

# Semantic Soundness for Language Interoperability

Daniel Patterson  
Northeastern University  
Boston, MA, USA  
dbp@dbpmail.net

Andrew Wagner  
Northeastern University  
Boston, MA, USA  
ahwagner@ccs.neu.edu

Noble Mushtak  
Northeastern University  
Boston, MA, USA  
mushtak.n@northeastern.edu

Amal Ahmed  
Northeastern University  
Boston, MA, USA  
amal@ccs.neu.edu

## Abstract

Programs are rarely implemented in a single language, and thus questions of type soundness should address not only the semantics of a single language, but how it interacts with others. Even between type-safe languages, disparate features can frustrate interoperability, as invariants from one language can easily be violated in the other. In their seminal 2007 paper, Matthews and Findler proposed a multi-language construction that augments the interoperating languages with a pair of *boundaries* that allow code from one language to be embedded in the other. While this technique has been widely applied, their syntactic source-level interoperability doesn't reflect practical implementations, where the behavior of interaction is only defined after compilation to a common target, and any safety must be ensured by target invariants or inserted target-level "glue code."

In this paper, we present a novel framework for the design and verification of sound language interoperability that follows an interoperation-after-compilation strategy. Language designers specify what data can be converted between types of the two languages via a convertibility relation  $\tau_A \sim \tau_B$  (" $\tau_A$  is convertible to  $\tau_B$ ") and specify target-level glue code implementing the conversions. Then, by giving a semantic model of source-language types as sets of target-language terms, they can establish not only the meaning of the source types, but also *soundness of conversions*: i.e., whenever  $\tau_A \sim \tau_B$ , the corresponding pair of conversions (glue code) convert target terms that behave like  $\tau_A$  to target terms that behave like  $\tau_B$ , and vice versa. With this, they can prove semantic type soundness for the entire system. We illustrate our framework

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
PLDI '22, June 13–17, 2022, San Diego, CA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9265-5/22/06...\$15.00

<https://doi.org/10.1145/3519939.3523703>

via a series of case studies that demonstrate how our semantic interoperation-after-compilation approach allows us both to account for complex differences in language semantics and make efficiency trade-offs based on particularities of compilers or targets.

**CCS Concepts:** • Software and its engineering → General programming languages.

**Keywords:** language interoperability, type soundness, semantics, logical relations

## ACM Reference Format:

Daniel Patterson, Noble Mushtak, Andrew Wagner, and Amal Ahmed. 2022. Semantic Soundness for Language Interoperability. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '22)*, June 13–17, 2022, San Diego, CA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3519939.3523703>

## 1 Introduction

All practical language implementations come with some way of interoperating with code written in a different language, usually via a foreign-function interface (FFI). This enables development of software systems with components written in different languages, whether to support legacy libraries or different programming paradigms. For instance, you might have a system with a high-performance data layer written in Rust interoperating with business logic implemented in OCaml. Sometimes, this interoperability is realized by targeting a common platform (e.g., Scala [40] and Clojure [23] for the JVM, or SML [10] and F# [48] for .NET). Other times, it is supported by libraries that insert boilerplate or "glue code" to mediate between the two languages (such as the binding generator SWIG [7], C->Haskell [16], OCaml-ctypes [54], NLFFI [13], Rust's bindgen [55], etc). While interoperability can be achieved in other ways—via the network, inter-process communication, or dispatching between interpreters and compiled code—we focus in this paper on the case when both languages are compiled to a shared intermediate or target language.

In 2007, Matthews and Findler [33] observed that while there were numerous FFIs that supported interoperation

between languages, there had been no effort to study the *semantics* of interoperability. They proposed a simple and elegant system for abstractly modeling interactions between languages  $A$  and  $B$  by embedding the existing operational syntax and semantics into a multi-language  $AB$  and adding boundaries to mediate between the two. Specifically, a boundary  ${}^{\tau_A}\mathcal{A}\mathcal{B}^{\tau_B}(\cdot)$  allows a term  $e_B$  of type  $\tau_B$  to be embedded in an  $A$  context that expects a term of type  $\tau_A$ , and likewise for the boundary  ${}^{\tau_B}\mathcal{B}\mathcal{A}^{\tau_A}(\cdot)$ . Operationally, the term  ${}^{\tau_A}\mathcal{A}\mathcal{B}^{\tau_B}(e_B)$  evaluates  $e_B$  using the  $B$ -language semantics to  ${}^{\tau_A}\mathcal{A}\mathcal{B}^{\tau_B}(v_B)$  and then a type-directed conversion takes the value  $v_B$  of type  $\tau_B$  to an  $A$ -language term of type  $\tau_A$ . There are often interesting design choices in deciding what conversions are available for a type, if any at all. One can then prove that the entire multi-language type system is sound by proving type safety for the multi-language, which includes the typing rules of both the embedded languages and the boundaries. This multi-language framework has inspired a significant amount of work on interoperability: between simple and dependently typed languages [41], between languages with unrestricted and substructural types [45, 50], between a high-level functional language and assembly [43], and between source and target languages of compilers [2, 37, 44].

Unfortunately, while Matthews-Findler-style boundaries give an elegant, abstract model for interoperability, they do not reflect reality. Indeed, a decade and a half later, there is little progress on assigning semantics to real multi-language systems. In the actual implementations we study, the source languages are compiled to components in a common target and glue code is inserted at the boundaries between them to account for different data representations or calling conventions. While one could try to approach this problem by defining source-level boundaries, building a compiler for the multi-language, and then showing that the entire system is realized correctly, there are serious downsides to this approach. One is that if the two languages differ significantly, the multi-language may be significantly more than just an embedding of the evaluation rules of both languages (c.f. our last case study, as an implicitly garbage-collected language interoperating with a manually managed language may need to make the garbage collection explicit). And that doesn't even consider the fact that in practice, we usually have *existing* compiler implementations for one or both languages and wish to add (or extend) support for interoperability. Here, language designers' understanding of what datatypes *should* be convertible at the source level very much depends on how the sources are compiled and how data is (or could be) represented in the target, all information that is ignored by the multi-language approach. Moreover, certain conversions, even if possible, might be undesirable because the glue code needed to realize *safe* interoperability imposes too much runtime overhead.

In this paper, we present a framework for the *design* and *verification* of sound language interoperability, where both

activities are connected to the actual implementation (of compilers and conversions). At the source, we still use Matthews-Findler-style boundaries, as our approach differs not in the source syntax but rather that instead of proving operational properties of that source, we instead prove semantic type soundness by defining a model of *source* types as sets of (or relations on) *target* terms. That is, the interpretation of a source type is the set of target terms that *behave as that type*. Guiding the design of these type interpretations are the compilers. This kind of model, often called a realizability model, is not a new idea — for instance, Benton and Zarfaty [12] and Benton and Tabareau [11] used such models to prove type soundness, but their work was limited to a single source language. By interpreting the types of two source languages as sets of terms in a common target, we enable rich reasoning about interoperability. Using the model, we can then give meaning to a boundary  ${}^{\tau_B}\mathcal{B}\mathcal{A}^{\tau_A}(\cdot)$ : there is a bit of target code that, when given a target term that is in the model of the type  $\tau_A$ , results in a target term in the model of type  $\tau_B$ .

A realizability model is valuable not only for proving soundness, but for reasoning about the *design* of interoperability. For example, we can ask if a particular type in one language is *the same* as a type in the other language. This is true if the same set of target terms inhabits both types, and in this case conversions between the types should do nothing. More generally, opportunities for efficient conversions may only become apparent upon looking at how source types and invariants are represented (or realized) in the target. Since interoperability is a design challenge, with tradeoffs just like any other—performance high among them—working with the ability to understand all the pieces is a tremendous advantage.

**Contributions.** To demonstrate the use and benefits of our framework, we present three case studies that illustrate different kinds of challenges for interoperability. In each case, we compile to an untyped target language.

1. **Shared-Memory Interoperability (§3):** We consider how mutable references can be exchanged between two languages and what properties must hold of stored data for aliasing to be safe. We show that to avoid copying mutable data — without having to wrap references in guards or chaperones [47] — convertible reference types must be inhabited by the *very same* set of target terms.

2. **Affine & Unrestricted (§4):** We consider how `MiniML`, a standard functional language with mutable references, can interact with `AFFI`, an affine language. We show that affine code can be safely embedded in unrestricted code and vice versa by using runtime checks (only where necessary) to ensure that affine resources are used at most once.

3. **Memory Management & Polymorphism (§5):** We consider how `MiniML`, whose references are garbage collected, can interact with `L3` [3], a language that uses linear

capabilities to support safe strong updates to a manually managed mutable heap. We demonstrate not only when memory can be moved between languages, but also a type-level form of interoperability that allows generics to be used with  $\mathbf{L}^3$  (which lacks type polymorphism) without violating any invariants of either language.

For each case study, we devise a novel realizability model. An interesting aspect of these models is that, since the target languages are untyped, statically enforced source invariants must be captured using either dynamic enforcement in target code or via invariants in the model. This demonstrates that our approach is viable even when working with existing target languages without rich static reasoning principles.

We chose these three case studies to exercise our framework both in terms of type system invariants (substructural types, polymorphism) but also properly handling mutable state (exchanging pointers and garbage collection). Interesting challenges for the future could include differences of control-flow and concurrency.

*Definitions and proofs elided from this paper are provided in our technical appendix [42].*

## 2 The Framework

Before diving into the case studies that serve as evidence of its efficacy, we first describe, in step-by-step fashion, the framework for proving type soundness in the presence of interoperability that is the primary contribution of this paper. The inputs to the framework are two source languages, language  $\mathbf{A}$  and language  $\mathbf{B}$ , a target language  $\mathbf{T}$ , and compilers  $e^+ = e$  and  $e^+ = e$ . This section serves both as a roadmap of what is to come and a reference to return to. The first two steps (§2.1 and §2.2) must be performed by the **designer** of the interoperability system, whereas the last three (§2.3, §2.4, and §2.5) should be performed by the **verifier** of the system. Note that, as with type soundness, partial verification is still potentially useful, and so the first two steps should be seen as what needs to be implemented, and the last three as what should be aspired to, if not formally carried out.

### 2.1 Boundary Syntax

To include code from another language, the designer requires some way of invoking such code. While there are various ways of doing this in real toolchains, here she adopts a general approach based on a notion of *language boundaries*.

If a language  $\mathbf{A}$  is to include code from language  $\mathbf{B}$ , the  $\mathbf{A}$  designer should add a boundary form  $(e)_{\tau_A}$ . This allows a term  $e : \tau_B$  to be used in an  $\mathbf{A}$  context at type  $\tau_A$ , for some  $\tau_A$  and  $\tau_B$ . This boundary strategy is very general: it allows both inline code, a strategy adopted by many FFI libraries for C, but also the more typical import/export style of linking. In that case, what is compiled would be an open term with a  $\mathbf{B}$  binding  $f : \tau \rightarrow \tau'$  free. Then, the use of the imported term would be  $(f)_{\tau_A \rightarrow \tau'_A}$  for appropriate types  $\tau_A$  and  $\tau'_A$ .

Note that while in our examples, we equip both languages with boundaries, the framework does not require this.

### 2.2 Convertibility Rules

To know whether a term  $(e)_{\tau_A}$  is well-typed, the designer needs to know if a  $\mathbf{B}$  term  $e : \tau_B$  can be converted to an  $\mathbf{A}$  type  $\tau_A$ . There is no way to know, a priori, what types can be converted, and thus the framework requires that the designer specify this explicitly. In particular, she must provide judgments of the form  $\tau_A \sim \tau_B$  to indicate that these two types are interconvertible, allowing for the possibility of dynamic conversion errors. Since our notion of linking depends upon both language  $\mathbf{A}$  and  $\mathbf{B}$  being compiled to a common target  $\mathbf{T}$ , this conversion needs to be witnessed by  $\mathbf{T}$  code that performs the conversion.  $C_{\tau_A \mapsto \tau_B}$  denotes the code that performs a target-level conversion from  $\tau_A$  to  $\tau_B$ . For example, if  $\mathbf{bool} \sim \mathbf{int}$ , and the former compiles to the integers 0 and 1, then the conversion  $C_{\mathbf{bool} \mapsto \mathbf{int}}$  is a no-op (since compiled booleans are already  $\mathbf{T}$  language integers), but  $C_{\mathbf{int} \mapsto \mathbf{bool}}$  must do something different. It could raise a dynamic conversion error if given a  $\mathbf{T}$  int other than 0 or 1, or it could collapse all other numbers into one of those, or something else. The particular choice depends on the languages in question, and what the designer of the interoperability system thinks makes sense: the framework only requires that the decision made preserves type soundness.

### 2.3 Realizability Models for Both Languages

In order to prove type soundness, and in particular, account for the boundaries and convertibility rules from §2.1 and §2.2, the verifier needs to build a logical relation for both languages. This relation is atypical in two ways. First, it is a *realizability* model, which means that while it is indexed by source types, it is inhabited by target terms. That is, the verifier must first define an interpretation of values for each source type  $\tau$ , written  $\mathcal{V}[\tau]$ , as the set of  $\mathbf{T}$  language values  $v$  that behave as  $\tau$ . That is,  $\mathcal{V}[\mathbf{bool}]$  is not the set of  $\mathbf{A}$  language booleans (i.e., `true` and `false`), but rather, the  $\mathbf{T}$  values that behave as  $\mathbf{A}$  booleans (likely, 0 and 1). In particular, the compiler from  $\mathbf{A}$  to  $\mathbf{T}$  must send `true` and `false` into  $\mathcal{V}[\mathbf{bool}]$ , but the latter can include more values. There is also an expression relation, written  $\mathcal{E}[\tau_A]$ , that is the set of  $\mathbf{T}$  language terms that evaluate to values in  $\mathcal{V}[\tau_A]$  (or diverge, or run to a well-defined error). The second atypical, and novel, aspect is that the relation is indexed with the types of *both* of our source languages; in this example,  $\mathbf{A}$  and  $\mathbf{B}$ . Since they compile to the same target, this works: the inhabitants of  $\mathcal{V}[\mathbf{bool}]$  and  $\mathcal{V}[\mathbf{int}]$  are both  $\mathbf{T}$  values. By bringing the types of both languages into a common setting, the verifier gains powerful reasoning principles; for example, we can ask if  $\mathcal{V}[\mathbf{bool}] = \mathcal{V}[\mathbf{int}]$ .

## 2.4 Soundness of Conversions

Using the realizability models defined in §2.3, the verifier can prove that the convertibility rules defined in §2.2 are sound. In particular, if  $\tau_A \sim \tau_B$ , then she should show that the conversions  $C_{\tau_A \mapsto \tau_B}$  and  $C_{\tau_B \mapsto \tau_A}$  actually translate expressions between the types correctly. This is done by showing for any term  $e$  in  $\mathcal{E}[\tau_A]$ , that  $C_{\tau_A \mapsto \tau_B}(e)$  is in  $\mathcal{E}[\tau_B]$ , and similarly for  $C_{\tau_B \mapsto \tau_A}$ . Since the model defines type interpretations, this ensures that the conversions do exactly what is expected.

## 2.5 Soundness of Entire Languages

Proving the conversions sound (§2.4) is the central goal, of course, but the verifier also needs to ensure that the model defined in §2.3 is actually faithful to the languages. She does this by following the standard approach for proving semantic type soundness. First, for each typing rule in both source languages, she proves that a corresponding lemma holds in terms of the model. For example, for pairs she proves that if  $e \in \mathcal{E}[\tau_1 \times \tau_2]$  then  $\text{fst}^+ e \in \mathcal{E}[\tau_1]$ —note we write  $\text{fst}^+$ , which is T code (and could be an array projection, or some other T operation), since what is in  $\mathcal{E}[\tau_1]$  are T terms.

## 3 Shared Memory

Aliased mutable data is challenging to deal with no matter the context, but aliasing across languages is especially difficult because giving a pointer to a foreign language can allow for *unknown* data to be written to its address. Specifically, if the pointer has a particular type in the host language, then only certain data should be written to it, but the foreign language may not respect or even know about that restriction. One existing approach to this problem is to create proxies, where data is guarded or converted before being read or written [17, 32, 47]. While sound, this comes with significant runtime overhead. Here, our framework allows a different approach.

**Languages.** In this case study, we explore this problem using two simply-typed functional source languages with dynamically allocated mutable references, **RefHL** and **RefLL** (for “higher-level” and “lower-level”). **RefHL** has boolean, sum, and product types, whereas **RefLL** has arrays ( $[e_1, \dots, e_n]$ :  $[\tau]$ ). Their syntax is given in Fig. 1 and their static semantics — which are entirely standard — are elided (see [42]). These two languages are compiled (Fig. 3—note that we write  $e^+$  to indicate  $e'$ , where  $e \rightsquigarrow e'$ ) into an untyped stack-based language called **StackLang** (inspired by [29]), whose syntax and small-step operational semantics — a relation on configurations  $\langle H; S; P \rangle$  comprised of a heap, stack, and program — are given in Fig. 2; here we describe a few highlights. First, we note that **StackLang** values include not only numbers, thunks, and locations, but arrays of values, a simplification we made for the sake of presentation. Second, notice the interplay between **thunk** and **lam**: thunks are suspended computations, whereas **lam** is an instruction (not a value)

```

RefHL  Type  $\tau$  ::= unit | bool |  $\tau + \tau$  |  $\tau \times \tau$  |  $\tau \rightarrow \tau$  | ref  $\tau$ 
Expr  $e$  ::= () | true | false | x | inl e | inr e
        | (e, e) | fst e | snd e | if e e e |  $\lambda x : \tau. e$  | e e
        | match e x{e} y{e} | ref e | !e | e := e | (e) $_{\tau}$ 

RefLL  Type  $\tau$  ::= int |  $[\tau]$  |  $\tau \rightarrow \tau$  | ref  $\tau$ 
Expr  $e$  ::= n | x | [e, ...] | e[e] |  $\lambda x : \tau. e$  | e e
        | e + e | if0 e e e | ref e | !e | e := e | (e) $_{\tau}$ 

```

Figure 1. Syntax for **RefHL** and **RefLL**.

```

Program P ::=  $\cdot$  | i, P   Value v ::= n | thunk P |  $\ell$  | [v, ...]
Instruction i ::= push v | add | less? | if0 P P | lam x.P | call
                | idx | len | alloc | read | write | fail c
Error Code c ::= TYPE | IDX | CONV
Heap H ::= { $\ell$ :v, ...}   Stack S ::= v, ..., v | Fail c

⟨H; S; push v, P⟩           → ⟨H; S, v; P⟩           (S ≠ Fail c)
⟨H; S, n', n; add, P⟩       → ⟨H; S, (n + n'); P⟩
⟨H; S, n', n; less?, P⟩     → ⟨H; S, b; P⟩       (b=0 if n < n' else 1)
⟨H; S, n; if0 P1 P2, P⟩   → ⟨H; S, Pi; P⟩     (i=1 if n=0 else 2)
⟨H; S; if0 P1 P2, P⟩     → ⟨H; S; fail TYPE⟩   (S ≠ S', n)
⟨H; S, v; lam x.P1, P2⟩ → ⟨H; S; [x ↦ v]P1, P2⟩
⟨H; S, thunk P1; call, P2⟩ → ⟨H; S; P1, P2⟩
⟨H; S, [v0, ..., vn'], n; idx, P⟩ → ⟨H; S, vn; P⟩       (n ∈ [0, n'])
⟨H; S, [v0, ..., vn'], n; idx, P⟩ → ⟨H; S; fail IDX⟩     (n ∉ [0, n'])
⟨H; S, [v0, ..., vn]; len, P⟩ → ⟨H; S, (n + 1); P⟩
⟨H; S, v; alloc, P⟩         → ⟨H ⊔ { $\ell$ :v}; S,  $\ell$ ; P⟩
⟨H ⊔ { $\ell$ :v}; S,  $\ell$ ; read, P⟩ → ⟨H ⊔ { $\ell$ :v}; S, v; P⟩
⟨H ⊔ { $\ell$ :_}; S,  $\ell$ , v; write, P⟩ → ⟨H ⊔ { $\ell$ :v}; S; P⟩
⟨H; S; fail c, P⟩          → ⟨H; Fail c;  $\cdot$ ⟩

```

Figure 2. Syntax and selected operational semantics for **StackLang** (most fail **TYPE** cases elided).

responsible solely for substitution<sup>1</sup>. We can see how these features are combined, or used separately, in our compilers (Fig. 3). Finally, note that for any instruction where the precondition on the stack is not met, the configuration steps to a program with fail **TYPE** (a dynamic type error), although we elide most of these reduction rules for space.

**Convertibility.** In our source languages, we may syntactically embed a term from one language into the other using the boundary forms  $(e)_{\tau_A}$  and  $(e)_{\tau_B}$ . The typing rules for boundary terms require that the boundary types be convertible, written  $\tau_A \sim \tau_B$ . Those typing rules are:

$$\frac{\Gamma; \Gamma \vdash e : \tau_A \quad \tau_A \sim \tau_B}{\Gamma; \Gamma \vdash (e)_{\tau_B} : \tau_B} \quad \frac{\Gamma; \Gamma \vdash e : \tau_B \quad \tau_A \sim \tau_B}{\Gamma; \Gamma \vdash (e)_{\tau_A} : \tau_A}$$

Note that the convertibility judgment is a declarative, extensible judgment that describes closed types in one language that are interconvertible with closed types in the other, allowing for the possibility of well-defined runtime errors. By separating this judgment from the rest of the type system,

<sup>1</sup>À la Levy’s Call-by-push-value [31].

$\text{SWAP} \triangleq \text{lam } x.(\text{lam } y.\text{push } x, \text{push } y)$	
$\text{DROP} \triangleq \text{lam } x.()$	$\text{DUP} \triangleq \text{lam } x.(\text{push } x, \text{push } x)$
$() \rightsquigarrow \text{push } 0$	$x \rightsquigarrow \text{push } x$
$\text{true} \mid \text{false} \rightsquigarrow \text{push } \langle 0 \mid 1 \rangle$	
$\text{inl } e \mid \text{inr } e \rightsquigarrow e^+, \text{lam } x.(\text{push } \langle 0 \mid 1 \rangle, x)$	
$\text{if } e_1 \ e_2 \rightsquigarrow e^+, \text{if0 } e_1^+ \ e_2^+$	
$\text{match } e \rightsquigarrow e^+, \text{DUP}, \text{push } 1, \text{idx}, \text{SWAP}, \text{push } 0,$	
$\quad x\{e_1\} \ y\{e_2\} \rightsquigarrow \text{idx}, \text{if0 } (\text{lam } x.e_1^+) \ (\text{lam } y.e_2^+)$	
$(e_1, e_2) \rightsquigarrow e_1^+, e_2^+, \text{lam } x_2, x_1.(\text{push } [x_1, x_2])$	
$\text{fst } e \mid \text{snd } e \rightsquigarrow e^+, \text{push } \langle 0 \mid 1 \rangle, \text{idx}$	
$e_1 \ e_2 \rightsquigarrow e_1^+, e_2^+, \text{SWAP}, \text{call}$	
$\text{ref } e \rightsquigarrow e^+, \text{alloc}$	
$e_1 := e_2 \rightsquigarrow e_1^+, e_2^+, \text{write}, \text{push } 0$	
$(e)_\tau \rightsquigarrow e^+, C_{\tau \mapsto \tau}$	
$n \rightsquigarrow \text{push } n$	$e_1 + e_2 \rightsquigarrow e_1^+, e_2^+, \text{SWAP}, \text{add}$
$[e_1, \dots, e_n] \rightsquigarrow e_1^+, \dots, e_n^+, \text{lam } x_n, \dots, x_1.$	
	$\quad (\text{push } [x_1, \dots, x_n])$
$e_1[e_2] \rightsquigarrow e_1^+, e_2^+, \text{idx}$	
$\lambda x : \tau. e \rightsquigarrow \text{push } (\text{thunk } \text{lam } x.e^+)$	
$!e \rightsquigarrow e^+, \text{read}$	
$(e)_\tau \rightsquigarrow e^+, C_{\tau \mapsto \tau}$	

**Figure 3.** Selections from compilers for **RefHL** and **RefLL**.

the language designer can allow additional conversions to be added later, whether by implementers or even end-users. The second thing to note is that this presentation allows for open terms to be converted, so we must maintain a type environment for both languages during typechecking (both  $\Gamma$  and  $\Gamma$ ), as we have to carry information from the site of binding—possibly through conversion boundaries—to the site of variable use. A simpler system, which we have explored, would only allow closed terms to be converted. In that case, the typing rules still use the  $\tau_A \sim \tau_B$  judgment but do not thread foreign environments (using only  $\Gamma$  for **RefHL** and only  $\Gamma$  for **RefLL**).

We present, in Fig. 4, some of the convertibility rules we have defined for this case study (we elide  $\tau_1 \times \tau_2 \sim [\tau]$ ), which come with target-language instruction sequences that perform the conversions, written  $C_{\tau_A \mapsto \tau_B}$  (some are no-ops). An instruction sequence  $C_{\tau_A \mapsto \tau_B}$ , while ordinary target code, when appended to a program in the model at type  $\tau_A$ , should result in a program in the model at type  $\tau_B$ . An implementer can write these conversions based on understanding of the sets of target terms that inhabit each source type, before defining a proper semantic model (or possibly, without defining one, if formal soundness is not required). They would do this based on inspection of the compiler and the target.

From Fig. 3, we see that **bool** and **int** both compile to target integers, and importantly, that **if** compiles to **if0**, which means the compiler interprets **false** as any non-zero integer. Hence, conversions between **bool** and **int** are identities.

For sums, we use the tags 0 and 1, and as for **if**, we use **if0** to branch in the compilation of **match**. Therefore, we can

$\frac{}{C_{\text{bool} \mapsto \text{int}}, C_{\text{int} \mapsto \text{bool}} : \text{bool} \sim \text{int}}$
$\frac{}{C_{\text{ref bool} \mapsto \text{ref int}}, C_{\text{ref int} \mapsto \text{ref bool}} : \text{ref bool} \sim \text{ref int}}$
$\frac{C_{\tau_1 \mapsto \text{int}}, C_{\text{int} \mapsto \tau_1} : \tau_1 \sim \text{int} \quad C_{\tau_2 \mapsto \text{int}}, C_{\text{int} \mapsto \tau_2} : \tau_2 \sim \text{int}}{C_{\tau_1 + \tau_2 \mapsto [\text{int}]}, C_{[\text{int}] \mapsto \tau_1 + \tau_2} : \tau_1 + \tau_2 \sim [\text{int}]}$
$C_{\text{bool} \mapsto \text{int}} \triangleq C_{\text{int} \mapsto \text{bool}} \triangleq C_{\text{ref bool} \mapsto \text{ref int}} \triangleq C_{\text{ref int} \mapsto \text{ref bool}} \triangleq \cdot$
$C_{\tau_1 + \tau_2 \mapsto [\text{int}]} \triangleq \text{DUP}, \text{push } 1, \text{idx}, \text{SWAP}, \text{push } 0, \text{idx}, \text{DUP},$
$\quad \text{if0 } (\text{SWAP}, C_{\tau_1 \mapsto \text{int}})$
$\quad (\text{SWAP}, C_{\tau_2 \mapsto \text{int}}), \text{lam } x_v. \text{lam } x_t. \text{push } [x_t, x_v]$
$C_{[\text{int}] \mapsto \tau_1 + \tau_2} \triangleq$
$\quad \text{DUP}, \text{len}, \text{push } 2, \text{SWAP}, \text{less?}, \text{if0 fail CONV},$
$\quad \text{DUP}, \text{push } 1, \text{idx}, \text{SWAP}, \text{push } 0, \text{idx}, \text{DUP},$
$\quad \text{if0 } (\text{SWAP}, C_{\text{int} \mapsto \tau_1})(\text{DUP}, \text{push } -1, \text{add},$
$\quad \text{if0 } (\text{SWAP}, C_{\text{int} \mapsto \tau_2})(\text{fail CONV}), \text{lam } x_v. \text{lam } x_t. \text{push } [x_t, x_v]$

**Figure 4.** Conversions for **RefHL** and **RefLL**.

choose if the **inl** and **inr** tags should be represented by 0 and 1, or by 0 and any other integer  $n$ . Given that tags could be added later, we choose the former, thus converting a sum to an array of integers is mostly a matter of converting the payload. In the other direction, we have to handle the case that the array is too short, and error.

The final case, between **ref bool** and **ref int**, is the reason for this case study. Intuitively, if you exchange pointers, any value of the new type can now be written at that address, and thus must have been compatible with the old type (as aliases could still exist). Thus, we require that **bool** and **int** are somehow “identical” in the target, so conversions are unnecessary.

**Semantic Model.** Declaring that a type **bool** is “identical” to **int** or that  $\tau$  is convertible to  $\tau$  and providing the conversion code is not sufficient for soundness. In order to show that these conversions are sound, and indeed to understand which conversions are even possible, we define a model for source types that is inhabited by target terms. Since both languages compile to the same target, the range of their relations will be the same (i.e., composed of terms and values from StackLang), and thus we will be able to easily and directly compare the inhabitants of two types, one from each language.

Our model, which aside from the use of StackLang is a standard step-indexed unary logical relation for a language with mutable state (essentially following Ahmed [4]), is presented with some parts elided in Fig. 5 (see [42]).

We give value interpretations for each source type  $\tau$ , written  $\mathcal{V}[\tau]$  as sets of target *values*  $v$  paired with *worlds*  $W$  that inhabit that type. A world  $W$  is comprised of a step index  $k$  and a *heap typing*  $\Psi$ , which maps locations to type interpretations in *Typ*. As is standard, *Typ* is the set of valid type interpretations, which must be closed under world extension.

A future world  $W'$  extends  $W$ , written  $W' \sqsupseteq W$ , if  $W'$  has a potentially lower step budget  $j \leq W.k$  and all locations in  $W.\Psi$  still have the same types (to approximation  $j$ ).

Intuitively,  $(W, v) \in \mathcal{V}[\tau]$  says that the target value  $v$  belongs to (or behaves like a value of) type  $\tau$  in world  $W$ . For example,  $\mathcal{V}[\text{unit}]$  is inhabited by 0 in any world. A more interesting case is  $\mathcal{V}[\text{bool}]$ , which is the set of all target integers, not just 0 and 1, though we could choose to define our model that way (provided we compiled `bool`s to 0 or 1). An array  $\mathcal{V}[\tau]$  is inhabited by an array of target values  $v_i$  in world  $W$  if each  $v_i$  is in  $\mathcal{V}[\tau]$  with  $W$ .

Functions follow the standard pattern for logical relations, appropriately adjusted for our stack-based target language:  $\mathcal{V}[\tau_1 \rightarrow \tau_2]$  is inhabited by values `thunk lam x.P` in world  $W$  if, for any future world  $W'$  and argument  $v$  in  $\mathcal{V}[\tau_1]$  at that world, the result of substituting the argument into the body ( $[x \mapsto v]P$ ) is in the expression relation at the result type  $\mathcal{E}[\tau_2]$ . Reference types  $\mathcal{V}[\text{ref } \tau]$  are inhabited by a location  $\ell$  in world  $W$  if the current world's heap typing  $W.\Psi$  maps  $\ell$  to the value relation  $\mathcal{V}[\tau]$  approximated to the step index in the world  $W.k$ . (The  $j$ -approximation of a type, written  $\lfloor \mathcal{V}[\tau] \rfloor_j$ , restricts  $\mathcal{V}[\tau]$  to inhabitants with worlds in  $\text{World}_j$ .)

Our expression relation  $\mathcal{E}[\tau]$  defines when a program  $P$  in world  $W$  behaves as a computation of type  $\tau$ . It says that for any heap  $H$  that satisfies the current world  $W$ , written  $H : W$ , and any non-Fail stack  $S$ , if the machine  $\langle H; S; P \rangle$  terminates in  $j$  steps (where  $j$  is less than our step budget  $W.k$ ), then either it ran to a non-type error or there exists some value  $v$  and some future world  $W'$  such that the resulting stack  $S'$  is the original stack with  $v$  on top, the resulting heap  $H'$  satisfies the future world  $W'$  and  $W'$  and  $v$  are in  $\mathcal{V}[\tau]$ .

At the bottom of Fig. 5, we show a syntactic shorthand,  $\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket$ , for showing that well-typed source programs, when compiled and closed off with well-typed substitutions  $\gamma$  that map variables to target values, are in the expression relation. Note  $\mathcal{G}[\Gamma]$  contains closing substitutions  $\gamma$  in world  $W$  that assign every  $x : \tau \in \Gamma$  to a  $v$  such that  $(W, v) \in \mathcal{V}[\tau]$ .

With our logical relation in hand, we can now state formal properties about our convertibility judgments.

**Lemma 3.1** (Convertibility Soundness).

*If  $\tau \sim \tau$ , then  $\forall (W, P) \in \mathcal{E}[\tau]. (W, (P, C_{\tau \mapsto \tau})) \in \mathcal{E}[\tau] \wedge \forall (W, P) \in \mathcal{E}[\tau]. (W, (P, C_{\tau \mapsto \tau})) \in \mathcal{E}[\tau]$ .*

*Proof.* We sketch the `ref bool`  $\sim$  `ref int` case; (rest elided, see [42]). For `ref bool`  $\sim$  `ref int`, what we need to show is that given any expression in  $\mathcal{E}[\text{ref bool}]$ , if we apply the conversion (which does nothing), the result will be in  $\mathcal{E}[\text{ref int}]$ . That requires  $\mathcal{V}[\text{ref bool}] = \mathcal{V}[\text{ref int}]$ .

The value relation at a reference type says that if you look up the location  $\ell$  in the heap typing of the world  $(W, \Psi)$ , you will get the value interpretation of the type. That means a `ref bool` must be a location  $\ell$  that, in the model, points to the value interpretation of `bool` (i.e.,  $\mathcal{V}[\text{bool}]$ ). In our model,

this must be true for all future worlds, which makes sense for ML-style references. Thus, for this proof to go through,  $\mathcal{V}[\text{bool}]$  must be the same as  $\mathcal{V}[\text{int}]$ , which it is.  $\square$

Once we have proved Lemma 3.1, we can prove semantic type soundness in the standard two-step way for our entire system. First, for each source typing rule, we define a compatibility lemma that is a semantic analog to that rule. For example, the compatibility lemma for the conversion typing rule, shown here, requires the proof of Lemma 3.1 to go through:

$$\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket \wedge \tau \sim \tau \implies \llbracket \Gamma; \Gamma \vdash (e)_{\tau} : \tau \rrbracket$$

Once we have all compatibility lemmas we can prove the following theorems as a consequence:

**Theorem 3.2** (Fundamental Property).

*If  $\Gamma; \Gamma \vdash e : \tau$  then  $\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket$  and if  $\Gamma; \Gamma \vdash e : \tau$  then  $\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket$ .*

**Theorem 3.3** (Type Safety for `RefLL`). *If  $\cdot; \cdot \vdash e : \tau$  then for any  $H : W$ , if  $\langle H; \cdot; e^+ \rangle \xrightarrow{*} \langle H'; S'; P' \rangle$ , then either  $\langle H'; S'; P' \rangle \rightarrow \langle H''; S''; P'' \rangle$ , or  $P' = \cdot$  and either  $S' = \text{Fail } c$  for some  $c \in \{\text{CONV}, \text{IDX}\}$  or  $S' = v$ .*

**Theorem 3.4** (Type Safety for `RefHL`). *If  $\cdot; \cdot \vdash e : \tau$  then for any  $H : W$ , if  $\langle H; \cdot; e^+ \rangle \xrightarrow{*} \langle H'; S'; P' \rangle$ , then either  $\langle H'; S'; P' \rangle \rightarrow \langle H''; S''; P'' \rangle$ , or  $P' = \cdot$  and either  $S' = \text{Fail } c$  for some  $c \in \{\text{CONV}, \text{IDX}\}$  or  $S' = v$ .*

**Discussion.** In addition to directly passing across pointers, there are two alternative conversion strategies, both of which our framework would accommodate. First, we could create a new location and copy and convert the data. This would allow the more flexible convertibility which does not require references to “identical” types, but would not allow aliasing, which may be desirable. Second, we could convert  $(\text{unit} \rightarrow \tau) \times (\tau \rightarrow \text{unit})$  and  $(\text{unit} \rightarrow \tau) \times (\tau \rightarrow \text{unit})$  instead `ref  $\tau$`  and `ref  $\tau$`  (assuming we had pairs)—i.e., read/write proxies to the reference (similar to that used in [17]). This allows aliasing, i.e., both languages reading / writing to the same location, and is sound for arbitrary convertibility relations, but comes at a runtime cost at each read / write.

The choice to use the encoding described in this case study, or either of these options, is not exclusive—we could provide different options for different types in the same system, depending on the performance characteristics we need.

## 4 Affine & Unrestricted

In our second case study, we consider an affine language, `AFFI`, interacting with an unrestricted one, `MiniML`. We enforce `AFFI`'s at-most-once variable use dynamically in the target using the well-known technique described, e.g., in [50], where affine resources are protected behind `thunks` with stateful flags that raise runtime errors the second time

$$\begin{aligned}
\text{AtomVal}_n &= \{(W, v) \mid W \in \text{World}_n\} \\
\text{World}_n &= \{(k, \Psi) \mid k < n \wedge \Psi \subset \text{HeapTy}_k\} \\
\text{HeapTy}_n &= \{\ell \mapsto \text{Typ}_n, \dots\} \\
\text{Typ}_n &= \{R \in 2^{\text{AtomVal}_n} \mid \forall (W, v) \in R. \\
&\quad \forall W'. W \sqsubseteq W' \implies (W', v) \in R\} \\
\mathcal{V}[\text{bool}] &= \{(W, n)\} & \mathcal{V}[\text{unit}] &= \{(W, 0)\} \\
\mathcal{V}[\tau_1 + \tau_2] &= \{(W, [0, v]) \mid (W, v) \in \mathcal{V}[\tau_1]\} \\
&\quad \cup \{(W, [1, v]) \mid (W, v) \in \mathcal{V}[\tau_2]\} \\
\mathcal{V}[\tau_1 \rightarrow \tau_2] &= \{(W, \text{thunk lam } x.P) \mid \\
&\quad \forall v, W' \sqsupseteq W. (W', v) \in \mathcal{V}[\tau_1] \\
&\quad \implies (W', [x \mapsto v]P) \in \mathcal{E}[\tau_2]\} \\
\mathcal{V}[\text{ref } \tau] &= \{(W, \ell) \mid W.\Psi(\ell) = \lfloor \mathcal{V}[\tau] \rfloor_{W,k}\} \\
\mathcal{V}[\text{int}] &= \{(W, n)\} \\
\mathcal{V}[\tau] &= \{(W, [v_1, \dots, v_n]) \mid (W, v_i) \in \mathcal{V}[\tau]\} \\
\mathcal{V}[\tau_1 \rightarrow \tau_2] &= \{(W, \text{thunk lam } x.P) \mid \\
&\quad \forall v, W' \sqsupseteq W. (W', v) \in \mathcal{V}[\tau_1] \\
&\quad \implies (W', [x \mapsto v]P) \in \mathcal{E}[\tau_2]\} \\
\mathcal{V}[\text{ref } \tau] &= \{(W, \ell) \mid W.\Psi(\ell) = \lfloor \mathcal{V}[\tau] \rfloor_{W,k}\} \\
\mathcal{E}[\tau] &= \{(W, P) \mid \forall H: W, S \neq \text{Fail } \_ , H', S', j < W.k. \\
&\quad \langle H; S; P \rangle \xrightarrow{j} \langle H'; S'; \_ \rangle \implies S' = \text{Fail } c \wedge c \in \{\text{CONV}, \text{IDX}\} \\
&\quad \vee \exists v, W' \sqsupseteq W. (S' = S, v \wedge H' : W' \wedge (W', v) \in \mathcal{V}[\tau])\} \\
\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket &\equiv \forall W \gamma_{\Gamma} \gamma_{\Gamma}. (W, \gamma_{\Gamma}) \in \mathcal{G}[\Gamma] \wedge (W, \gamma_{\Gamma}) \in \mathcal{G}[\Gamma] \\
&\quad \implies (W, \text{close}(\gamma_{\Gamma}, \text{close}(\gamma_{\Gamma}, e^+))) \in \mathcal{E}[\tau] \\
\llbracket \Gamma; \Gamma \vdash e : \tau \rrbracket &\equiv \forall W \gamma_{\Gamma} \gamma_{\Gamma}. (W, \gamma_{\Gamma}) \in \mathcal{G}[\Gamma] \wedge (W, \gamma_{\Gamma}) \in \mathcal{G}[\Gamma] \\
&\quad \implies (W, \text{close}(\gamma_{\Gamma}, \text{close}(\gamma_{\Gamma}, e^+))) \in \mathcal{E}[\tau]
\end{aligned}$$

Figure 5. Logical relation for RefHL and RefLL.

the `thunk` is forced. However, an interesting and challenging aspect of our case study is that we only want to use dynamic enforcement when we lack static assurance that an affine variable will be use at most once.

**Languages.** We present the syntax of **AFFI**, **MiniML**, and our untyped Scheme-like functional target **LCVM** in Fig. 6 and selected static semantics in Fig. 7 (see supplementary material [42]). Our target **LCVM** is untyped, with functions, pattern matching, mutable references, and a standard operational semantics defined via steps  $\langle H, e \rangle \rightarrow \langle H', e' \rangle$  over heap and expression pairs. As in the previous case study, we will support open terms across language boundaries, and thus need to carry environments for both languages throughout our typing judgments.

To avoid unnecessary dynamic enforcement, we have two kinds of affine function types in **AFFI**:  $\rightarrow$  and  $\rightarrow\bullet$ .<sup>2</sup> We introduce a distinction between **AFFI** functions (and thus bindings) that may be passed across the boundary (our “dynamic” affine arrows  $\rightarrow$ , written with a hollow circle and bind dynamic affine variables  $a_{\circ}$ ), and ones that will only ever be

<sup>2</sup>In our supplementary materials [42], we also present a complete case study with a simpler variant of **AFFI**, which does not distinguish  $\rightarrow$ / $\rightarrow\bullet$  and thus does dynamic enforcement even on affine variables that have no interaction with unrestricted code.

$$\begin{aligned}
\text{AFFI} \\
\text{Type } \tau &::= \text{unit} \mid \text{bool} \mid \text{int} \mid \tau \rightarrow \tau \mid \tau \rightarrow\bullet \tau \mid !\tau \mid \tau \&\tau \mid \tau \otimes \tau \\
\text{Expr. } e &::= () \mid \text{true} \mid \text{false} \mid n \mid x \mid a_{\circ} \mid \lambda a_{\circ}. \tau.e \\
&\quad \mid e.e \mid (e)_{\tau} \mid !v \mid \text{let } !x = e \text{ in } e' \mid \langle e, e' \rangle \\
&\quad \mid e.1 \mid e.2 \mid (e, e) \mid \text{let } (a_{\bullet}, a'_{\bullet}) = e \text{ in } e' \\
\text{Value } v &::= () \mid \lambda a_{\circ}. \tau.e \mid !v \mid \langle e, e' \rangle \mid (v, v') \\
\text{Mode } \bullet &::= \circ \mid \bullet \\
\text{MiniML} \\
\text{Type } \tau &::= \text{unit} \mid \text{int} \mid \tau \times \tau \mid \tau + \tau \mid \tau \rightarrow \tau \mid \forall \alpha. \tau \mid \alpha \mid \text{ref } \tau \\
\text{Expr. } e &::= () \mid n \mid x \mid (e, e) \mid \text{fst } e \mid \text{snd } e \mid \text{inl } e \mid \text{inr } e \\
&\quad \mid \text{match } e \text{ x}\{e\} \text{ y}\{e\} \mid \lambda x : \tau.e \mid e.e \mid \Lambda \alpha.e \mid e[\tau] \\
&\quad \mid \text{ref } e \mid !e \mid e := e \mid (e)_{\tau} \\
\text{LCVM} \\
\text{Expr } e &::= () \mid n \mid \ell \mid x \mid (e, e) \mid \text{fst } e \mid \text{snd } e \mid \text{inl } e \mid \text{inr } e \\
&\quad \mid \text{if } e \{e\} \{e\} \mid \text{match } e \text{ x}\{e\} \text{ y}\{e\} \mid \text{let } x = e \text{ in } e \\
&\quad \mid \lambda x \{e\} \mid e.e \mid \text{ref } e \mid !e \mid e := e \mid \text{fail } c \\
\text{Values } v &::= () \mid n \mid \ell \mid (v, v) \mid \lambda x.e \\
\text{Err } c &::= \text{TYPE} \mid \text{CONV}
\end{aligned}$$

Figure 6. Syntax for MiniML, AFFI, and LCVM.

$$\begin{aligned}
&\frac{a_{\circ} : \tau \in \Omega}{\Delta; \Gamma; \Omega \vdash a_{\circ} : \tau} \quad \frac{\Delta; \Gamma; \Omega[a_{\circ} := \tau_1] \vdash e : \tau_2 \quad \text{no}_{\bullet}(\Omega)}{\Delta; \Gamma; \Omega \vdash \lambda a_{\circ} : \tau_1. e : \tau_1 \rightarrow \tau_2} \\
&\frac{\Delta; \Gamma; \Omega[a_{\circ} := \tau_1] \vdash e : \tau_2}{\Delta; \Gamma; \Omega \vdash \lambda a_{\bullet} : \tau_1. e : \tau_1 \rightarrow\bullet \tau_2} \\
&\frac{\Omega = \Omega_1 \uplus \Omega_2 \quad \Delta; \Gamma; \Omega_1 \vdash e_1 : \tau_1 \rightarrow\bullet \tau_2 \quad \Delta; \Gamma; \Omega_2 \vdash e_2 : \tau_1}{\Delta; \Gamma; \Omega \vdash e_1 e_2 : \tau_2} \\
&\frac{\Omega = \Omega_1 \uplus \Omega_2 \quad \Delta; \Gamma; \Omega_1 \vdash e : \tau_1 \otimes \tau_2 \quad \Delta; \Gamma; \Omega_2[a_{\bullet} := \tau_1, a'_{\bullet} := \tau_1] \vdash e' : \tau'}{\Delta; \Gamma; \Omega \vdash \text{let } (a_{\bullet}, a'_{\bullet}) = e \text{ in } e' : \tau'} \\
&\frac{\Omega = \Omega_e \uplus \Omega' \quad \text{no}_{\bullet}(\Omega_e) \quad \Delta; \Gamma; \Omega_e \vdash e : \tau \quad \_ : \tau \sim \tau}{\Gamma; \Omega; \Delta; \Gamma \vdash (e)_{\tau} : \tau} \\
&\frac{\Omega = \Omega_1 \uplus \Omega_2 \quad \Gamma; \Omega_1; \Delta; \Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma; \Omega_2; \Delta; \Gamma \vdash e_2 : \tau_1}{\Delta; \Gamma; \Omega \vdash e_1 e_2 : \tau_2}
\end{aligned}$$

Figure 7. Selected statics for AFFI and MiniML.

used within **AFFI** (our “static” affine arrows  $\rightarrow\bullet$ , written with a solid circle and bind static affine variables  $a_{\bullet}$ ).

We can see in Fig. 7 how **AFFI**’s affine-variable environment  $\Omega$  is maintained: variables are introduced by lambda and tensor-destructuring `let`, and environments are split across subterms, but all bindings are not required to be used, as we can see, in the variable rule. (In the full rules in supplementary material [42], a similar pattern shows up for base types). Since affine resources can exist within unrestricted **MiniML** terms, our affine environments  $\Omega$  need to be split, even in **MiniML** typing rules.

Note that we do not allow a dynamic function  $\lambda a_{\bullet} : \_ . e$  to close over static resources, as it may be duplicated if passed to

$\text{thunk}(e) \triangleq \text{let } r_{\text{fr}} = \text{ref } 1 \text{ in } \lambda\_.\{\text{if } !r_{\text{fr}} \{\text{fail CONV}\} \{r_{\text{fr}} := 0; e\}\}$   
 $() \rightsquigarrow () \quad n \rightsquigarrow n \quad \lambda x : \tau.e \rightsquigarrow \lambda x.\{e^+\} \quad \text{true/false} \rightsquigarrow 0/1$   
 $a_{\circ} \rightsquigarrow a_{\circ} () \quad a_{\bullet} \rightsquigarrow a_{\bullet} \quad \lambda a_{\circ/\bullet} : \tau.e \rightsquigarrow \lambda a_{\circ/\bullet}.\{e^+\}$   
 $(e_1 : \tau_1 \multimap \tau_2) e_2 \rightsquigarrow e_1^+ (\text{let } x = e_2^+ \text{ in } \text{thunk}(x))$   
 $(e_1 : \tau_1 \multimap \tau_2) e_2 \rightsquigarrow e_1^+ e_2^+$   
 $\text{let } (a_{\bullet}, a'_{\bullet}) = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x_{\text{fresh}} = e_1^+,$   
 $a_{\bullet} = \text{fst } x_{\text{fresh}},$   
 $a'_{\bullet} = \text{snd } x_{\text{fresh}} \text{ in } e_2^+$

Figure 8. Selected cases for **MiniML** and **AFFI** compilers.

**MiniML**, and thus the static resources would be unprotected. However, we do allow a dynamic function to accept a static closure as argument. This is safe because the dynamic guards will ensure that the static closure is called at most once. Once called, any static resources in its body will be used safely because the static closure typechecked.

We present selections of our compilers in Fig. 8 that highlight the interesting cases: how we compile variables, binders, and application. In the application cases, we can see that static variables do not introduce the overhead that dynamic variables have (see the `thunk` macro at the top of the figure that errors on second invocation).

**Convertibility.** We define convertibility relations and conversions for **AFFI** and **MiniML**, highlighting selections in Fig. 9 (see supplementary material for elided  $\text{unit} \sim \text{unit}$  and  $\tau_1 \otimes \tau_2 \sim \tau_1 \times \tau_2$ ). We focus on the conversion between  $\rightarrow$  and  $\multimap$  (note, of course, that it is impossible to safely convert  $\multimap$  to **MiniML**). Our compiler is designed to support affine code being mixed directly with unrestricted code. Intuitively, an affine function should be able to behave as an unrestricted one, but the other direction is harder to accomplish, and higher-order functions mean both must be addressed at once. In order to account for this, we convert  $\tau_1 \multimap \tau_2$  not to  $\tau_1 \rightarrow \tau_2$ , but rather to  $(\text{unit} \rightarrow \tau_1) \rightarrow \tau_2$ . That is, to a **MiniML** function that expects its argument to be a `thunk` containing a  $\tau_1$  rather than a  $\tau_1$  directly. Provided that the `thunk` fails if invoked more than once, we can ensure, dynamically, that a **MiniML** function with that type behaves as an **AFFI** function of a related type. These invariants are ensured by appropriate wrapping and use of the compiler macro `thunk(·)` (see top of Fig. 8).

**Semantic Model.** The most interesting part of this case study is the logical relation because we must build a model that allows us to show that the dynamic and static affine bindings within **AFFI** are used at most once. For a dynamic binding, this is tracked in target code by the dynamic reference flag created by the macro `thunk`. For a static binding, we use a similar strategy of tracking use via a flag, but rather than a target-level dynamic runtime flag, we create a *phantom* flag that exists only within our model. Specifically, we define an augmented target operational semantics that exists

$$\begin{array}{c}
 \hline
 C_{\text{int} \mapsto \text{bool}}, C_{\text{bool} \mapsto \text{int}} : \text{int} \sim \text{bool} \\
 \hline
 \frac{C_{\tau_1 \mapsto \tau_1}, C_{\tau_1 \mapsto \tau_1} : \tau_1 \sim \tau_1 \quad C_{\tau_2 \mapsto \tau_2}, C_{\tau_2 \mapsto \tau_2} : \tau_2 \sim \tau_2}{C_{\rightarrow}, C_{\rightarrow} : \tau_1 \multimap \tau_2 \sim (\text{unit} \rightarrow \tau_1) \rightarrow \tau_2} \\
 C_{\text{bool} \mapsto \text{int}}(e) \triangleq e \quad C_{\text{int} \mapsto \text{bool}}(e) \triangleq \text{if } e \ 0 \ 1 \\
 C_{\tau_1 \multimap \tau_2 \mapsto (\text{unit} \rightarrow \tau_1) \rightarrow \tau_2}(e) \triangleq \\
 \text{let } x = e \text{ in } \lambda x_{\text{thnk}}.\text{let } x_{\text{conv}} = C_{\tau_1 \mapsto \tau_1}(x_{\text{thnk}} ()) \text{ in} \\
 \text{let } x_{\text{acc}} = \text{thunk}(x_{\text{conv}}) \text{ in } C_{\tau_2 \mapsto \tau_2}(x \ x_{\text{acc}}) \\
 C_{(\text{unit} \rightarrow \tau_1) \rightarrow \tau_2 \mapsto \tau_1 \multimap \tau_2}(e) \triangleq \text{let } x = e \text{ in} \\
 \lambda x_{\text{thnk}}.\text{let } x_{\text{acc}} = \text{thunk}(C_{\tau_1 \mapsto \tau_1}(x_{\text{thnk}} ())) \text{ in } C_{\tau_2 \mapsto \tau_2}(x \ x_{\text{acc}})
 \end{array}$$

Figure 9. Selected convertibility rules for **MiniML** and **AFFI**.

solely for the model, and any program that runs without getting stuck under the augmented semantics has a trivial erasure to a program that runs under the standard semantics. This means we are using the model to identify a subset of target programs (the erasures of well-behaved augmented programs) that do not violate source type constraints (i.e., do not use static variables more than once), even if there is nothing in the target programs that actually witnesses those constraints (i.e., dynamic checks or static types).

We build the model as follows. First, we extend our machine configurations to keep track of *phantom flags*  $f$  – i.e., in addition to a heap  $H$  and term  $e$ , we have a *phantom flag set*  $\Phi$ . Second, the augmented semantics uses one additional term, `protect`, which consumes one of the aforementioned phantom flags when it reduces:

$$\begin{array}{l}
 \text{Expressions } e ::= \dots \text{protect}(e, f) \\
 \langle \Phi \uplus \{f\}, H, \text{protect}(e, f) \rangle \rightarrow \langle \Phi, H, e \rangle
 \end{array}$$

And finally, we modify the two rules that introduce bindings such that whenever a binding in the syntactic category  $\bullet$  is introduced, we create a new phantom flag (where “ $f$  fresh” means  $f$  is disjoint from all flags generated in this execution):

$$\begin{array}{c}
 f \text{ fresh} \\
 \hline
 \langle \Phi, H, \text{let } a_{\bullet} = v \text{ in } e \rangle \rightarrow \langle \Phi \uplus \{f\}, H, [a_{\bullet} \mapsto \text{protect}(v, f)]e \rangle \\
 f \text{ fresh} \\
 \hline
 \langle \Phi, H, \lambda a_{\bullet}.e \ v \rangle \rightarrow \langle \Phi \uplus \{f\}, H, [a_{\bullet} \mapsto \text{protect}(v, f)]e \rangle
 \end{array}$$

Note that we write  $\rightarrow$  for a step in this augmented semantics, to distinguish it from the true operational step  $\rightarrow$ . While phantom flags in the augmented operational semantics play a similar role in protecting static affine resources as dynamic reference flags in the dynamic case, the critical difference is that in the augmented semantics, a `protect(·)`-ed resource for which there is no phantom flag will get stuck, and thus be excluded from the logical relation by construction. This is very different from the dynamic case, where we want – and, in fact, need – to include terms that can fail in order to mix **MiniML** and **AFFI** without imposing an affine type



system on `MiniML` itself. What this means for the model is that dynamic reference flags are a *shared resource* that can be accessed from many parts of the program and therefore tracked in the world, while phantom flags are an *unique resource* which our type system ensures is owned/used by at most one part of the program, which is what allows us to prove that the augmented semantics will not get stuck.

While the full definitions are in our supplementary materials [42], we give a high-level description of our expression and value relations, shown in Fig. 10, noting that the high-level structure is similar to the first case study.

Our expression relation,  $\mathcal{E}[\tau]_\rho$ , is made up of tuples of worlds  $W$  and phantom flag stores / term pairs  $(\Phi_i, e_i)$ , where each flag store represents the phantom variables owned by the expression. Our worlds  $W$  keep the step index, a standard heap typing  $\Psi$  (see §3), but also an affine flag store  $\Theta$ , which maps dynamic flags  $\ell$  to *either* a marker that indicates a dynamic affine variable has been used (0, written `USED`), or the phantom flags  $\Phi$  that it closes over if it has not been used (a set that can be empty, of course). These dynamic flags  $\ell$  are a subset of the heap, disjoint from  $\Psi$  (which tracks the rest of the heap, i.e., all the normal/non-dynamic-flag references). The expression relation then says that, given a heap that satisfies the world and arbitrary “rest” of phantom flag store  $\Phi_r$  (disjoint from that closed over by the world and the owned portion), the term  $e$  will either: (i) run longer than the step index accounts for, (ii) `fail CONV` (error while converting a value), or (iii) terminate at some value  $e'$ , where the flag store  $\Phi$  has been modified to  $\Phi_f \uplus \Phi_g$ , the heap has changed to  $H'$ , and the new world  $W'$  is an extension of  $W$ . World extension ( $\sqsubseteq_{\Phi_r}$ ) is defined over worlds that do not contain phantom flags from  $\Phi_r$ , since phantom flags are a local resource and the world contains what is global. It allows the step index to decrease, the heap typing to gain (but not overwrite or remove) entries, and the affine store to mark (but not unmark) dynamic bindings as `USED`.

At that future world, we know that the resulting value, along with their  $\Phi_f$ , will be in the value relation  $\mathcal{V}[\tau]_\rho$ . The phantom flag store  $\Phi_g$  is “garbage” that is no longer needed, and the “rest” is unchanged. Note that, while running, some phantom flags may have moved into the new world but the new world cannot have absorbed what was in the “rest”.

Our value relation cases are now mostly standard, so we will focus only on the interesting ones:  $\rightarrow$  and  $\rightarrow\bullet$ .  $\mathcal{V}[\tau_1 \rightarrow \tau_2]$  is defined to take an arbitrary argument from  $\mathcal{V}[\tau_1]$ , which may own static phantom flags in  $\Phi$ , and add a new location  $\ell$  that will be used in the thunk that prevents multiple uses, but also store the phantom flags in the affine store. The idea is that a function  $\lambda a_\bullet : \_ . e$  can be applied to an expression that closes over static phantom flags, like `let (b•, c•) = (1, 2) in  $\lambda a_\bullet . b_\bullet$` —the latter will have phantom flags for both  $b_\bullet$  and  $c_\bullet$ . The body is then run with the argument substituted with a guarded expression. Now, consider

$$\begin{aligned}
& \text{guard}(e, \ell) \triangleq \lambda \_ . \{\text{if } !\ell \ \{\text{fail CONV}\} \ \{\ell := \text{USED}; e\}\} \\
& \mathcal{V}[\tau_1 \rightarrow \tau_2]_\rho = \{(W, (\emptyset, \lambda x. \{e\})) \mid \forall v \ W'. \\
& \quad W \sqsubseteq_\emptyset W' \wedge (W', (\emptyset, v)) \in \mathcal{V}[\tau_1]_\rho \\
& \quad \implies (W', (\emptyset, [x \mapsto v]e)) \in \mathcal{E}[\tau_2]_\rho\} \\
& \mathcal{V}[\tau_1 \rightarrow\bullet \tau_2]_\bullet = \{(W, (\emptyset, \lambda x \{e\})) \mid \forall \Phi \ v \ W'. \\
& \quad W \sqsubseteq_\emptyset W' \wedge (W', (\Phi, v)) \in \mathcal{V}[\tau_1]_\bullet. \\
& \quad \implies ((W'.k, W'.\Psi, W'.\Theta \uplus \ell \mapsto \Phi), \\
& \quad \quad (\emptyset, [x \mapsto \text{guard}(v, \ell)]e)) \in \mathcal{E}[\tau_2]_\bullet.\} \\
& \mathcal{V}[\tau_1 \rightarrow\bullet \tau_2]_\bullet = \{(W, (\Phi, \lambda a_\bullet . \{e\})) \mid \\
& \quad \forall \Phi' \ f_1 \ v \ W'. W \sqsubseteq_\Phi W' \wedge (W', (\Phi', v)) \in \mathcal{V}[\tau_1]_\bullet. \\
& \quad \wedge \Phi \cap \Phi' = \emptyset \wedge f \notin \Phi \uplus \Phi' \uplus \text{flags}(W') \\
& \quad \implies (W', (\Phi \uplus \Phi' \uplus \{f\}, [a_\bullet \mapsto \text{protect}(v, f)]e)) \\
& \quad \in \mathcal{E}[\tau_2]_\bullet.\} \\
& \mathcal{E}[\tau]_\rho = \{(W, (\Phi, e)) \mid \text{freevars}(e) = \emptyset \wedge \\
& \quad \forall \Phi_r, H:W, e', H', j < W.k. \Phi_r \# \Phi \wedge \Phi_r \uplus \Phi : W \wedge \\
& \quad \langle \Phi_r \uplus \text{flags}(W) \uplus \Phi, H, e \rangle \xrightarrow{j} \langle \Phi', H', e' \rangle \rightsquigarrow \\
& \quad \implies e' = \text{fail CONV} \vee (\exists \Phi_f \ \Phi_g \ W'. \\
& \quad \quad \Phi' = \Phi_r \uplus \text{flags}(W') \uplus \Phi_f \uplus \Phi_g \\
& \quad \quad \wedge W \sqsubseteq_{\Phi_r} W' \wedge H' : W' \wedge (W', (\Phi_f, e')) \in \mathcal{V}[\tau]_\rho)\} \\
& (k, \Psi, \Theta) \sqsubseteq_\Phi (j, \Psi', \Theta') \triangleq (j, \Psi', \Theta') \in \text{World}_j \wedge \\
& \quad j \leq k \wedge \Phi \# \text{flags}(k, \Psi, \Theta) \wedge \Phi \# \text{flags}(j, \Psi', \Theta') \\
& \quad \wedge \forall \ell \in \text{dom}(\Psi). \lfloor \Psi(\ell) \rfloor_j = \Psi'(\ell) \wedge \\
& \quad \forall \ell \in \text{dom}(\Theta). (\ell) \in \text{dom}(\Theta') \wedge \\
& \quad \quad (\Theta(\ell) = \text{USED} \implies \Theta'(\ell) = \text{USED}) \\
& \quad \wedge (\Theta(\ell) = \Phi \implies \Theta'(\ell) = (\text{USED} \vee \Phi))
\end{aligned}$$

Figure 10. Selections of `MiniML` & `Affi` Logical Relation.

what happens when the variable is used: the `guard(·)` wrapper will update the location to `USED`, which means that in the world, the phantom flags that were put at that location are no longer there – i.e., they are no longer returned by `flags(W')`, which returns all phantom flags closed over by dynamic flags. That means, for the reduction to be well-formed, the phantom flags have to move somewhere else—either back to being owned by the term (in  $\Phi_f$ ) or in the discarded “garbage”  $\Phi_g$ . Once the phantom flag set has been moved back out of the world, the flags can again be used by `protect(·)` expressions.

The static function,  $\mathcal{V}[\tau_1 \rightarrow\bullet \tau_2]_\bullet$ , has a similar flavor, but it may itself own static phantom flags. That means that the phantom flag set for the argument must be disjoint, and when we run the body, we combine the set along with a fresh phantom flags  $f$  for the argument, which are then put inside the `protect(·)` expressions.

With the logical relation in hand, we can prove analogous theorems to Lemma 3.1 (Convertibility Soundness), Theorem 3.2 (Fundamental Property), Theorem 3.3 (Type Safety for Lang A), and Theorem 3.4 (Type Safety for Lang B).

Note that to prove our type safety theorems, we prove a lemma which states that, if  $\langle H, e \rangle \xrightarrow{*} \langle H', e' \rangle \rightsquigarrow$ , then for any  $\Phi$ ,  $\langle \Phi, H, e \rangle \xrightarrow{*} \langle \Phi', H', e' \rangle \rightsquigarrow$ . This lemma is necessary because the given assumption of the type safety theorem is that the configuration  $\langle H, e \rangle$  steps under the normal operational

semantics, but to apply the expression relation, we need that a corresponding configuration steps to an irreducible configuration under the phantom operational semantics.

Although our phantom flag realizability model was largely motivated by efficiency concerns with the dynamic enforcement of affinity, more broadly, it demonstrates how one can build complex static reasoning into the model even if such reasoning is absent from the target. Indeed, the actual target language, which source programs are compiled to and run in, has not changed; the augmentations exist *only in the model*. In this way, the preservation of source invariants is subtle: it is not that the types actually exist in the target (via runtime invariants or actual target types), but rather that the operational behavior of the target is exactly what the type interpretations characterize.

## 5 Memory Management & Polymorphism

For our third case study, we consider how `MiniML`, whose references are garbage collected, can interoperate with core  $L^3$ , a language with safe strong updates despite memory aliasing, supported via linear capabilities [3]. This case study primarily highlights how different memory management strategies can interoperate safely, in particular, that manually managed linear references can be converted to garbage-collected references without copying. This is of particular interest as more low-level code is written in Rust, a language with an ownership discipline on memory that similarly could allow safe transfer of memory to garbage-collected languages.

We also use this case study to explore how polymorphism/generics in one language can be used, via a form of interoperability, from the other. This is interesting because significant effort has gone into adding generics to languages that did not originally support them, in order to more easily build certain re-usable libraries.<sup>3</sup> While we are not claiming that interoperability could entirely replace built-in polymorphism, sound support for cross-language type instantiation and polymorphic libraries presents a possible alternative, especially for smaller, perhaps more special-purpose, languages. This would allow us to write something like:

```
map((λx : int.x + 1))⟨int⟩→⟨int⟩([1, 2, 3])list ⟨int⟩
```

where the `blue` language supports polymorphism, and has a generic `map` function, while the `pink` language does not. Of course, since convertibility is still driving this, in addition to using a concrete `intlist`, `[1, 2, 3]`, as above, the language without polymorphism could convert entirely different (non-list) concrete representations into similar polymorphic ones — i.e., implementing a sort of polymorphic interface at the boundary. For example, rather than an `intlist` (or a `stringlist`), in the example above, one could start with an `intarray` or

`intbtree`, or any number of other traversable data structures that could be converted to `list int` (or any `list α`).

**Languages.** We present the syntax of  $L^3$ , augmented with forms for interoperability, in Fig. 11.  $L^3$  has linear capability types `cap ζ τ` (capability for abstract location  $\zeta$  storing data of type  $\tau$ ), unrestricted pointer types `ptr ζ` to support aliasing, and location abstraction ( $\Lambda \zeta.e : \forall \zeta.\tau$  and  $\ulcorner \zeta, v \urcorner : \exists \zeta.\tau$ ). The key insight to  $L^3$  is that the pointer can be separated from the capability and passed around in the program separately. At runtime, the capabilities will be erased, but the static discipline only allows pointers to be used with their capabilities (tied together with the type variables  $\zeta$ ), and only allows capabilities to be used linearly. This enables safe in-place updates and low-level manual memory management while still supporting some flexibility in terms of pointer manipulation. We refer the reader to our supplementary materials [42], or the original paper on  $L^3$  ([3]) for more details on its precise static semantics, but present highlights here. In particular, `new` allocates memory and returns an existential package containing a capability and pointer ( $\exists \zeta.\text{cap } \zeta \tau \otimes \text{ptr } \zeta$ ). `swap` takes a matching capability (`cap ζ τ1`) and pointer `ptr ζ` and a value (of a possibly different type  $\tau_2$ ) and replaces what is stored, returning the capability and old value `cap ζ τ2 ⊗ τ1`. Note that since capabilities record the type of what is in the heap and are unique, strong updates are safe. Finally, `free` takes a package of a capability and pointer ( $\exists \zeta.\text{cap } \zeta \tau \otimes \text{ptr } \zeta$ ) and frees the memory, consuming both in the process and returning what was stored there—any lingering pointers are harmless, as the necessary capability is now gone.

We compile both  $L^3$  and `MiniML` to an extension of the Scheme-like target LCVM that we used in the previous case study (see Fig. 13 for  $L^3$ ; `MiniML` is standard). Our additions to LCVM, shown in Fig. 12, add manual memory allocation (`alloc`), `free` (which will error on a garbage-collected location), an instruction (`gcmov`) to convert a manually managed location to garbage collected, and an instruction (`callgc`) to explicitly invoke the garbage collector. The last allows the compiler to decide where the GC can intercede (before allocation, in our compiler), and in doing so simplifies our model slightly. The memory management itself is captured in our heap definition, which allows the same location names to be used as either GC'd ( $\xrightarrow{gc}$ ) or manually managed ( $\xrightarrow{m}$ ), and re-used after garbage collection or manual free. Dereference (`!e`) and assignment (`e := e`) work on both types of reference (failing, of course, if it is manually managed and has been freed). This strategy of explicitly invoking the garbage collector and using a single pool of locations retains significant challenging aspects about garbage collectors while remaining simple enough to expose the interesting aspects of interoperation.

As in the previous case study, we have boundary terms,  $(e)_\tau$  and  $(e)^\tau$ , for *converting* a term and using it in the other

<sup>3</sup>e.g., Java 1.5/5, C# 2.0 [28] and more recently, in the Go programming language

$L^3$	
Type $\tau$	$::= \text{unit} \mid \text{bool} \mid \tau \otimes \tau \mid \tau \multimap \tau \mid !\tau$ $\mid \text{ptr } \zeta \mid \text{cap } \zeta \tau \mid \forall \zeta. \tau \mid \exists \zeta. \tau$
Value $v$	$::= \lambda x : \tau. e \mid () \mid \mathbb{B} \mid (v, v) \mid !v \mid \Lambda \zeta. e \mid \ulcorner \zeta, v \urcorner$
Expr. $e$	$::= v \mid x \mid (e, e) \mid e e \mid \text{let } () = e \text{ in } e \mid \text{if } e e e$ $\mid \text{let } (x, x) = e \text{ in } e \mid \text{let } !x = e \text{ in } e \mid \text{dupl } e$ $\mid \text{drop } e \mid \text{new } e \mid \text{free } e \mid \text{swap } e e e \mid e [\zeta]$ $\mid \ulcorner \zeta, e \urcorner \mid \text{let } \ulcorner \zeta, x \urcorner = e \text{ in } e \mid (e)_{\tau} \mid (e)_{\tau}$
DUPLICABLE	$= \{\text{unit}, \text{bool}, \text{ptr } \zeta, !\tau\}$

Figure 11. Syntax for  $L^3$ .

Expr $e$	$::= \dots \mid \text{alloc } e \mid \text{free } e \mid \text{gcmov } e \mid \text{callgc}$
Heap $H$	$::= \ell \xrightarrow{m} v, H \mid \ell \xrightarrow{gc} v, H \mid \cdot$
Err Code $c$	$::= \dots \mid \text{PTR}$

Figure 12. Additions to LCVM (see Fig. 6 for base LCVM).

language. Now, we also add new types  $\langle \tau \rangle$ , pronounced “foreign type”, and allow conversions from  $\tau$  to  $\langle \tau \rangle$  for *opaquely embedding*<sup>4</sup> types for use in polymorphic functions.

If a language supports polymorphism, then its type abstractions should be agnostic to the types that instantiate them, allowing them to range over not only host types, but indeed any foreign types as well. Doing so should not violate parametricity. However, the non-polymorphic language may need to make restrictions on how this power can be used, so as to not allow the polymorphic language to violate its invariants. To make this challenge material, our non-polymorphic language in this case study has linear resources (heap capabilities) that cannot, if we are to maintain soundness, be duplicated. This means, in particular, that whatever interoperability strategy we come up with cannot allow a linear capability from  $L^3$  to flow over to a *MiniML* function that duplicates it, even if such function is well-typed (and parametric) in *MiniML*.

**Convertibility.** The first conversion that we want to highlight is between references. In  $L^3$ , pointers have capabilities that convey ownership, and thus to convert a pointer we also need the corresponding capability. For brevity, we may use  $\text{REF } \tau$  to abbreviate a capability+pointer package type.

$$\frac{C_{\tau \mapsto \tau}, C_{\tau \mapsto \tau} : \tau \sim \tau}{C_{\text{REF } \tau \mapsto \text{ref } \tau}, C_{\text{ref } \tau \mapsto \text{REF } \tau} : \text{ref } \tau \sim \exists \zeta. \text{cap } \zeta \tau \otimes \text{ptr } \zeta}$$

$$C_{\text{REF } \tau \mapsto \text{ref } \tau}(e) \triangleq \text{let } x = \text{snd } e \text{ in}$$

$$\quad \text{let } \_ = (x := C_{\tau \mapsto \tau}(!x)) \text{ in gcmov } x$$

$$C_{\text{ref } \tau \mapsto \text{REF } \tau}(e) \triangleq \text{let } x = \text{alloc } C_{\tau \mapsto \tau}(!e) \text{ in } ((), x)$$

The glue code itself is quite interesting: going from  $L^3$  to *MiniML*, since the  $L^3$  type system guarantees that this is the only capability to this pointer, we can safely directly convert the pointer into a *MiniML* pointer with *gcmov* after in-place

<sup>4</sup>Similar to “lumps” in Matthews-Findler[33], though they give a *single* lump type for all foreign types, i.e., they would have only  $\langle \rangle$ , rather than  $\langle \tau \rangle$ .

$x \rightsquigarrow x$	$() \rightsquigarrow ()$	$\text{true/false} \rightsquigarrow 0/1$	$!v \rightsquigarrow v^+$	$\lambda x : \tau. e \rightsquigarrow \lambda x. e^+$
$e_1 e_2$	$\rightsquigarrow e_1^+ e_2^+$			
$\text{if } e_1 e_2 e_3$	$\rightsquigarrow \text{if } e_1^+ e_2^+ e_3^+$			
$(e_1, e_2)$	$\rightsquigarrow (e_1^+, e_2^+)$			
$\text{dupl } e$	$\rightsquigarrow \text{let } x = e^+ \text{ in } (x, x)$			
$\text{drop } e$	$\rightsquigarrow \text{let } \_ = e^+ \text{ in } ()$			
$\text{new } e$	$\rightsquigarrow \text{let } \_ = \text{callgc} \text{ in } \text{let } x_{\ell} = \text{alloc } e^+$ $\text{in } ((), x_{\ell})$			
$\text{free } e$	$\rightsquigarrow \text{let } x = e^+ \text{ in } \text{let } x_r = !( \text{snd } x) \text{ in}$ $\text{let } \_ = \text{free } (\text{snd } x) \text{ in } x_r$			
$\text{swap } e_c e_p e_v$	$\rightsquigarrow \text{let } x_p = e_p^+ \text{ in } \text{let } \_ = e_c \text{ in } \text{let } x_v = !x_p$ $\text{in } \text{let } \_ = (x_p := e_v^+) \text{ in } ((), x_v)$			
$\Lambda \zeta. e$	$\rightsquigarrow \lambda \_ . e^+$			
$e [\zeta]$	$\rightsquigarrow e^+ ()$			
$\ulcorner \zeta, e \urcorner$	$\rightsquigarrow e^+$			
$(e)_{\tau}$	$\rightsquigarrow C_{\tau \mapsto \tau}(e^+)$			
$\text{let } () = e_1 \text{ in } e_2$	$\rightsquigarrow \text{let } \_ = e_1^+ \text{ in } e_2^+$			
$\text{let } (x_1, x_2) = e_1 \text{ in } e_2$	$\rightsquigarrow \text{let } p = e_1^+ \text{ in } \text{let } x_1 = \text{fst } p \text{ in}$ $\text{let } x_2 = \text{snd } p \text{ in } e_2^+$			
$\text{let } !x = e_1 \text{ in } e_2$	$\rightsquigarrow \text{let } x = e_1^+ \text{ in } e_2^+$			
$\text{let } \ulcorner \zeta, x \urcorner = e_1 \text{ in } e_2$	$\rightsquigarrow \text{let } x = e_1^+ \text{ in } e_2^+$			

Figure 13. Compiler for  $L^3$ .

replacing the contents with the result of converting (a less general rule that had a different premise might not need to convert, e.g., if the data was already compatible—see the first case study for more details). Going the other direction, from *MiniML* to  $L^3$ , there is no way for us to know if there are other aliases to the reference, so we can’t re-use the pointer. While we could simply disallow this conversion, and error if it were attempted, instead we copy and convert data into a freshly allocated manually managed location (note how, in the target, capabilities are erased to unit). In this case, as in many, there are multiple sound ways of converting, and it may be that a particular one makes more sense for your use case: we took the position that it was useful to get a copy of the data, unaliased, but perhaps a language designer would rather force the pointer to be dereferenced on the *MiniML* side and the underlying data converted.

We account for interoperability of polymorphism in two parts. First, we have a *foreign type*,  $\langle \tau \rangle$ , which embeds an  $L^3$  type into the type grammar of *MiniML*. This foreign type, like any *MiniML* type, can be used to instantiate type abstractions, define functions, etc, but *MiniML* has no introduction or elimination rules for it—terms of foreign type must come across from, and then be sent back to,  $L^3$ . These come by way of the conversion rule  $\langle \tau \rangle \sim \tau$ , which allow terms of the form  $(e)_{\langle \tau \rangle}$  (to bring an  $L^3$  term to *MiniML*) and  $(e)_{\tau}$  (the reverse). Moreover, the conversion rule for foreign types restricts  $\tau$  to a safe *DUPLICABLE* subset of types, but has no runtime consequences:

$$\frac{\tau \in \text{DUPLICABLE}}{C_{\langle \tau \rangle \mapsto \tau}, C_{\tau \mapsto \langle \tau \rangle} : \langle \tau \rangle \sim \tau} \quad C_{\langle \tau \rangle \mapsto \tau}(e) \triangleq e$$

$$C_{\tau \mapsto \langle \tau \rangle}(e) \triangleq e$$

To prove soundness we need to show that `DUPLICABLE` types are indeed safe to embed. The soundness condition depends on the expressive power of the two languages when viewed through the lens of polymorphism. In our case, since the non-polymorphic language is linear but the polymorphic one is not, we need to show that a `DUPLICABLE` type can be copied (i.e., none of its values own linear capabilities)—this includes `unit` and `bool`, but also `ptr ζ` and any type of the form `!τ`. Now, consider examples using this:

$$\begin{aligned} & (\lambda\alpha.\lambda x:\alpha.\lambda y:\alpha.y) [\langle \mathbf{bool} \rangle] (\mathbf{true})_{\langle \mathbf{bool} \rangle} (\mathbf{false})_{\langle \mathbf{bool} \rangle} \quad (1) \\ & (\lambda x : \mathbf{BOOL}.x) (\mathbf{true})_{\mathbf{BOOL}} \text{ where } \mathbf{BOOL} \triangleq \forall\alpha.\alpha \rightarrow \alpha \rightarrow \alpha \quad (2) \end{aligned}$$

In (1), the leftmost expression is a polymorphic `MiniML` function that returns the second of its two arguments. It is instantiated it with a foreign type, `⟨bool⟩`. Next, two terms of type `bool` in  $\mathbf{L}^3$  are embedded via the foreign conversion, `(·)_{⟨bool⟩}`, which requires that `bool`  $\in$  `DUPLICABLE`. Not only does this mechanism allow  $\mathbf{L}^3$  programmers to use polymorphic functions, but also `MiniML` programmers to use new base types. Of course, we could also convert the actual values, as in (2). To do so, we can define conversions between Church booleans in `MiniML` (which has no booleans) and ordinary booleans in  $\mathbf{L}^3$ :

$$\frac{}{\forall\alpha.\alpha \rightarrow \alpha \rightarrow \alpha \sim \mathbf{bool}} \quad \begin{array}{l} \mathbf{C}_{\mathbf{BOOL} \mapsto \mathbf{bool}}(e) \triangleq e () 0 1 \\ \mathbf{C}_{\mathbf{bool} \mapsto \mathbf{BOOL}}(e) \triangleq \text{if0 } e \{ \lambda\alpha.\lambda x:\alpha.\lambda y:\alpha.x \\ \quad \{ \lambda\alpha.\lambda x:\alpha.\lambda y:\alpha.y \} \end{array}$$

**Semantic Model.** In Fig. 14, we present parts of the logical relation that we use to prove our conversions and entire languages sound (see supplementary material [42]).

Our model is inspired by that of core  $\mathbf{L}^3$  [3], though ours is significantly more complex to account for garbage collection and interoperability with `MiniML`. The key is a careful distinction between owned (linear) manual memory, which is *local* and described by heap fragments associated with terms, and garbage-collected memory, which is *global* and described by the world  $W$ . Since memory can be freed (via garbage collection or manual free), reused, and moved from manual memory to garbage-collected memory, there are several constraints on how heap fragments and worlds may evolve so we can ensure safe memory usage.

With that in mind, our value interpretation of source types  $\mathcal{V}[\tau]_\rho$  are sets of worlds and related heap-fragments-and-values  $(H, v)$ , where the heap fragment  $H$  paired with value  $v$  is the portion of the manually managed heap that  $v$  owns.

The relational substitution  $\rho$  maps type variables  $\alpha$  to arbitrary type interpretations  $R$  and location variables  $\zeta$  to concrete locations  $\ell$ . Since `MiniML` cannot own manual (linear) memory, all cases of  $\mathcal{V}[\tau]_\rho$  have empty  $\emptyset$  heap fragments. However, during evaluation, memory could be allocated and subsequently freed so the expression relation does not have

that restriction. In  $\mathbf{L}^3$ , pointer types `ptr ζ` do not own locations, so they can be freely copied. Rather, linear capabilities `cap ζ τ` convey ownership of the location  $\ell$  that  $\zeta$  maps to and the heap fragment  $H$  pointed to by  $\ell$ .

In the expression relation  $\mathcal{E}[\tau]_\rho$ , we run the expression with a set of pinned locations ( $L$ ) that the garbage collector should not touch (which may come from an outer context if we are evaluating a subterm), a garbage-collected heap fragment that satisfies the world ( $H_{g+}$ ), an arbitrary disjoint manually allocated (*MHeap*) “rest” of the heap ( $H_r$ ), composed with the owned fragment ( $H$ ). Then, assuming  $e$  terminates at  $v$ , we expect the “rest” heap is unchanged, the garbage-collected portion has been transformed to  $H'_g$ , the owned portion has been transformed into  $H'$ , and that  $(W', (H', v)) \in \mathcal{V}[\tau]_\rho$ , where  $W'$  is a world the transformed GC'd portion of the heap  $H'_g$  must satisfy.

Critical to the relation is world extension, written  $\sqsubseteq_{\mathbb{L}, \eta}$ , which indicates how our logical worlds can evolve over time. In typical logical relations for state, the heap grows monotonically and no location is ever overwritten, which world extension captures. But, in our setting, the future heap might have deallocated, overwritten, re-used memory (and re-used it between the GC and manual allocation). We can't just allow arbitrary future states, however, as the semantics of types do dictate restrictions on what has to happen in the heap. In particular, there are two sets of locations that we need to keep careful track of: the rest can change freely. The first are manually managed locations that we can't disturb, which index  $\mathbb{L}$  captures. Those are generally just the owned locations of term that we are currently running. The second are the garbage collected locations that we must preserve in the heap, at the same type (but we can change the value of), captured by  $\eta$ . We also have a syntactic shorthand, denoted by  $\sqsubseteq$ , that is indexed by the heap  $H$  and the expressions  $e$ . This syntactic shorthand is defined so that  $\mathbb{L}$  takes its manually managed locations from the domain of  $H$  while  $\eta$  takes its garbage collected locations as the locations in the original world that are present in either some value in the heap  $H$  or the expression  $e$ . Finally, we often use `rchgclocs` in order to compute  $\eta$  when using world extension. `rchgclocs(W, S)` is the set of locations in the world  $W$  that are actually mentioned in the set  $S$ ; i.e., `rchgclocs(W, S) = dom(W) ∩ S`.

While our target supports dynamic failure (in the form of the fail term), our logical relation rules out that possibility, ensuring that there are no errors from the source nor from the conversion. This is, of course, a choice we made, which may be stronger than desired for some languages (and, indeed, for our previous two case studies), but given our choice of conversions, it is possible.

With the logical relation in hand, we prove analogous theorems to Lemma 3.1 (Convertibility Soundness), Theorem 3.2 (Fundamental Property), Theorem 3.3 (Type Safety for Lang A), and Theorem 3.4 (Type Safety for Lang B).

$$\begin{aligned}
\mathcal{V}[\alpha]_\rho &= \rho.F(\alpha) \\
\mathcal{V}[\mathbf{unit}]_\rho &= \{(W, (\theta, ()))\} \\
\mathcal{V}[\tau_1 \rightarrow \tau_2]_\rho &= \{(W, (\theta, \lambda x.e)) \mid \forall W', v. W \sqsubseteq_{\theta, e} W' \wedge \\
&\quad (W', (\theta, v)) \in \mathcal{V}[\tau_1]_\rho \implies (W', (\theta, [x \mapsto v]e)) \in \mathcal{E}[\tau_2]_\rho\} \\
\mathcal{V}[\forall \alpha. \tau]_\rho &= \{(W, (\theta, \lambda_.e),) \mid \forall R \in RelT, W' \\
&\quad W \sqsubseteq_{\theta, e} W' \implies (W', (\theta, e)) \in \mathcal{E}[\tau]_\rho[F(\alpha) \mapsto R]\} \\
\mathcal{V}[\mathbf{ref} \tau]_\rho &= \{(W, (\theta, \ell)) \mid W.\Psi(\ell) = \lfloor \mathcal{V}[\tau]_\rho \rfloor_{W.k}\} \\
\mathcal{V}[\langle \tau \rangle]_\rho &= \mathcal{V}[\tau]_\rho \\
\mathcal{V}[\mathbf{unit}]_\rho &= \{(W, (\theta, ()))\} \\
\mathcal{V}[\mathbf{bool}]_\rho &= \{(W, (\theta, b)) \mid b \in \{0, 1\}\} \\
\mathcal{V}[\tau_1 \otimes \tau_2]_\rho &= \{(W, (H_1 \uplus H_2, (v_1, v_2))) \mid \\
&\quad (W, (H_1, v_1)) \in \mathcal{V}[\tau_1]_\rho \wedge (W, (H_2, v_2)) \in \mathcal{V}[\tau_2]_\rho\} \\
\mathcal{V}[\tau_1 \multimap \tau_2]_\rho &= \{(W, (H, \lambda x.e)) \mid \forall W', H_v, v. \\
&\quad W \sqsubseteq_{H, e} W' \wedge (W', (H_v, v)) \in \mathcal{V}[\tau_1]_\rho \implies \\
&\quad (W', (H \uplus H_v, [x \mapsto v]e)) \in \mathcal{E}[\tau_2]_\rho\} \\
\mathcal{V}[\mathbf{!} \tau]_\rho &= \{(W, (\theta, v)) \mid (W, (\theta, v)) \in \mathcal{V}[\tau]_\rho\} \\
\mathcal{V}[\mathbf{ptr} \zeta]_\rho &= \{(W, (\theta, \ell)) \mid \rho.L3(\zeta) = \ell\} \\
\mathcal{V}[\mathbf{cap} \zeta \tau]_\rho &= \{(W, (H \uplus \{\ell \mapsto v\}, ())) \mid \\
&\quad \rho.L3(\zeta) = \ell \wedge (W, (H, v)) \in \mathcal{V}[\tau]_\rho\} \\
\mathcal{V}[\forall \zeta. \tau]_\rho &= \{(W, (H, \lambda_.e)) \mid \\
&\quad \forall \ell. (W, (H, e)) \in \mathcal{E}[\tau]_\rho[L3(\zeta) \mapsto \ell]\} \\
\mathcal{V}[\exists \zeta. \tau]_\rho &= \{(W, (H, v)) \mid \exists \ell. (W, (H, v)) \in \mathcal{V}[\tau]_\rho[L3(\zeta) \mapsto \ell]\} \\
\mathcal{E}[\tau]_\rho &= \{(W, (H, e)) \mid \forall L, v, H_{g+} : W, H_r : MHeap, H_* \\
&\quad (H_{g+} \uplus H \uplus H_r, e) \xrightarrow{*}_L (H_*, v) \rightarrow_L \\
&\quad \implies \exists H', H'_g, \exists W'. H_* = H'_g \uplus H' \uplus H_r \wedge H'_g : W' \wedge \\
&\quad W \sqsubseteq_{(\text{dom}(H_r), \text{rchgloc}(W, L \cup FL(\text{cod}(H_r)))} W' \\
&\quad \wedge (W', (H', v)) \in \mathcal{V}[\tau]_\rho \wedge H_{g'} = \emptyset\} \\
(k, \Psi) \sqsubseteq_{L, \eta} (j, \Psi') &= j \leq k \wedge L \# \text{dom}(\Psi') \\
&\quad \wedge \forall \ell \in \eta. \Psi'(\ell) = \lfloor \Psi(\ell) \rfloor_j
\end{aligned}$$

Note the **highlighted parts** only apply to **MiniML** types.

**Figure 14.** Logical Relation for **MiniML** and **L<sup>3</sup>**.

Our convertibility soundness result proves that our conversions above between garbage-collected and manual references, as well as **L<sup>3</sup>** booleans and **MiniML** Church booleans (described above) are sound. We also show that  $\tau_1 \rightarrow \tau_2 \sim \mathbf{!}(\tau_1 \multimap \tau_2)$  assuming  $\tau_1 \sim \tau_1$  and  $\tau_2 \sim \tau_2$ .

**Discussion.** While we showed how to handle universal types, handling existential types is another question. With our existing “foreign type” mechanism, we can support defining data structures and operations over them and passing both. For example, we could pass an expression of type  $\langle \text{int} \rangle \times \langle \text{int} \rangle \rightarrow \langle \text{int} \rangle \times \langle \text{int} \rangle \rightarrow \text{int}$ , for a counter defined as an integer. That provides some degree of abstraction, but doesn’t, for example, disallow passing the  $\langle \text{int} \rangle$  back to some other code that expects that type. We could, however, in the language with existential types, pack that to  $\exists \alpha. \alpha \times \alpha \rightarrow \alpha \times \alpha \rightarrow \text{int}$ .

More interesting is the question when both languages have polymorphism. In that case, if we wanted to convert abstract types, we would need to generalize our convertibility rules to handle open types, i.e.,  $\Delta \vdash \tau \sim \tau'$ . If the interpretation of type variables were the same in both languages

(i.e., in our model this would mean that both were drawn from the same relation), this would be sufficient. If, however, the interpretation of type variables were different in the two languages (we do this in the case study in §4, see our supplemental materials [42] for the use of *UnrTyp* in  $\mathcal{V}[\forall \alpha. \tau]_\rho$ ), we would need, in our source type systems, some form of bounded polymorphism in order to restrict the judgment to variables that were equivalent. Otherwise, it would be impossible to prove convertibility rules sound.

## 6 Related Work and Conclusion

Most research on interoperability has focused either on reducing boilerplate or improving performance. We will not discuss those, focusing on work addressing soundness.

*Multi-language semantics.* Matthews and Findler [33] studied the question of the interoperability of source languages, developing the idea of a syntactic multi-language with *boundary terms* (c.f., contracts [18, 19]) that mediate between the two languages. They focused on a static language interacting with a dynamic one, but similar techniques have been applied widely (e.g., object-oriented [20, 21], affine and unrestricted [50], simple and dependently typed [41], functional language and assembly [43], linear and unrestricted [45]) and used to prove compiler properties (e.g., correctness [44], full abstraction [2, 37]). More recently, there has been an effort to understand this construction from a denotational [15] and categorical [14] perspective. While the last may seem particularly relevant to our work, they still firmly root the multi-language as a source-language construct, rather than building it out of a common substrate, our key divergence from this prior work.

Barrett et al. [6] take a slightly different path, directly mixing languages (PHP and Python) and allowing bindings from one to be used in the other, though to similar ends.

*Interoperability via typed targets.* Shao and Trifonov [46, 51] studied interoperability much earlier, and closer to our context: they consider interoperability mediated by translation to a common target. They tackle the problem that one language has access to control effects and the other does not. Their approach, however, is different: it relies upon a target language with an effect-based type system that is sufficient to capture the safety invariants, whereas while our realizability approach can certainly benefit from typed target languages, it doesn’t rely upon them. While typed intermediate languages obviously offer real benefits, there are also unaddressed problems, foremost of which is designing a usable type system that is sufficiently general to allow (efficient) compilation from all the languages you want to support. While there are ongoing attempts (probably foremost is the TruffleVM project [22]) to design such general intermediates, most have focused their attention on untyped or unsound languages, and in the particular case of TruffleVM, there is as-yet no meta-theory.

*An abstract framework for unsafe FFI.* Turcotte et al. [52] advocate a framework using an abstract version of the foreign language, so soundness can be proved without building a full multi-language. They demonstrate this by proving a modified type safety proof of Lua and C interacting via the C FFI, modeling the C as code that can do arbitrary unsound behavior and thus blamed for all unsoundness. While this approach seems promising in the context of unsound languages, it is less clear how it applies to sound languages.

*Semantic Models and Realizability Models* The use of semantic models to prove type soundness has a long history [34]. We make use of step-indexed models [4, 5], developed as part of the Foundational Proof-Carrying Code [1] project, which showed how to scale the semantic approach to complex features found in real languages such as recursive types and higher-order mutable state. While much of the recent work that uses step-indexed models is concerned with program equivalence, one recent project that focuses on type soundness is RustBelt [27]: they give a semantic model of  $\lambda_{Rust}$  types and use it to prove the soundness of  $\lambda_{Rust}$  typing rules, but also to prove that the  $\lambda_{Rust}$  implementation of standard library features (essentially unsafe code) are semantically sound inhabitants of their ascribed type specification.

Unlike the above, our realizability model interprets source types as sets of target terms. Our work takes inspiration from a line of work by Benton and collaborators on “low-level semantics for high-level types” (dubbed “realistic realizability”) [8]. Such models were used to prove type soundness of standalone languages, specifically, Benton and Zarfaty [12] proved an imperative while language sound and Benton and Tabareau [11] proved type soundness for a simply typed functional language, both times interpreting source types as relations on terms of an idealized assembly and allowing for compiled code to be linked with a verified memory allocation module implemented in assembly [8]. Krishnaswami et al. [30] make use of a realizability model to prove consistency of LNL<sub>D</sub> a core type theory that integrates linearity and full type dependency. The linear parts of their model, like our interpretation of  $L^3$  types, are directly inspired by the semantic model for  $L^3$  by Ahmed et al. [3]. While they consider interoperability and use realizability models, their approach is quite different from ours, as they introduce both term constructors and types ( $G$  and  $F$ ) that allow direct embedding into the other language, thereby changing it, rather than defining conversions into existing types (which, indeed, is probably impossible in their case). More generally, such realizability models have also been used by Jensen et al. [26] to verify low-level code using a high-level separation logic, and by Benton and Hur [9] to verify compiler correctness.

Finally, New et al. [36, 38, 39] make use of realizability models in their work on semantic foundations of gradual typing, work that we have drawn inspiration from, given

gradual typing is a special instance of language interoperability. They compile type casts in a surface gradual language to a target Call-By-Push-Value [31] language without casts, build a realizability model of gradual types and type precision as relations on target terms, and prove properties about the gradual surface language using the model.

*Verification-based Approaches* Much work has been done using high-level program logics to reason about target terms, which can be seen as analogous to the realizability approach. Perhaps most relevant, in the context of interoperability, is the Cito system of Wang et al. [53], where code to-be-linked is given a specification over the behavior of target code, and compilation can then proceed relying upon that specification. This clearly renders benefits in terms of language independence, since any compiled code that satisfied that specification could be used. However, there is a significant difference from our work: by incorporating the semantics of types of both languages we can prove that the *conversions* preserve those semantics, and thus allow an end user to gain the benefits of type soundness without having to do any verification. Indeed, proving the conversions sound (or, in the case that they can be no-ops, proving that is okay) is the central result of this paper, and such conversions are not a part of the setup of Wang et al. [53].

**Conclusion and Future Work.** We have presented a novel framework for the design and verification of sound language interoperability where that interoperability happens, as in practical systems, after compilation. The realizability models at the heart of our technique give us powerful reasoning tools, including the ability to encode static invariants that are otherwise impossible to express in often untyped or low-level target languages. Even when it is possible to turn static source-level invariants into dynamic target-level checks, the ability to instead move these invariants into the model allows for more performant (and perhaps, realistic) compilers without losing the ability to prove soundness.

In the future, we hope to apply the framework to further explorations of the interoperability design space, e.g., to investigate interactions between lazy and strict languages (compilation to Call-By-Push-Value [31] may illuminate conversions), between single-threaded and concurrent languages (session types [24, 25, 49] may help guide interoperability with process calculi like the  $\pi$ -calculus [35]), between different control effects, and between Rust and a GC’ed language such as ML, Java, or Haskell compiled to a low-level target.

## Acknowledgments

We thank the anonymous reviewers for their in-depth comments. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1816837 and CCF-1453796.

## References

- [1] Amal Ahmed, Andrew W. Appel, Christopher D. Richards, Kedar N. Swadi, Gang Tan, and Daniel C. Wang. 2010. Semantic Foundations for Typed Assembly Languages. *ACM Transactions on Programming Languages and Systems* 32, 3 (March 2010), 1–67. <https://doi.org/10.1145/1709093.1709094>
- [2] Amal Ahmed and Matthias Blume. 2011. An equivalence-preserving CPS translation via multi-language semantics. In *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, Manuel M. T. Chakravarty, Zhenjiang Hu, and Olivier Danvy (Eds.). ACM, 431–444. <https://doi.org/10.1145/2034773.2034830>
- [3] Amal Ahmed, Matthew Fluet, and Greg Morrisett. 2007. L3 : A Linear Language with Locations. *Fundamenta Informaticae* 77, 4 (June 2007), 397–449. <https://content.iospress.com/articles/fundamenta-informaticae/fi77-4-06>
- [4] Amal Jamil Ahmed. 2004. *Semantics of Types for Mutable State*. Ph.D. Dissertation. Princeton University. <https://www.cs.princeton.edu/research/techreps/TR-713-04>
- [5] Andrew W. Appel and David A. McAllester. 2001. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.* 23, 5 (2001), 657–683. <https://doi.org/10.1145/504709.504712>
- [6] Edd Barrett, Carl Friedrich Bolz, Lukas Diekmann, and Laurence Tratt. 2016. Fine-grained Language Composition: A Case Study. In *30th European Conference on Object-Oriented Programming (ECOOP 2016) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 56)*, Shriram Krishnamurthi and Benjamin S. Lerner (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 3:1–3:27. <https://doi.org/10.4230/LIPIcs.ECOOP.2016.3>
- [7] David M. Beazley. 1996. SWIG: An Easy to Use Tool for Integrating Scripting Languages with C and C++. In *Fourth Annual USENIX Tcl/Tk Workshop 1996, Monterey, California, USA, July 10-13, 1996*, Mark Diekhans and Mark Roseman (Eds.). USENIX Association. <https://www.usenix.org/legacy/publications/library/proceedings/tcl96/beazley.html>
- [8] Nick Benton. 2006. Abstracting allocation: The new new thing. In *Computer Science Logic (CSL)*. [https://doi.org/10.1007/11874683\\_12](https://doi.org/10.1007/11874683_12)
- [9] Nick Benton and Chung-Kil Hur. 2009. Biorthogonality, Step-indexing and Compiler Correctness. In *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (Edinburgh, Scotland) (ICFP '09)*. ACM, New York, NY, USA, 97–108. <https://doi.org/10.1145/1596550.1596567>
- [10] Nick Benton, Andrew Kennedy, and Claudio V Russo. 2004. Adventures in interoperability: the sml.net experience. In *Proceedings of the 6th ACM SIGPLAN International conference on Principles and Practice of Declarative Programming*, 215–226. <https://doi.org/10.1145/1013963.1013987>
- [11] Nick Benton and Nicolas Tabareau. 2009. Compiling functional types to relational specifications for low level imperative code. In *Proceedings of TLDI'09: 2009 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, Savannah, GA, USA, January 24, 2009*, 3–14. <https://doi.org/10.1145/1481861.1481864>
- [12] Nick Benton and Uri Zarfaty. 2007. Formalizing and Verifying Semantic Type Soundness of a Simple Compiler. In *Proceedings of the 9th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (Wroclaw, Poland) (PPDP '07)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/1273920.1273922>
- [13] Matthias Blume. 2001. No-longer-foreign: Teaching an ML compiler to speak C “natively”. *Electronic Notes in Theoretical Computer Science* 59, 1 (2001), 36–52. [https://doi.org/10.1016/s1571-0661\(05\)80452-9](https://doi.org/10.1016/s1571-0661(05)80452-9)
- [14] Samuele Buro, Roy Crole, and Isabella Mastroeni. 2020. Equational logic and categorical semantics for multi-languages. *Electronic Notes in Theoretical Computer Science* 352 (2020), 79–103. <https://doi.org/10.1016/j.entcs.2020.09.005>
- [15] Samuele Buro and Isabella Mastroeni. 2019. On the Multi-Language Construction.. In *ESOP*, 293–321. [https://doi.org/10.1007/978-3-030-65474-0\\_14](https://doi.org/10.1007/978-3-030-65474-0_14)
- [16] Manuel MT Chakravarty. 1999. C->HASKELL, or Yet Another Interfacing Tool. In *Symposium on Implementation and Application of Functional Languages*. Springer, 131–148. [https://doi.org/10.1007/10722298\\_8](https://doi.org/10.1007/10722298_8)
- [17] Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. 2012. Complete Monitors for Behavioral Contracts. In *European Symposium on Programming (ESOP)*. [https://doi.org/10.1007/978-3-642-28869-2\\_11](https://doi.org/10.1007/978-3-642-28869-2_11)
- [18] Robert Bruce Findler and Matthias Blume. 2006. Contracts as pairs of projections. In *International Symposium on Functional and Logic Programming*. Springer, 226–241. [https://doi.org/10.1007/11737414\\_16](https://doi.org/10.1007/11737414_16)
- [19] Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for higher-order functions. In *Proceedings of the seventh ACM SIGPLAN international conference on Functional programming*, 48–59. <https://doi.org/10.1145/583852.581484>
- [20] Kathryn E Gray. 2008. Safe cross-language inheritance. In *European Conference on Object-Oriented Programming*. Springer, 52–75. [https://doi.org/10.1007/978-3-540-70592-5\\_4](https://doi.org/10.1007/978-3-540-70592-5_4)
- [21] Kathryn E Gray, Robert Bruce Findler, and Matthew Flatt. 2005. Fine-grained interoperability through mirrors and contracts. *ACM SIGPLAN Notices* 40, 10 (2005), 231–245. <https://doi.org/10.1145/1103845.1094830>
- [22] Matthias Grimmer, Chris Seaton, Roland Schatz, Thomas Würthinger, and Hanspeter Mössenböck. 2015. High-performance cross-language interoperability in a multi-language runtime. In *Proceedings of the 11th Symposium on Dynamic Languages*, 78–90. <https://doi.org/10.1145/2816707.2816714>
- [23] Rich Hickey. 2020. A history of Clojure. *Proceedings of the ACM on programming languages* 4, HOPL (2020), 1–46. <https://doi.org/10.1145/3386321>
- [24] Kohei Honda. 1993. Types for dyadic interaction. In *International Conference on Concurrency Theory*. Springer, 509–523. [https://doi.org/10.1007/3-540-57208-2\\_35](https://doi.org/10.1007/3-540-57208-2_35)
- [25] Kohei Honda, Vasco T Vasconcelos, and Makoto Kubo. 1998. Language primitives and type discipline for structured communication-based programming. In *European Symposium on Programming*. Springer, 122–138. <https://doi.org/10.1007/bfb0053567>
- [26] Jonas B. Jensen, Nick Benton, and Andrew Kennedy. 2013. High-Level Separation Logic for Low-Level Code (POPL '13). Association for Computing Machinery, New York, NY, USA, 301–314. <https://doi.org/10.1145/2429069.2429105>
- [27] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the Foundations of the Rust Programming Language. In *ACM Symposium on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158154>
- [28] Andrew Kennedy and Don Syme. 2001. Design and Implementation of Generics for the .NET Common Language Runtime. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation (Snowbird, Utah, USA) (PLDI '01)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/378795.378797>
- [29] Robert Kleffner. 2017. *A Foundation for Typed Concatenative Languages*. Master’s thesis. Northeastern University. <https://www2.ccs.neu.edu/racket/pubs/dissertation-kleffner.pdf>
- [30] Neelakantan R. Krishnaswami, Pierre Pradic, and Nick Benton. 2015. Integrating Linear and Dependent Types. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 17–30. <https://doi.org/10.1145/2676726.2676969>

- [31] Paul Blain Levy. 2001. *Call-by-Push-Value*. Ph. D. Dissertation. Queen Mary, University of London, London, UK.
- [32] Phillip Mates, Jamie Perconti, and Amal Ahmed. 2019. Under Control: Compositionally Correct Closure Conversion with Mutable State. In *ACM Conference on Principles and Practice of Declarative Programming (PPDP)*. <https://doi.org/10.1145/3354166.3354181>
- [33] Jacob Matthews and Robert Bruce Findler. 2007. Operational semantics for multi-language programs. In *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*, Martin Hofmann and Matthias Felleisen (Eds.). ACM, 3–10. <https://doi.org/10.1145/1190216.1190220>
- [34] Robin Milner. 1978. A theory of type polymorphism in programming. *J. Comput. Syst. Sci.* 17 (1978), 348–375. [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- [35] Robin Milner, Joachim Parrow, and David Walker. 1992. A calculus of mobile processes, i. *Information and computation* 100, 1 (1992), 1–40. [https://doi.org/10.1016/0890-5401\(92\)90008-4](https://doi.org/10.1016/0890-5401(92)90008-4)
- [36] Max S. New and Amal Ahmed. 2018. Graduality from Embedding-Projection Pairs. In *ICFP. Proceedings of the ACM on Programming Languages* 2, 73:1–73:30. <https://doi.org/10.1145/3236768>
- [37] Max S. New, William J. Bowman, and Amal Ahmed. 2016. Fully abstract compilation via universal embedding. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 103–116. <https://doi.org/10.1145/2951913.2951941>
- [38] Max S. New, Dustin Jamner, and Amal Ahmed. 2020. Graduality and Parametricity: Together Again for the First Time. *Proceedings of the ACM on Programming Languages* 4, POPL, 46:1–46:32. <https://doi.org/10.1145/3371114>
- [39] Max S New, Daniel R Licata, and Amal Ahmed. 2019. Gradual type theory. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 15:1–15:31. <https://doi.org/10.1017/s0956796821000125>
- [40] Martin Odersky and Matthias Zenger. 2005. Scalable component abstractions. In *Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*. 41–57. <https://doi.org/10.1145/1103845.1094815>
- [41] Peter-Michael Osera, Vilhelm Sjöberg, and Steve Zdancewic. 2012. Dependent interoperability. In *Proceedings of the sixth workshop on Programming Languages meets Program Verification, PLPV 2012, Philadelphia, PA, USA, January 24, 2012*, Koen Claessen and Nikhil Swamy (Eds.). ACM, 3–14. <https://doi.org/10.1145/2103776.2103779>
- [42] Daniel Patterson, Noble Mushtak, Andrew Wagner, and Amal Ahmed. 2022. Semantic Soundness for Language Interoperability (Technical Appendix). (March 2022). Available at <https://dbp.io/pubs/2022/semint-tr.pdf>.
- [43] Daniel Patterson, Jamie Perconti, Christos Dimoulas, and Amal Ahmed. 2017. FunTAL: reasonably mixing a functional language with assembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, Albert Cohen and Martin T. Vechev (Eds.). ACM, 495–509. <https://doi.org/10.1145/3062341.3062347>
- [44] James T. Perconti and Amal Ahmed. 2014. Verifying an Open Compiler Using Multi-language Semantics. In *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings (Lecture Notes in Computer Science, Vol. 8410)*, Zhong Shao (Ed.). Springer, 128–148. [https://doi.org/10.1007/978-3-642-54833-8\\_8](https://doi.org/10.1007/978-3-642-54833-8_8)
- [45] Gabriel Scherer, Max S. New, Nick Rioux, and Amal Ahmed. 2018. FabULous Interoperability for ML and a Linear Language. In *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 10803)*, Christel Baier and Ugo Dal Lago (Eds.). Springer, 146–162. [https://doi.org/10.1007/978-3-319-89366-2\\_8](https://doi.org/10.1007/978-3-319-89366-2_8)
- [46] Zhong Shao and Valery Trifonov. 1998. Type-directed continuation allocation. In *International Workshop on Types in Compilation*. Springer, 116–135. <https://doi.org/10.1007/bfb0055515>
- [47] T. Stephen Strickland, Sam Tobin-Hochstadt, Robert Bruce Findler, and Matthew Flatt. 2012. Chaperones and Impersonators: Run-Time Support for Reasonable Interposition. In *ACM International Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA) (Tucson, Arizona, USA)*. Association for Computing Machinery, New York, NY, USA, 943–962. <https://doi.org/10.1145/2384616.2384685>
- [48] Don Syme. 2006. Leveraging .NET meta-programming components from F# integrated queries and interoperable heterogeneous execution. In *Proceedings of the 2006 workshop on ML*. 43–54. <https://doi.org/10.1145/1159876.1159884>
- [49] Kaku Takeuchi, Kohei Honda, and Makoto Kubo. 1994. An interaction-based language and its typing system. In *International Conference on Parallel Architectures and Languages Europe*. Springer, 398–413. [https://doi.org/10.1007/3-540-58184-7\\_118](https://doi.org/10.1007/3-540-58184-7_118)
- [50] Jesse Tov and Riccardo Pucella. 2010. Stateful Contracts for Affine Types. In *Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings* (Paphos, Cyprus). [https://doi.org/10.1007/978-3-642-11957-6\\_29](https://doi.org/10.1007/978-3-642-11957-6_29)
- [51] Valery Trifonov and Zhong Shao. 1999. Safe and principled language interoperation. In *European Symposium on Programming*. Springer, 128–146. [https://doi.org/10.1007/3-540-49099-x\\_9](https://doi.org/10.1007/3-540-49099-x_9)
- [52] Alexi Turcotte, Ellen Arteca, and Gregor Richards. 2019. Reasoning About Foreign Function Interfaces Without Modelling the Foreign Language. In *33rd European Conference on Object-Oriented Programming (ECOOP 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 134)*, Alastair F. Donaldson (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 16:1–16:32. <https://doi.org/10.4230/LIPIcs.ECOOP.2019.16>
- [53] Peng Wang, Santiago Cuellar, and Adam Chlipala. 2014. Compiler Verification Meets Cross-Language Linking via Data Abstraction. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications* (Portland, Oregon, USA) (OOPSLA '14). Association for Computing Machinery, New York, NY, USA, 675–690. <https://doi.org/10.1145/2660193.2660201>
- [54] Jeremy Yallop, David Sheets, and Anil Madhavapeddy. 2018. A modular foreign function interface. *Science of Computer Programming* 164 (2018), 82–97. <https://doi.org/10.1016/j.scico.2017.04.002>
- [55] Jyun-Yan You. 2021. Rust Bindgen. <https://github.com/rust-lang/rust-bindgen>