\bar{\tau}) \rightarrow \tau'} \quad \frac{\Delta; \Gamma \vdash t: \forall[\beta, \bar{\alpha}].(\bar{\tau}) \rightarrow \tau' \quad \Delta \vdash \tau_0}{\Delta; \Gamma \vdash t[\tau_0]: \forall[\bar{\alpha}].(\bar{\tau}[\tau_0/\beta]) \rightarrow \tau'[\tau_0/\beta]} \quad \dots$$

$\tau ::= \alpha \mid \text{unit} \mid \text{int} \mid \exists \alpha. \tau \mid \mu \alpha. \tau \mid \text{ref } \psi \mid \text{box } \psi$
 $\psi ::= \forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau \mid \langle \tau, \dots, \tau \rangle$
 $e ::= (t, H)$
 $t ::= x \mid () \mid n \mid \text{t pt} \mid \text{if0 t t t} \mid \ell \mid t[\bar{\tau}]\bar{t} \mid t[\bar{\tau}] \mid \text{pack}(\tau, t) \text{ as } \exists \alpha. \tau \mid \text{unpack} \langle \alpha, x \rangle = t \text{ in } t$
 $\quad \mid \text{fold}_{\mu \alpha. \tau} t \mid \text{unfold } t \mid \text{ralloc} \langle \bar{t} \rangle \mid \text{balloc} \langle \bar{t} \rangle \mid \text{read}[i] t \mid \text{write } t [i] \leftarrow t$
 $p ::= + \mid - \mid *$
 $v ::= () \mid n \mid \text{pack}(\tau, v) \text{ as } \exists \alpha. \tau \mid \text{fold}_{\mu \alpha. \tau} v \mid \ell \mid v[\bar{\tau}]$
 $E ::= (E_t, \cdot) \quad E_t ::= [\cdot] \mid \dots \mid \text{balloc} \langle \bar{v}, E_t, \bar{t} \rangle \mid \dots$
 $h ::= \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t \mid \langle v, \dots, v \rangle \quad H ::= \cdot \mid H, \ell \mapsto h$

$\langle H \mid e \rangle \mapsto \langle H' \mid e' \rangle$ Reduction Relation (selected cases)

$$\langle H \mid (t, (H', \ell \mapsto h)) \rangle \mapsto \langle H, \ell' \mapsto h \mid (t[\ell'/\ell], H'[e'/\ell]) \rangle \quad \text{if } \ell' \notin \text{dom}(H)$$

$$\langle H \mid E[\ell[\bar{\tau}]\bar{v}] \rangle \mapsto \langle H \mid E[t[\bar{\tau}'/\bar{\alpha}][\bar{v}/\bar{x}]] \rangle \quad \text{if } H(\ell) = \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).t$$

$\Psi \vdash h: \psi$ where $\Psi ::= \cdot \mid \Psi, \ell: \text{ref } \psi \mid \Psi, \ell: \text{box } \psi$

$\Psi \vdash H: \Psi'$ which implies $\text{dom}(\Psi) \cap \text{dom}(\Psi') = \emptyset$

$\Psi; \Delta; \Gamma \vdash e: \tau$ where $\Delta ::= \cdot \mid \Delta, \alpha$ and $\Gamma ::= \cdot \mid \Gamma, x: \tau$

$$\frac{\Psi \vdash H: \Psi' \quad (\Psi, \Psi'); \Delta; \Gamma \vdash t: \tau}{\Psi; \Delta; \Gamma \vdash (t, H): \tau} \quad \dots$$

$$\frac{\Psi; \Delta; \Gamma \vdash \bar{t}: \bar{\tau} \quad \Psi; \Delta; \Gamma \vdash t: \text{box} \langle \tau_0, \dots, \tau_1, \dots, \tau_n \rangle}{\Psi; \Delta; \Gamma \vdash \text{ralloc} \langle \bar{t} \rangle: \text{box} \langle \bar{\tau} \rangle} \quad \frac{\Psi; \Delta; \Gamma \vdash t: \text{balloc} \langle \bar{t} \rangle: \text{box} \langle \bar{\tau} \rangle \quad \Psi; \Delta; \Gamma \vdash t: \text{read}[i] t: \tau}{\Psi; \Delta; \Gamma \vdash \text{read}[i] t: \tau}$$

Fig. 1. Definition of F (top), C (middle), and A (bottom)

to denote the free term variables of e and $\text{ftv}(e)$ (or $\text{ftv}(\tau)$) to denote the free type variables of e (or of type τ). We use a line above a syntactic element to indicate a list of repeated instances of this element, e.g., $\bar{\alpha} = \alpha_1, \dots, \alpha_n$ for $n \geq 0$. When the arities of different lists are required to match up in a definition or inference rule, these constraints will usually be obvious from context. Whenever two environments (e.g. Δ or Γ or Ψ) are joined by a comma, this should be interpreted as a *disjoint* union.

Source Language. Our source language F is System F with recursive types, existential types, and tuples. The syntax of types and terms in F is shown in Figure 1 (top). We combine type- and term-level abstractions of arbitrary arity into a single binding form $\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'$, abbreviating $\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'$ as $\bar{\tau} \rightarrow \tau'$. We define a small-step operational semantics for F (written $e \mapsto e'$) using evaluation contexts E to lift the primitive reductions to a standard left-to-right call-by-value semantics for the language. The reduction rules are standard; we show only the application rule.

F's typing judgment has the form $\Delta; \Gamma \vdash e : \tau$. The type environment Δ tracks the type variables in scope. The value environment Γ tracks the term variables in scope along with their types τ , which must be well formed under Δ (written $\Delta \vdash \tau$ and defined as $\text{ftv}(\tau) \subseteq \Delta$). The typing rules are standard and hence omitted.

Intermediate Language. Our intermediate language C, shown in Figure 1 (middle), is nearly identical to F, with two exceptions. First, since this language is the target of closure conversion, functions are not allowed to contain free type or term variables. Second, we allow the partial application of a function to a type. Hence, C terms include $t[\tau]$ and we consider $v[\tau]$ to be a value.

The reduction relation $e \mapsto e'$ is identical to that of F, and the typing judgment $\Delta; \Gamma \vdash e : \tau$ differs only in the rules for abstraction and application which are shown in the figure. Note that the body of a C function must typecheck in an environment that contains only the function's formal arguments.

Target Language. Our target A must serve as a target for heap allocation. Its design is similar to the language λ^A from [13]. Since we are compiling a source language without mutable references, it would suffice for A to provide only immutable references to functions and tuples that must now live on the heap. However, to provide a concrete illustration of the ability to link with target code that cannot be expressed in the source language, we augment A with mutable references to tuples.

The language A is shown in Figure 1 (bottom). Functions in A are stored only in immutable cells on the heap, while tuples are stored in heap cells that can be either mutable or immutable. We use ψ for the types of these *heap values* h . Mutable and immutable references have types **ref** ψ and **box** ψ , respectively. The terms **ralloc** $\langle \bar{t} \rangle$ and **balloc** $\langle \bar{t} \rangle$ —which allocate mutable and immutable cells, respectively—each allocate a new location ℓ and initialize it to the given tuple. The instructions **read** $[i]$ ℓ and **write** $\ell [i] \leftarrow v$ respectively read from and write the value v to the i -th slot in the tuple (of length n) stored at ℓ , assuming $0 \leq i < n$. The type system ensures that writes are only performed on mutable tuples.

Unlike F and C, the syntax of A distinguishes components e from terms t . A component e pairs a term t with a *heap fragment* H . H can contain functions and tuples that t may use by referring to locations in H . Intuitively, we need this notion of components because a bare term t is not as expressive as C component. In particular, A does not

provide any way to dynamically allocate a location and initialize it to a function. We discuss how the compiler produces components with heap fragments in §4.

Heap fragments are assigned heap types Ψ . A heap fragment may reference locations that are to be linked in by another component, so the judgment $\Psi \vdash \mathbf{H} : \Psi'$ includes an external heap type Ψ as an environment used in assigning \mathbf{H} the type Ψ' . Here, Ψ' must provide types for exactly the locations in \mathbf{H} . Each h in \mathbf{H} must typecheck under the disjoint union of the two heap types (Ψ, Ψ') . Similarly, a component (\mathbf{t}, \mathbf{H}) can reference both external locations and those bound by \mathbf{H} , that is, locations in the domain of either the external heap type Ψ or of \mathbf{H} .

Our operational semantics for A is a relation between configurations $\langle \mathbf{H} \mid \mathbf{e} \rangle$. Any code or data in the internal heap fragment of component \mathbf{e} must be loaded into memory before it can be run. We formally capture this with a reduction rule that “loads” a component by merging its internal heap fragment with the external heap. When loading a component (\mathbf{t}, \mathbf{H}) , we must rename the locations bound in \mathbf{H} so that they do not conflict with the external heap. After the loading step, the term component \mathbf{t} can be evaluated using standard reduction rules.

The structure of A components also entails a small change to the structure of evaluation contexts, which are defined in two layers: contexts \mathbf{E} expect components \mathbf{e} , and term contexts \mathbf{E}_t expect terms \mathbf{t} . Terms are plugged into term contexts in the obvious way. Plugging a component-level evaluation context $\mathbf{E} = (\mathbf{E}_t, \cdot)$ with a component \mathbf{e} is defined by $(\mathbf{E}_t, \cdot)[(\mathbf{t}, \mathbf{H})] = (\mathbf{E}_t[\mathbf{t}], \mathbf{H})$

4 The Compiler

Compiling F to C. Closure conversion collects a function’s free term variables in a tuple called the *closure environment* that is passed as an additional argument to the function, thus turning the function into a closed term. The closed function is paired with its environment to create a *closure*. The basic idea of typed closure conversion goes back to Minamide *et al.* [14], whom we follow in using an existential type to abstract the type of the environment. This ensures that two functions with the same type but different free variables still have the same type after closure conversion: the abstract type hides the fact that the closures’ environments have different types.

We must also rewrite functions to take their free type variables as additional arguments. However, instead of collecting these types in a type environment as Minamide *et al.* do, we follow Morrisett *et al.* [13] and directly substitute the types into the function. Like the latter, we adopt a *type-erasure* interpretation, which means that since all types are erased at run time the substitution of types into functions has no run-time effect.

Our closure-conversion pass compiles F terms of type τ to C terms of type τ^C . Figure 2 (top) presents the type translation τ^C and some of the compilation rules. Since this is closure conversion, the only interesting parts are those that involve functions. The omitted rules are defined by structural recursion on terms.

Compiling C to A. Our second compiler pass combines hoisting of functions with explicit allocation of tuples. It takes a C component (that is, just a C term \mathbf{t}) of type τ , and produces an A term \mathbf{t} as well as a heap fragment \mathbf{H} with all the hoisted functions.




Fig. 2. Compiler from F to C (top) and from C to A (bottom)

The component (t, H) is the overall output, and has type τ^A under an empty external heap. The heap fragment generated by the compiler does not contain tuples: the compiler translates C tuples by generating **balloc** expressions, not by putting them in a static heap fragment. The type translation and interesting parts of the term translation are shown in Figure 2 (bottom).

5 F and C Interoperability

5.1 The Basics

We now present a formal semantics for interoperability between F and C. For now, we define a combined language FC; in §6, we will extend this to FCA. Our FC multi-language system embeds the languages F and C so that both languages have natural access to foreign values (i.e., values from the other language). In particular, we want F components of type τ to be usable as C components of type τ^C , and vice versa. To allow cross-language communication, FC extends the original F and C with syntactic boundaries, written ${}^{\tau}\mathcal{F}C\ e$ (C inside, F outside) and $C\mathcal{F}^{\tau}\ e$ (F inside, C outside).

The interesting cases in the semantics of boundaries are those that handle universal and existential types. These must be defined carefully to ensure that type abstraction is not broken as values pass between languages. First, though, we explain the general principles of our boundary semantics by looking at the cases for simple types and their translations.

CF Boundary Semantics. A term $C\mathcal{F}^{\tau}\ e$ has type τ^C if e has type τ . To evaluate this boundary term, FC's operational semantics require first that e be reduced to a value v (using F reduction rules). Then a type-directed meta-function is applied to v , yielding a value in C of type τ^C (written $CF^{\tau}(v) = v$). An important restriction on this meta-function, which we call the *value translation*, is that it is only defined for *closed* values. This is sufficient for our needs because it is used only by the FC operational semantics, and substitution-based reduction relations are defined only for closed programs. We can still write FC programs with free variables appearing under boundaries, but by the time we evaluate the boundary term, we will have supplied values for all of these free variables.

At base types, value translation is easy: for example, translating a value n of type `int` yields the same integer in C, `n`. Most of the other types are translated simply by structural recursion.

The interesting case is the case for function types. Consider the translation of a value v of type $\tau \rightarrow \tau'$. As per the type translation, this should produce a value of type $\exists\beta.(((\beta, \tau^C) \rightarrow \tau'^C), \beta)$. Since v is closed, we can simply use `unit` for the type β of the closure environment:

$$CF^{\tau \rightarrow \tau'}(v) = \text{pack}\langle \text{unit}, \langle v, () \rangle \rangle \text{ as } \exists\beta.(((\beta, \tau^C) \rightarrow \tau'^C), \beta)$$

We must still construct the underlying function v for this closure, which we can do using boundary terms and the original function v :

$$v = \lambda(z: \text{unit}, x: \tau^C). C\mathcal{F}^{\tau'}(v\ {}^{\tau}\mathcal{F}C\ x).$$

The function we build simply translates its argument from C to F, applies v to the translated argument, and finally translates the result back into C.

The full translation rule for functions must also handle type arguments and requires some additional machinery, which we will discuss momentarily.

FC Boundary Semantics. The term ${}^{\tau}\mathcal{F}C\ e$ has type τ when e has type τ^C . As before, to evaluate a boundary term, we first evaluate the component under the boundary, this time to a value v . Then we apply a value translation ${}^{\tau}\mathcal{F}C(v) = v$ that yields an F value v of type τ . Again, this translation is only defined for closed values of translation type.

Let us consider the type $\tau \rightarrow \tau'$ again. A closure \mathbf{v} of type $(\tau \rightarrow \tau')^{\mathbf{C}}$ must be translated to an F function that first translates its argument from F to C, then unpacks the closure \mathbf{v} and applies the code to its environment and the translated argument, and finally translates the result back from C to F:

$$\tau \rightarrow \tau' \mathbf{FC}(\mathbf{v}) = \lambda(x: \tau). \tau' \mathcal{FC}(\mathbf{unpack} \langle \beta, y \rangle = \mathbf{v} \text{ in } \pi_1(y) \pi_2(y) \mathcal{CF}^\tau x)$$

In both function cases, notice that the direction of the conversion (and the boundary used) reverses for function arguments.

5.2 Handling Abstract Types

Now that we have established the general structure of boundary rules, we come to the interesting cases, those for abstract types.

FC Type Abstraction. Consider the type $\forall[\alpha].(\alpha) \rightarrow \alpha$. Since $\alpha^{\mathbf{C}} = \alpha$, the translation of this type is

$$(\forall[\alpha].(\alpha) \rightarrow \alpha)^{\mathbf{C}} = \exists \beta. \langle (\forall[\alpha].(\beta, \alpha) \rightarrow \alpha), \beta \rangle.$$

If we naively try to extend the function case of the value translation given above, we get the following:

$$\forall[\alpha].(\alpha) \rightarrow \alpha \mathbf{FC}(\mathbf{v}) = \lambda[\alpha](x: \alpha). \alpha \mathcal{FC}(\mathbf{unpack} \langle \beta, y \rangle = \mathbf{v} \text{ in } \pi_1(y) [\alpha^{\mathbf{C}}] \pi_2(y) \mathcal{CF}^{\alpha} x)$$

Note that we have not expanded $\alpha^{\mathbf{C}}$ in the application produced by this translation. It would expand to a C type variable α , but we cannot allow this, because that α would be unbound! What we really want is that when α is instantiated with a concrete type τ , the positions inside language C where that type is needed receive $\tau^{\mathbf{C}}$.

We resolve this by making two changes to our system: first, we add a type $[\alpha]$ (which may be read as “ α suspended in C”) that allows an F type variable to appear in a C type. The F type variable α needs to be translated, but the translation is *delayed* until α is instantiated with a concrete type. We enforce this semantics in the definition of type substitution: $[\alpha][\tau/\alpha] = \tau^{\mathbf{C}}$.

Second, we adjust the type translation to turn F type variables into suspended type variables instead of C type variables. We call this modified version of the type translation the *boundary type translation*, and notate it by $\tau^{(\mathbf{C})}$. Formally, the rule for type variables in the compiler’s type translation is replaced by the rule $\alpha^{(\mathbf{C})} = [\alpha]$ in the boundary type translation. We only want to suspend free type variables, so when we translate a type that contains bound variables, we need to restore the behavior of the compiler’s type translation when we translate the binding position. We can do this using a substitution, e.g., $(\exists \alpha. \tau)^{(\mathbf{C})} = \exists \alpha. (\tau^{(\mathbf{C})} [\alpha/[\alpha]])$. Thus the boundary type translation preserves the binding structure of the type to which it is applied.

With these two changes, we can correct the example above by replacing the appearance of $\alpha^{\mathbf{C}}$ with $\alpha^{(\mathbf{C})}$, and we get a sensible translation from C to F for values of type $(\forall[\alpha].(\alpha) \rightarrow \alpha)^{\mathbf{C}}$.

CF Type Abstraction. Next, consider translating values of type $\forall[\alpha].(\alpha) \rightarrow \alpha$ from F into C. Once again, the existing machinery is not quite sufficient. Here is a naive attempt:

$$\mathbf{CF}^{\forall[\alpha].(\alpha) \rightarrow \alpha}(\mathbf{v}) = \mathbf{pack} \langle \mathbf{unit}, \langle \mathbf{v}, () \rangle \rangle \text{ as } (\forall[\alpha].(\alpha) \rightarrow \alpha)^{(\mathbf{C})}$$

where $\mathbf{v} = \lambda[\alpha](z: \mathbf{unit}, x: \alpha). \mathcal{CF}^{\alpha}(\mathbf{v} [\alpha]^{\alpha} \mathcal{FC} x)$.

$\tau ::= \dots \mid \mathbf{L}(\tau)$	$\tau ::= \dots \mid \lceil \alpha \rceil$	$\tau ::= \tau \mid \tau$
$\mathbf{t} ::= \dots \mid \mathcal{F}\mathcal{C}\mathbf{e}$	$\mathbf{t} ::= \dots \mid \mathcal{C}\mathcal{F}\tau\mathbf{e}$	$e ::= \mathbf{e} \mid \mathbf{e}$
$\mathbf{v} ::= \dots \mid \mathbf{L}(\tau)\mathcal{F}\mathcal{C}\mathbf{v}$	$\mathbf{v} ::= \dots$	$v ::= \mathbf{v} \mid \mathbf{v} \quad \Delta ::= \cdot \mid \Delta, \alpha \mid \Delta, \alpha$
$\mathbf{E} ::= \dots \mid \mathcal{F}\mathcal{C}\mathbf{E}$	$\mathbf{E} ::= \dots \mid \mathcal{C}\mathcal{F}\tau\mathbf{E}$	$E ::= \mathbf{E} \mid \mathbf{E} \quad \Gamma ::= \cdot \mid \Gamma, \mathbf{x} : \tau \mid \Gamma, \mathbf{x} : \tau$

$\tau^{(C)}$	Boundary Type Translation
$\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau^{(C)} = \exists\beta. \langle (\forall[\bar{\alpha}].(\beta, \overline{\tau^{(C)}[\alpha/\lceil\alpha\rceil]}) \rightarrow \tau^{(C)}[\alpha/\lceil\alpha\rceil]), \beta \rangle$	
$\alpha^{(C)} = \lceil \alpha \rceil$	$\text{unit}^{(C)} = \text{unit} \quad \text{int}^{(C)} = \text{int} \quad \exists\alpha.\tau^{(C)} = \exists\alpha.\tau^{(C)}[\alpha/\lceil\alpha\rceil]$
$\mu\alpha.\tau^{(C)} = \mu\alpha.\tau^{(C)}[\alpha/\lceil\alpha\rceil]$	$\langle \bar{\tau} \rangle^{(C)} = \overline{\tau^{(C)}} \quad \mathbf{L}\langle \tau \rangle^{(C)} = \tau$

Type Substitution: $\lceil \alpha \rceil[\tau/\alpha] = \tau^{(C)}$

$\Delta; \Gamma \vdash e : \tau$	Include F and C rules, with environments replaced by $\Delta; \Gamma$
$\frac{\Delta; \Gamma \vdash e : \tau^{(C)}}{\Delta; \Gamma \vdash \mathcal{F}\mathcal{C}\mathbf{e} : \tau}$	$\frac{\Delta; \Gamma \vdash \mathbf{e} : \tau}{\Delta; \Gamma \vdash \mathcal{C}\mathcal{F}\tau\mathbf{e} : \tau^{(C)}}$

$\mathbf{CF}^\tau(\mathbf{v}) = \mathbf{v}$	Value Translation	$\mathbf{CF}^{\text{unit}}(\cdot) = ()$	$\mathbf{CF}^{\text{int}}(n) = n$	$\mathbf{CF}^{\mathbf{L}(\tau)}(\mathbf{L}(\tau)\mathcal{F}\mathcal{C}\mathbf{v}) = \mathbf{v}$
$\mathbf{CF}^{\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'}(\mathbf{v}) = \text{pack}(\text{unit}, \langle \mathbf{v}, () \rangle) \text{ as } \overline{\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'}^{(C)}$				
where $\mathbf{v} = \lambda[\bar{\alpha}](\mathbf{z} : \text{unit}, \mathbf{x} : \tau^{(C)}[\alpha/\lceil\alpha\rceil]). \mathcal{C}\mathcal{F}\tau'[\mathbf{L}(\alpha)/\alpha](\mathbf{v}[\mathbf{L}(\alpha)] \tau[\mathbf{L}(\alpha)/\alpha]\mathcal{F}\mathcal{C}\mathbf{x})$				
$\mathbf{CF}^{\exists\alpha.\tau}(\text{pack}(\tau', \mathbf{v}) \text{ as } \exists\alpha.\tau) = \text{pack}(\tau'^{(C)}, \mathbf{v}) \text{ as } \exists\alpha.\tau^{(C)}$			where $\mathbf{CF}^{\tau'/\alpha}(\mathbf{v}) = \mathbf{v}$	
$\mathbf{CF}^{\mu\alpha.\tau}(\text{fold}_{\mu\alpha.\tau}\mathbf{v}) = \text{fold}_{\mu\alpha.\tau}(\mathbf{v})$			where $\mathbf{CF}^{\tau[\mu\alpha.\tau/\alpha]}(\mathbf{v}) = \mathbf{v}$	
$\mathbf{CF}^{\langle \tau_1, \dots, \tau_n \rangle}(\langle \mathbf{v}_1, \dots, \mathbf{v}_n \rangle) = \langle \mathbf{v}_1, \dots, \mathbf{v}_n \rangle$			where $\mathbf{CF}^{\tau_i}(\mathbf{v}_i) = \mathbf{v}_i$	

$\tau\mathbf{FC}(\mathbf{v}) = \mathbf{v}$	Value Translation	$\text{unit}\mathbf{FC}(\cdot) = ()$	$\text{int}\mathbf{FC}(n) = n$	$\mathbf{L}(\tau)\mathbf{FC}(\mathbf{v}) = \mathbf{L}(\tau)\mathcal{F}\mathcal{C}\mathbf{v}$
$\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau' \mathbf{FC}(\mathbf{v}) = \lambda[\bar{\alpha}](\bar{\mathbf{x}} : \bar{\tau}). \tau' \mathcal{F}\mathcal{C}(\text{unpack} \langle \beta, \mathbf{y} \rangle = \mathbf{v} \text{ in } \pi_1(\mathbf{y}) \lceil \lceil \alpha \rceil \rceil \pi_2(\mathbf{y}), \overline{\mathcal{C}\mathcal{F}\tau\bar{\mathbf{x}}})$				
$\exists\alpha.\tau\mathbf{FC}(\text{pack}(\tau', \mathbf{v}) \text{ as } \exists\alpha.\tau^{(C)}) = \text{pack}(\mathbf{L}(\tau'), \mathbf{v}) \text{ as } \exists\alpha.\tau$			where $\tau[\mathbf{L}(\tau')/\alpha]\mathbf{FC}(\mathbf{v}) = \mathbf{v}$	
$\mu\alpha.\tau\mathbf{FC}(\text{fold}_{\mu\alpha.\tau}(\mathbf{v})) = \text{fold}_{\mu\alpha.\tau}(\mathbf{v})$			where $\tau[\mu\alpha.\tau/\alpha]\mathbf{FC}(\mathbf{v}) = \mathbf{v}$	
$\langle \tau_1, \dots, \tau_n \rangle \mathbf{FC}(\langle \mathbf{v}_1, \dots, \mathbf{v}_n \rangle) = \langle \mathbf{v}_1, \dots, \mathbf{v}_n \rangle$			where $\tau_i \mathbf{FC}(\mathbf{v}_i) = \mathbf{v}_i$	

$e \mapsto e'$	Include F and C rules, replacing eval. contexts \mathbf{E}, \mathbf{E} with E .
$\frac{\mathbf{CF}^\tau(\mathbf{v}) = \mathbf{v}}{E[\mathcal{C}\mathcal{F}\tau\mathbf{v}] \mapsto E[\mathbf{v}]}$	$\frac{\tau\mathbf{FC}(\mathbf{v}) = \mathbf{v} \quad \tau \neq \mathbf{L}(\tau)}{E[\tau\mathcal{F}\mathcal{C}\mathbf{v}] \mapsto E[\mathbf{v}]}$

Fig. 3. FC multi-language system (extends F and C from Figure 1)

This time, we have translated the binder for α into a C binder for α , but we are left with free occurrences of α in the result! This is not a suitable translation, as we must produce a closed value. Note that the boundary terms in the body of \mathbf{v} expect to be annotated with a type that translates to α .

To fix this problem, we introduce a *lump type* $\mathbf{L}(\tau)$ that allows us to pass C values to F terms as opaque lumps. The introduction form for the lump type is the boundary term $\mathbf{L}(\tau)\mathcal{F}\mathcal{C}\mathbf{e}$, and the elimination form is $\mathcal{C}\mathcal{F}\tau\mathbf{e}$. A pair of opposite boundaries at lump type cancel, to yield the underlying C value. We extend the boundary type translation by defining $\mathbf{L}(\tau)^{(C)} = \tau$.

Now the three free occurrences of α in \mathbf{v} can be replaced with $\mathbf{L}(\alpha)$, yielding a well-typed translation.

Summary. With the additional tools of lumps, suspensions, and the boundary type translation, we have now developed everything needed for the FC multi-language system. Figure 3 presents more of the details, including the complete value translations.

The syntax of FC simply combines the syntax of F with that of C, and adds boundaries, lumps, and suspensions. The type judgment combines the type rules for F and C, but with the environments replaced by environments that can contain variables from both languages. We also add rules to typecheck boundary terms.

The cases of the value translations we have not yet covered mostly proceed by structural recursion, but note that the cases for existential types need to make use of lumps and suspensions (the suspensions are introduced by the boundary type translation) in ways that are dual to the function cases.

The reduction relation combines the reduction rules from F and C and adds rules for boundaries. The boundary reduction rules use the value translations to produce a value in the other language.

$$\begin{array}{lcl}
 \tau ::= \dots & | & \mathbf{L}\langle\tau\rangle \\
 \mathbf{t} ::= \dots & | & \mathcal{TC}\mathcal{A}\mathbf{e} \\
 \mathbf{v} ::= \dots & | & \mathbf{L}\langle\tau\rangle\mathcal{C}\mathcal{A}\mathbf{v} \\
 \mathbf{E} ::= \dots & | & \mathcal{TC}\mathcal{A}\mathbf{E}
 \end{array}
 \quad
 \begin{array}{lcl}
 \tau ::= \dots & | & [\alpha] \mid [\alpha] \\
 \mathbf{t} ::= \dots & | & \mathcal{AC}^\tau\mathbf{e} \\
 \mathbf{v} ::= \dots & & \\
 \mathbf{E}_t ::= \dots & | & \mathcal{AC}^\tau\mathbf{E}
 \end{array}
 \quad
 \begin{array}{lcl}
 \tau ::= \dots & | & \tau \\
 \mathbf{e} ::= \dots & | & \mathbf{e} \\
 \mathbf{v} ::= \dots & | & \mathbf{v} \\
 \mathbf{E} ::= \dots & | & \mathbf{E}
 \end{array}
 \quad
 \begin{array}{lcl}
 \Delta ::= \dots & | & \Delta, \alpha \\
 \Gamma ::= \dots & | & \Gamma, \mathbf{x} : \tau
 \end{array}$$

$\tau^{\langle\mathcal{A}\rangle}$	Boundary Type Translation
$\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'^{\langle\mathcal{A}\rangle} = \mathbf{box} \forall[\bar{\alpha}].(\tau^{\langle\mathcal{A}\rangle}[\bar{\alpha}/[\bar{\alpha}]] \rightarrow \tau'^{\langle\mathcal{A}\rangle}[\bar{\alpha}/[\bar{\alpha}]])$ $\alpha^{\langle\mathcal{A}\rangle} = [\alpha] \quad \dots \quad \mathbf{L}\langle\tau\rangle^{\langle\mathcal{A}\rangle} = \tau \quad [\alpha]^{\langle\mathcal{A}\rangle} = [\alpha]$	
Type Substitution: $[\alpha][\tau/\alpha] = (\tau^{\langle\mathcal{C}\rangle})^{\langle\mathcal{A}\rangle} \quad [\alpha][\tau/\alpha] = \tau^{\langle\mathcal{A}\rangle}$	
$\Psi; \Delta; \Gamma \vdash e : \tau$	Include A rules and add Ψ to existing rules
$\frac{\Psi; \Delta; \Gamma \vdash e : \tau^{\langle\mathcal{A}\rangle}}{\Psi; \Delta; \Gamma \vdash \mathcal{TC}\mathcal{A}\mathbf{e} : \tau} \qquad \frac{\Psi; \Delta; \Gamma \vdash \mathbf{e} : \tau}{\Psi; \Delta; \Gamma \vdash \mathcal{AC}^\tau\mathbf{e} : \tau^{\langle\mathcal{A}\rangle}}$	
$\mathcal{AC}^\tau(\mathbf{v}, \mathbf{H}) = (\mathbf{v}, \mathbf{H}')$	Value Translation (selected cases)
$\mathcal{AC}^{\mathbf{unit}}((\), \mathbf{H}) = ((\), \mathbf{H})$	
$\mathcal{AC}^{\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau'}(\mathbf{v}, \mathbf{H}) = (\ell, (\mathbf{H}, \ell \mapsto \mathbf{h}))$ <p style="margin-left: 40px;">where $\mathbf{h} = \lambda[\bar{\alpha}](\mathbf{x} : \tau^{\langle\mathcal{A}\rangle}[\bar{\alpha}/[\bar{\alpha}]]) . \mathcal{AC}^{\tau'[\bar{\mathbf{L}}\langle\alpha\rangle/\bar{\alpha}]}\mathbf{v}[\bar{\mathbf{L}}\langle\alpha\rangle] \tau^{\overline{[\bar{\mathbf{L}}\langle\alpha\rangle/\bar{\alpha}]}}\mathcal{C}\mathcal{A}\mathbf{x}$</p>	
$\mathcal{AC}^{\langle\bar{\tau}\rangle}(\langle\bar{\mathbf{v}}\rangle, \mathbf{H}_1) = (\ell, (\mathbf{H}_{n+1}, \ell \mapsto \langle\bar{\mathbf{v}}\rangle))$ <p style="margin-left: 40px;">where $\mathcal{AC}^{\tau_i}(\mathbf{v}_i, \mathbf{H}_i) = (\mathbf{v}_i, \mathbf{H}_{i+1})$</p>	
$\mathcal{TC}\mathcal{A}(\mathbf{v}, \mathbf{H}) = (\mathbf{v}, \mathbf{H}')$	Value Translation (selected cases)
$\mathbf{unit}\mathcal{CA}((\), \mathbf{H}) = ((\), \mathbf{H})$	
$\forall[\bar{\alpha}].(\bar{\tau}) \rightarrow \tau' \mathcal{CA}(\mathbf{v}, \mathbf{H}) = (\lambda[\bar{\alpha}](\bar{\mathbf{x}} : \bar{\tau}) . \mathcal{TC}\mathcal{A}(\mathbf{v}[\bar{\alpha}], \mathcal{AC}^\tau\bar{\mathbf{x}}), \mathbf{H})$ <p style="margin-left: 40px;">where $\mathbf{H}_1(\ell) = \langle\bar{\mathbf{v}}\rangle$ and $\mathcal{TC}\mathcal{A}(\mathbf{v}_i, \mathbf{H}_i) = (\mathbf{v}_i, \mathbf{H}_{i+1})$</p>	
$\langle\bar{\tau}\rangle\mathcal{CA}(\ell, \mathbf{H}_1) = (\langle\bar{\mathbf{v}}\rangle, \mathbf{H}_{n+1})$	
$\langle\mathbf{H} \mid e\rangle \mapsto \langle\mathbf{H}' \mid e'\rangle$	Lift FC rules to new config.; replace \mathbf{E} with \mathbf{E}
$\frac{\mathcal{AC}^\tau(\mathbf{v}, \mathbf{H}) = (\mathbf{v}, \mathbf{H}')}{\langle\mathbf{H} \mid E[\mathcal{AC}^\tau\mathbf{v}]\rangle \mapsto \langle\mathbf{H}' \mid E[\mathbf{v}]\rangle} \quad \frac{\mathcal{TC}\mathcal{A}(\mathbf{v}, \mathbf{H}) = (\mathbf{v}, \mathbf{H}') \quad \tau \neq \mathbf{L}\langle\tau\rangle}{\langle\mathbf{H} \mid E[\mathcal{TC}\mathcal{A}\mathbf{v}]\rangle \mapsto \langle\mathbf{H}' \mid E[\mathbf{v}]\rangle}$	

Fig. 4. FCA multi-language system (extends Figures 1 and 3)

6 C and A Interoperability

The extensions to FC for interoperability with A are given in Figure 4. The principles discussed in the development of FC still apply, but here we need to handle the presence of the heap. Specifically, since functions and tuples in A are contained in the heap, the value translations need access to the program’s memory. Going from C to A, the value translation may allocate new memory for functions and tuples; going from A to C requires looking up the contents of locations and translating those contents to functions or tuples in C. Thus, we pass the current memory as an argument to the translations, and return a memory that may have had additional locations allocated. Memory cells allocated by boundaries are always immutable.

Aside from this change, the extension for the new language mostly follows what we did for FC: we augment the syntax with boundaries between C and A, a lump type $L\langle\tau\rangle$ for opaquely embedding A values into C, and suspensions of type variables into A. Note that we need the boundary type translation from C to A to handle both C type variables α and suspended F type variables $[\alpha]$. Thus A has both $[\alpha]$ and $[\alpha]$ as suspension types. The boundary type translation $\tau^{(A)}$ works similarly to $\tau^{(C)}$. The figure shows the function case and the cases involving lumps and suspensions. The type judgment merges the A type rules with the FC type rules, but where the latter are modified to add the extra environment Ψ , and adds type rules for boundaries. Finally, the reduction relation for FCA lifts the FC reductions to use the configuration from A, with a program heap. We also add the reduction rules from A and a pair of boundary reduction rules that utilize the value translations.

7 Compiler Correctness

As mentioned in §1, we state compiler correctness in terms of FCA contextual equivalence. Below, we formally define contextual equivalence for FCA components and then present our compiler correctness theorems. We discuss how to prove these theorems in §9 and give a longer discussion and the full proofs in the technical report [9].

7.1 FCA Contextual Equivalence

A general context C is an FCA component with a hole. A component e can be plugged into the context only if it is from the same language as the hole. Since contexts can contain boundaries, e need not be from the same language as the outermost layer of C . The syntax of general contexts is given in Figure 5 (top). Contexts for F and C forms are standard. In A, we need contexts to be able to have their hole in either the term part of a component, or in the body of a function contained in the heap fragment. So in addition to contexts C that produce components, we have context forms C_t and C_H that produce terms and heap fragments, respectively.

When plugging an A component (t, H) into a context C , the heap fragment H is placed at the innermost component-level layer of C —that is, at the language boundary closest to the hole—and merged with the heap fragment already in that position. To formalize this, the A portion of the definition of plugging a component into a context is given in Figure 5 (middle). The definition of plugging for F and C contexts is standard.

$$\begin{aligned}
C &::= [\cdot] \mid \mathbf{C} \mathbf{p} \mathbf{t} \mid \dots \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).C \mid \dots \mid {}^{\tau}\mathcal{F}C \\
\mathbf{C} &::= [\cdot] \mid \dots \mid \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).\mathbf{C} \mid \dots \mid \mathcal{C}\mathcal{F}^{\tau}\mathbf{C} \mid {}^{\tau}\mathcal{C}\mathcal{A}\mathbf{C} \\
\mathbf{C} &::= (\mathbf{C}_t, \mathbf{H}) \mid (t, \mathbf{C}_H) \\
\mathbf{C}_t &::= [\cdot] \mid \dots \mid \mathcal{A}\mathcal{C}^{\tau}\mathbf{C} \quad \mathbf{C}_H ::= \mathbf{C}_H, \ell \mapsto \mathbf{h} \mid \mathbf{H}, \ell \mapsto \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).\mathbf{C}_t \\
C &::= \mathbf{C} \mid \mathbf{C} \mid \mathbf{C}
\end{aligned}$$

$C[e]$ Context Plugging (A cases shown)

$$\begin{aligned}
(\mathbf{C}_t, \mathbf{H})[e] &= \begin{cases} (\mathbf{C}_t[t], (\mathbf{H}, \mathbf{H}')) & e = (t, \mathbf{H}') \wedge \mathbf{C}_t \text{ contains no language boundaries} \\ (\mathbf{C}_t[e], \mathbf{H}) & \text{otherwise} \end{cases} \\
(t, \mathbf{C}_H)[e] &= \begin{cases} (t, (\mathbf{C}_H[t'], \mathbf{H}')) & e = (t', \mathbf{H}') \wedge \mathbf{C}_H \text{ contains no language boundaries} \\ (t, \mathbf{C}_H[e]) & \text{otherwise} \end{cases} \\
[\cdot][t] &= t & (\mathbf{C}_t \mathbf{p} \mathbf{t})[e] &= (\mathbf{C}_t[e]) \mathbf{p} \mathbf{t} & \dots \\
& & (\mathbf{C}_H, \ell \mapsto \mathbf{h})[e] &= (\mathbf{C}_H[e]), \ell \mapsto \mathbf{h} \\
& & (\mathbf{H}, \ell \mapsto \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).\mathbf{C}_t)[e] &= \mathbf{H}, \ell \mapsto \lambda[\bar{\alpha}](\bar{x}:\bar{\tau}).(\mathbf{C}_t[e])
\end{aligned}$$

$\vdash C : (\Psi; \Delta; \Gamma \vdash \tau) \rightsquigarrow (\Psi'; \Delta'; \Gamma' \vdash \tau')$ Context Typing (omitted)

Contextual Equivalence

$$\begin{aligned}
\Psi; \Delta; \Gamma \vdash e_1 \approx^{ctx} e_2 : \tau &\stackrel{\text{def}}{=} \Psi; \Delta; \Gamma \vdash e_1 : \tau \wedge \Psi; \Delta; \Gamma \vdash e_2 : \tau \wedge \\
&\forall C, \mathbf{H}, \Psi', \tau'. \vdash C : (\Psi; \Delta; \Gamma \vdash \tau) \rightsquigarrow (\Psi'; \cdot; \cdot \vdash \tau') \wedge \vdash \mathbf{H} : \Psi' \\
&\implies \langle \langle \mathbf{H} \mid C[e_1] \rangle \downarrow \rangle \iff \langle \langle \mathbf{H} \mid C[e_2] \rangle \downarrow \rangle
\end{aligned}$$

Fig. 5. General Contexts & Contextual Equivalence for FCA

Given this notion of general contexts, contextual equivalence for FCA is standard (see Figure 5, bottom). It says that two components e_1 and e_2 are contextually equivalent under environments Ψ, Δ, Γ and at type τ if the following hold: First, both components must typecheck under Ψ, Δ, Γ at type τ . Second, if C is a context that expects to be given a component that typechecks under Ψ, Δ, Γ at type τ , and produces a resulting program that is closed but expects to be run with a heap of type Ψ' , then $C[e_1]$ and $C[e_2]$ have the same termination behavior when we run them with any initial heap \mathbf{H} that has type Ψ' .

7.2 Compiler Correctness

We can now state our main result: compiler-correctness theorems for both passes of our compiler.

Theorem 1 (Closure Conversion is Semantics-Preserving). *If $\bar{\alpha}; \overline{x:\tau'} \vdash e : \tau \rightsquigarrow e$, then $\cdot; \bar{\alpha}; \overline{x:\tau'} \vdash e \approx^{ctx} {}^{\tau}\mathcal{F}C(e[\overline{[\alpha]}/\alpha] \overline{[\mathcal{C}\mathcal{F}^{\tau'}x/x]}): \tau$.*

Theorem 2 (Allocation is Semantics-Preserving). *If $\bar{\alpha}; \overline{x:\tau'} \vdash e : \tau \rightsquigarrow (t, \mathbf{H} : \Psi)$, then $\cdot; \bar{\alpha}; \overline{x:\tau'} \vdash e \approx^{ctx} {}^{\tau}\mathcal{C}\mathcal{A}(t[\overline{[\alpha]}/\alpha] \overline{[\mathcal{A}\mathcal{C}^{\tau'}x/x]}, \mathbf{H}): \tau$.*

The formal theorems are essentially as we described our compiler correctness results in §1, with only one additional subtlety: we need to perform a substitution so that

the free variables of the original component match those of the compiled component. Recall that the compiler turns free type and term variables α and x into type and term variables $\bar{\alpha}$ and \bar{x} from the next language, whereas FCA needs the binding structure of components to be preserved, including free variables being in the language prescribed by the type environments Δ and Γ . To get the free variables of the two components back into sync, we substitute suspended type variables for translated type variables, and we substitute boundary terms for translated term variables. Note that we do not need to perform a substitution in the heap fragment produced by the allocation pass, since heap values must be closed anyway.

We could equivalently have stated these theorems with the substitution on the other side, and the environments correspondingly translated; e.g.

$$\cdot; \bar{\alpha}^c; \bar{x}; \tau^c \vdash e[\overline{L\langle\alpha\rangle/\alpha}] [\hat{\tau}' \mathcal{F}C \bar{x}/\bar{x}] \approx^{ctx} \hat{\tau} \mathcal{F}C e; \hat{\tau},$$

where $\hat{\tau} = \tau[\overline{L\langle\alpha\rangle/\alpha}]$ and $\hat{\tau}' = \tau'[\overline{L\langle\alpha\rangle/\alpha}]$.

It also does not matter which side the boundary term is placed on: boundary cancellation lemmas allow us to prove as a corollary that, for example,

$$\cdot; \bar{\alpha}; \bar{x}; \bar{\tau} \vdash \mathcal{C}\mathcal{F}^T e \approx^{ctx} e[\overline{[\bar{\alpha}]/\alpha}] [\mathcal{C}\mathcal{F}^T \bar{x}/\bar{x}]; \tau^{(c)}.$$

Since we want to ensure that type variables in the environment remain tied to their free occurrences in the result type, this version of the theorem uses the boundary type translation $\tau^{(c)}$ for the result type (instead of the compiler's type translation τ^c).

Contextual equivalence is transitive, so we can easily chain these theorems together to prove correctness for the full compiler:

Corollary 1 (Compiler Correctness). *If $\bar{\alpha}; \bar{x}; \bar{\tau}' \vdash e; \tau \rightsquigarrow e \rightsquigarrow e$, then*

$$\cdot; \bar{\alpha}; \bar{x}; \bar{\tau}' \vdash e \approx^{ctx} \tau \mathcal{F}C \mathcal{A}(e[\overline{[\bar{\alpha}]/\alpha}] [\mathcal{A}\mathcal{C}\mathcal{F}^T \bar{x}/\bar{x}]); \tau.$$

8 An Example

We can use our compiler correctness theorem to make statements about linking with arbitrary A components, as long as they have translation type. In this section, we present an example showing how our framework allows linking both with A components that cannot be expressed in F, and with those that can. To keep our example concise, we use variable substitution as a simple notion of linking.

Consider the component

$$e = (\lambda g: \text{unit} \rightarrow \text{int}. (g ()) * (g ())) x,$$

where $\cdot; \cdot; (x: \text{unit} \rightarrow \text{int}) \vdash e: \text{int}$. In F alone, only divergent or constant functions can have type $\text{unit} \rightarrow \text{int}$, but if we are compiling to A before linking, we could be given a component that makes use of A's mutable references.

Putting e through the first compiler pass, we get a C component that contains several administrative reductions. The complete result of compilation is shown in the technical report, but for readability, we pretend that e compiles to

$$e = (\lambda g: \exists \alpha. \langle (\alpha, \text{unit}) \rightarrow \text{int}, \alpha \rangle. (\text{unpack } \langle \beta, z \rangle = g \text{ in } (\pi_1(z) \pi_2(z) ())) \\ * (\text{unpack } \langle \beta, z \rangle = g \text{ in } (\pi_1(z) \pi_2(z) ()))) x,$$

which is equivalent to the actual result of compilation, and has exactly the same function body as the closure produced by the compiler.

The second pass brings us to an A component $\mathbf{e} = (\mathbf{t}, \mathbf{H})$, where $\mathbf{t} = \ell \mathbf{x}$ and $\mathbf{H} = \ell \mapsto \lambda \mathbf{g} : \exists \alpha. \mathbf{box} \langle \mathbf{box} (\alpha, \mathbf{unit}) \rightarrow \mathbf{int}, \alpha \rangle$.

$$\begin{aligned} & ((\mathbf{unpack} \langle \beta, z \rangle = \mathbf{g} \text{ in } ((\mathbf{read}[1] z) (\mathbf{read}[2] z) ())) * \\ & (\mathbf{unpack} \langle \beta, z \rangle = \mathbf{g} \text{ in } ((\mathbf{read}[1] z) (\mathbf{read}[2] z) ())))). \end{aligned}$$

By compiler correctness, we know that

$$;\cdot;\cdot; (\mathbf{x} : \mathbf{unit} \rightarrow \mathbf{int}) \vdash \mathbf{e} \approx^{ctx} \mathbf{int} \mathcal{FCA}(\mathbf{e}[\mathcal{ACF}^{\mathbf{unit} \rightarrow \mathbf{int}} \mathbf{x}/\mathbf{x}]) : \mathbf{int}.$$

Equivalently,

$$;\cdot;\cdot; (\mathbf{x} : \tau) \vdash \mathcal{ACF}^{\mathbf{int}}(\mathbf{e}[\mathbf{unit} \rightarrow \mathbf{int} \mathcal{FCA} \mathbf{x}/\mathbf{x}]) \approx^{ctx} \mathbf{e} : \mathbf{int},$$

where $\tau = \mathbf{unit} \rightarrow \mathbf{int} \langle \mathcal{C} \rangle \langle \mathcal{A} \rangle = \exists \alpha. \mathbf{box} \langle \mathbf{box} (\alpha, \mathbf{unit}) \rightarrow \mathbf{int}, \alpha \rangle$.

Suppose we want to instantiate \mathbf{x} with the following A component, which creates a function that uses a mutable reference to return the number of times it has been called:

$$\mathbf{e}' = (\mathbf{pack} \langle \mathbf{ref} \mathbf{int}, \mathbf{balloc} \langle \ell, \mathbf{ralloc} \langle 0 \rangle \rangle \rangle) \text{ as } \tau,$$

$$\ell \mapsto \lambda (\mathbf{x} : \mathbf{ref} \mathbf{int}, \mathbf{z} : \mathbf{unit}). \text{ let } \mathbf{y} = \mathbf{read}[1] \mathbf{x} \text{ in let } \mathbf{z} = \mathbf{write} \mathbf{x} [1] \leftarrow \mathbf{y} + 1 \text{ in } \mathbf{y} + 1).$$

We would then have

$$;\cdot;\cdot; \vdash \mathcal{ACF}^{\mathbf{int}}(\mathbf{e}[\mathbf{unit} \rightarrow \mathbf{int} \mathcal{FCA} \mathbf{e}'/\mathbf{x}]) \approx^{ctx} \mathbf{e}'[\mathbf{x}] : \mathbf{int},$$

The right-hand side of this equivalence is exactly the pure-A program that we would ultimately run, and the left-hand side is an FCA program that models it. Note that on either side of the equation, the function exported by \mathbf{e}' will be applied to the unit value twice, returning **1** the first time and **2** the second time. An F function could not exhibit this behavior. This demonstrates how our framework allows for linking with components that are not expressible in F.

If we want instead to link with a different A component $\hat{\mathbf{e}}$ that was compiled from an F component $\hat{\mathbf{e}}$, we can still make the statement

$$;\cdot;\cdot; \vdash \mathcal{ACF}^{\mathbf{int}}(\mathbf{e}[\mathbf{unit} \rightarrow \mathbf{int} \mathcal{FCA} \hat{\mathbf{e}}/\mathbf{x}]) \approx^{ctx} \mathbf{e}[\hat{\mathbf{e}}/\mathbf{x}] : \mathbf{int},$$

but we can also simplify this statement using our additional knowledge of $\hat{\mathbf{e}}$. Our compiler correctness theorem tells us that

$$;\cdot;\cdot; \vdash \mathcal{ACF}^{\mathbf{unit} \rightarrow \mathbf{int}} \hat{\mathbf{e}} \approx^{ctx} \hat{\mathbf{e}} : \tau.$$

From this, we can infer that

$$;\cdot;\cdot; \vdash \mathcal{ACF}^{\mathbf{int}}(\mathbf{e}[\mathbf{unit} \rightarrow \mathbf{int} \mathcal{FCA}(\mathcal{ACF}^{\mathbf{unit} \rightarrow \mathbf{int}} \hat{\mathbf{e}})/\mathbf{x}]) \approx^{ctx} \mathbf{e}[\hat{\mathbf{e}}/\mathbf{x}] : \mathbf{int}.$$

Applying boundary cancellation yields

$$;\cdot;\cdot; \vdash \mathcal{ACF}^{\mathbf{int}}(\mathbf{e}[\hat{\mathbf{e}}/\mathbf{x}]) \approx^{ctx} \mathbf{e}[\hat{\mathbf{e}}/\mathbf{x}] : \mathbf{int}.$$

Now we are essentially equating the pure-A program with a pure-F program, since the only multi-language element in this statement is the integer boundary at the outermost level, which merely converts an \mathbf{n} to \mathbf{n} . This demonstrates that when we do have source-language equivalents for all our target-level components, our framework allows us to model target-level linking with source-level linking.

9 Proving Compiler Correctness

To prove the compiler correctness theorem, we design a step-indexed Kripke logical relation as a sound and complete model of contextual equivalence in FCA. Our logical relation extends that of Dreyer *et al.* [15] with the ability to handle multi-language type

abstraction. We give an overview of the logical relation and a more detailed discussion of its novel features in the technical report [9]. In this section, we briefly discuss the high-level ideas behind our model’s novel elements.

A logical-relations model provides a *relational value interpretation* of each type τ . This relation, which we denote $\mathcal{V}[\tau]$, specifies when two values of type τ should be considered related or equivalent. When τ has free type variables, an environment ρ holds *arbitrary relational interpretations* for those abstract types. The relations in ρ capture the invariants of different instantiations of polymorphic values, which allows us to prove parametricity properties.

The interpretation $\mathcal{V}[\alpha]\rho$ is defined by just looking up $\rho(\alpha)$. To prove important properties of $\mathcal{V}[\tau]\rho$ for all types, we must ensure those properties hold in the α case by constraining the relations we can put into ρ to require these properties to hold upfront. Interpretations that satisfy these properties are called *candidates* or *admissible relations*.

In our multi-language setting, the two key properties we need to require for admissibility are boundary cancellation and the *bridge lemma*. The bridge lemma states that, given a pair of values \mathbf{v}_1 and \mathbf{v}_2 related according to the interpretation $\mathcal{V}[\tau]\rho$, the \mathbf{CF}^τ translations of those values must be related according to $\mathcal{V}[\tau^{\langle \mathbf{C} \rangle}]\rho$. Similarly, given values \mathbf{v}_1 and \mathbf{v}_2 related according to $\mathcal{V}[\tau^{\langle \mathbf{C} \rangle}]\rho$, their ${}^\tau\mathbf{FC}$ translations must be related according to $\mathcal{V}[\tau]\rho$. (We also require the analogous properties for the second pass.)

The type translation of α is $\lceil \alpha \rceil$, so in order for the bridge lemma to hold at type α , we need a suitable definition of $\mathcal{V}[\lceil \alpha \rceil]\rho$, which necessarily will depend on $\rho(\alpha)$. One naïve definition we tried is the set of translations of values from $\rho(\alpha)$, roughly:

$$\mathcal{V}[\lceil \alpha \rceil]\rho = \{(\mathbf{v}_1, \mathbf{v}_2) \mid (\mathbf{v}_1, \mathbf{v}_2) \in \rho(\alpha) \wedge \mathbf{CF}(\mathbf{v}_i) = \mathbf{v}_i\}.$$

While this definition does let us prove the bridge lemma at type α , it does not satisfy boundary cancellation: if \mathbf{v}_1 and \mathbf{v}_2 are related according to this definition of $\mathcal{V}[\lceil \alpha \rceil]\rho$, it is not necessarily the case that $\mathbf{CA}(\mathbf{AC}(\mathbf{v}_1))$ and \mathbf{v}_2 are related.

All the ways we tried to define $\mathcal{V}[\lceil \alpha \rceil]\rho$ by a simple formula in terms of $\rho(\alpha)$ failed for similar reasons. Instead of giving a uniform definition, we took the viewpoint that if the properties of $\rho(\alpha)$ must be given *a priori*, then the particular relations with those properties that instantiate $\mathcal{V}[\alpha]\rho$ and $\mathcal{V}[\lceil \alpha \rceil]\rho$ should be given *a priori* as well. Specifically, in our model, an interpretation $\rho(\alpha)$ is not just given by a relation on F values, but by a triple containing the relation on F values, a relation on C values to serve as its “translation” and instantiate $\mathcal{V}[\lceil \alpha \rceil]\rho$, and a relation on A values to instantiate $\mathcal{V}[\alpha]\rho$. Similarly, an interpretation $\rho(\alpha)$ is given by a pair containing a C-level relation and an A-level relation. For $\rho(\alpha)$, since A is the target language, only one relation is needed.

This strategy moves the burden for defining the “translations” of candidate relations to the places in our proof development where individual candidates are needed. But in all these places, there is some specific information available about the relation, so it was not difficult to construct them.

10 Discussion and Future Work

Software is composed from components written in different languages because different languages are suited to different tasks. We have provided a novel methodology for verifying *open, multi-pass* compilers, one that yields a stronger theorem than any existing

work, allowing target-level linking with components of arbitrary provenance regardless of whether the component can be expressed in the source language compiled by the verified compiler.

Adding Compiler Passes. Adding more intermediate languages to our compiler pipeline requires extending the multi-language model with new boundary forms and translation rules, and extending the logical relation with new clauses. Our aim is that the proof structure should be as modular as possible, so that the major lemmas and the correctness proof for one compiler pass can be completed independently of the rest of the pipeline. Presently, since our admissible relations design requires relations from multiple languages, we have a small number of places where a proof about one pass is affected by the other languages and passes. We hope to improve our proof engineering so that proofs for existing passes are unaffected when the compiler pipeline is changed.

Compiling to Assembly. We have extended our compiler with a code-generation pass that translates A components to a stack-based typed assembly language, T. The latter is similar to Morrisett *et al.*'s stack-based TAL [6] but with a type system that tracks more information. Informally, the T type system allows us to track calls and returns of semantic “functions” that may span multiple basic blocks, and to determine the “return type” of such functions. With this information, we are able to give a formal definition of contextual equivalence for T that makes distinctions about assembly at an appropriate level of granularity. That is, we relate assembly language components comprised of any number of basic blocks, rather than relating individual basic blocks. An equivalence relation based on individual blocks would be too fine grained; for instance, it would be unable to relate two components with an unequal number of basic blocks that may have been produced by compiling two equivalent source terms. We are working on the proofs for this pass and will report on it in a future paper.

Mutable References. Consider adding mutable references to F and C. For the first compiler pass, we would extend the type translation with $(\text{ref } \tau)^C = \text{ref } \tau^C$. When defining interoperability at type $\text{ref } \tau$, it doesn't make sense to convert an F location ℓ into a fresh C location ℓ (and vice versa) since it would lead to duplication of mutable cells in the interoperating languages and these would be impossible to keep in sync. One solution is to treat a wrapped location (e.g., ${}^{\text{ref } \tau} \mathcal{F}C\ell$) as a value form. Operations on these wrapped locations can be performed by reduction rules such as these:

$$!(^{\text{ref } \tau} \mathcal{F}C\ell) \mapsto {}^{\tau} \mathcal{F}C(!\ell) \quad (^{\text{ref } \tau} \mathcal{F}C\ell) := v \mapsto \text{unit } \mathcal{F}C(\ell := C\mathcal{F}^{\tau}v),$$

where $!v$ is a dereference and $v := v'$ is an assignment. Passing references between C and A can be done analogously. While these interoperability semantics are straightforward, we expect to find nontrivial challenges in designing a logical relation to properly handle the wrapped-location value forms they introduce.

Supporting Realistic Interoperability. We are particularly interested in supporting target-level interoperability between a language with parametric polymorphism such as ML and languages without type abstraction such as Scheme or C. For instance, given a generic tree library compiled from ML, we want to allow code compiled from Scheme or C to be able to use the library but ensure that such use cannot invalidate ML's parametricity guarantees by inspecting values that have abstract type on the ML side. In this paper, we have shown how to preserve ML's parametricity guarantees part-way through the compiler. Going forward we wish to develop a gradually typed assembly language

that, following Matthews and Ahmed [16], uses dynamic sealing on the untyped side to enforce parametricity guarantees provided by type abstraction on the typed side.

11 Related Work

The literature on compiler verification spans over four decades but is mostly limited to whole-program compilation; we refer the reader to the bibliography by Dave [17] for compilers for first-order languages, and to Chlipala [18] for compilers for higher-order functional languages. We have already discussed the existing work [10,12] on compositional compiler correctness in §2. Here we focus on other closely related work.

Dreyer et al. have recently been working on Relational Transition Systems (RTS's) [19] that may provide an alternative cross-language specification technique that is designed to make it possible to prove transitivity. Regardless, it is still not easy to do: see their technical report [20] where they prove transitivity for their *single-language* RTS system for an idealized ML. It is a non-trivial task to do this for multiple cross-language RTS's. Additionally, even if the RTS approach proves effective for verifying a multi-pass compiler, it still does not address the problem of linking with a component e'_T for which there is no related source-level e'_S .

The design of our multi-language system builds on that of Ahmed and Blume [8], who developed a boundary-based multi-language system embedding the source (STLC) and target (System F) of CPS translation. Ahmed and Blume did not have type abstraction in the source language, which meant that they did not have to make use of lumps or suspensions, nor design a logical relation to handle these. Our semantics preservation proof is analogous to theirs. However, since they were interested in fully abstract CPS translation, they designed their type translation to disallow linking compiled code with target components whose behavior cannot be expressed at the source level. The additional work that they do to prove full abstraction provides a roadmap for how to extend our methodology to prove full abstraction in a setting where the type translation enforces it.

Acknowledgements. We would like to thank Nick Benton, whose views on compositional compiler correctness have been an inspiration to us. In particular, our thinking has been influenced by Benton and Hur's introduction [11], which eloquently lays out desirable features of a compiler correctness specification. We would also like to thank Aaron Turon for helpful feedback on an earlier version of this paper. This research was supported by the National Science Foundation (grant CCF-1203008).

References

1. Leroy, X.: Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In: POPL (2006)
2. Leroy, X.: A formally verified compiler back-end. J. Automated Reasoning 43(4), 363–446 (2009)
3. Lochbihler, A.: Verifying a compiler for Java threads. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 427–447. Springer, Heidelberg (2010)

4. Myreen, M.O.: Verified just-in-time compiler on x86. In: POPL 2010 (2010)
5. Sevcik, J., Vafeiadis, V., Nardelli, F.Z., Jagannathan, S., Sewell, P.: Relaxed-memory concurrency and verified compilation. In: POPL 2011 (2011)
6. Morrisett, G., Crary, K., Glew, N., Walker, D.: Stack-based typed assembly language. *J. Functional Programming* 12(1), 43–88 (2002)
7. Matthews, J., Findler, R.B.: Operational semantics for multi-language programs. In: POPL 2007 (2007)
8. Ahmed, A., Blume, M.: An equivalence-preserving CPS translation via multi-language semantics. In: ICFP 2011 (2011)
9. Perconti, J.T., Ahmed, A.: Verifying an open compiler using multi-language semantics (technical report) (January 2014), <http://ccs.neu.edu/home/amal/voc/>
10. Benton, N., Hur, C.K.: Biorthogonality, step-indexing and compiler correctness. In: ICFP 2009 (2009)
11. Benton, N., Hur, C.K.: Realizability and compositional compiler correctness for a polymorphic language. Technical Report MSR-TR-2010-62, Microsoft Research (April 2010)
12. Hur, C.K., Dreyer, D.: A Kripke logical relation between ML and assembly. In: POPL 2011 (2011)
13. Morrisett, G., Walker, D., Crary, K., Glew, N.: From System F to typed assembly language. *ACM TOPLAS* 21(3), 527–568 (1999)
14. Minamide, Y., Morrisett, G., Harper, R.: Typed closure conversion. In: POPL 1996 (1996)
15. Dreyer, D., Neis, G., Birkedal, L.: The impact of higher-order state and control effects on local relational reasoning. *J. Functional Programming* 22(4&5), 477–528 (2012)
16. Matthews, J., Ahmed, A.: Parametric polymorphism through run-time sealing, or, theorems for low, low prices! In: Drossopoulou, S. (ed.) ESOP 2008. LNCS, vol. 4960, pp. 16–31. Springer, Heidelberg (2008)
17. Dave, M.A.: Compiler verification: A bibliography. *ACM SIGSOFT Software Engineering Notes* 28(6) (2003)
18. Chlipala, A.: A verified compiler for an impure functional language. In: POPL 2010 (2010)
19. Hur, C.K., Dreyer, D., Neis, G., Vafeiadis, V.: The marriage of bisimulations and Kripke logical relations. In: POPL 2012 (2012)
20. Hur, C.K., Dreyer, D., Neis, G., Vafeiadis, V.: The marriage of bisimulations and Kripke logical relations. Technical report, Max Planck Institute for Software Systems (January 2012)