Fully Abstract Compilation via Universal Embedding^{*}

Max S. New

Northeastern University, USA

maxnew@ccs.neu.edu

William J. Bowman

Northeastern University, USA

wjb@williamjbowman.com

Amal Ahmed

Northeastern University, USA

amal@ccs.neu.edu

Abstract

A *fully abstract* compiler guarantees that two source components are observationally equivalent in the source language if and only if their translations are observationally equivalent in the target. Full abstraction implies the translation is *secure*: target-language attackers can make no more observations of a compiled component than a source-language attacker interacting with the original source component. Proving full abstraction for realistic compilers is challenging because realistic target languages contain features (such as control effects) unavailable in the source, while proofs of full abstraction require showing that every target context to which a compiled component may be linked can be *back-translated* to a behaviorally equivalent source context.

We prove the first full abstraction result for a translation whose target language contains exceptions, but the source does not. Our translation—specifically, closure conversion of simply typed λ -calculus with recursive types—uses types at the target level to ensure that a compiled component is never linked with attackers that have more distinguishing power than source-level attackers. We present a new back-translation technique based on a shallow embedding of the target language into the source language at a dynamic type. Then boundaries are inserted that mediate terms between the untyped embedding and the strongly-typed source. This technique allows back-translating non-terminating programs, target features that are untypeable in the source, and well-bracketed effects.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory—Semantics

General Terms Languages, Security, Theory

Keywords Secure compilation, fully abstract compilation, universal embedding, universal domain, parametricity, logical relations, multi-language semantics, back-translation.

1. Introduction

When building secure software systems, programmers rely on language-provided abstractions and on the assumption that any

attacker—i.e., any code that their software component might be linked with—will be bound by the rules of the programming language. However, after the component is compiled, it may be linked with arbitrary target-level attackers that violate source-level abstractions, thus invalidating source-level security guarantees. Target attackers may be able to do things impossible in the source, such as read the compiled component’s private data, modify the component’s control flow, and even modify code implementing the component’s methods.

To guarantee that target attackers respect source-language rules, a compiler must be *fully abstract*—that is, it should *preserve* and *reflect* observational equivalence [1, 21, 7, 16, 30, 11, 14]. We use the standard notion of observational equivalence, also known as *contextual equivalence*: two components are contextually equivalent if they are indistinguishable in any valid (appropriately typed) program context. Fully abstract compilation ensures that when a source component e compiles to a target component e a valid target-language context C (attacker) does not have the power to observe anything more from interacting with e than a source-language context C interacting with e . Note that ensuring fully abstract compilation is only important when compiling *components* (not whole programs) since it is a property that ensures a component is protected from an attacker context (but whole programs have no context).

Achieving Secure Compilation There are three complementary techniques that one could use to achieve secure compilation. The first and most basic is to change the source program’s notion of program equivalence to be exactly the equivalence of compiled code. In this way all compilers can be seen to be secure, but at the cost of introducing extra-linguistic reasoning. A good example of this is programming in low-level languages like C where one reasons about programs by understanding their compilation.

The second way is to introduce mediating constructs such as dynamic checks or encodings that coerce low-level attackers into behaving like high-level code. This comes at the cost of runtime overhead in time and space. This is the approach taken in much of the literature (e.g., [16, 3, 29, 30, 12]).

The final way is to statically ensure that low-level program interfaces are secure, i.e., low-level programs operating at the specified interface must be in some way verified to act like high-level programs. This comes at the cost of disallowing linking with unverified/untyped code.

In this paper we focus on the final approach. In particular, we use types to statically verify that the target-language contexts that a compiled component is linked with have no more observational power than source-language contexts. Of course, taking this statically-enforced secure compilation approach in a realistic compiler will require designing *statically sound* (e.g., typed) intermediate languages for the compiler, as described in Ahmed [5]. Since fully abstract compilation is only important when compiling components, as noted above, we only need to perform type-preserving compilation until

^{*} In electronic versions of this paper, we use a blue sans-serif font to typeset our source language and a bold red serif font to typeset the target. The paper will be much easier to read if viewed/printed in color.

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 ACM 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.

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

ICFP’16, September 18–24, 2016, Nara, Japan
ACM. 978-1-4503-4219-3/16/09...\$15.00
<http://dx.doi.org/10.1145/2951913.2951941>

the intermediate language in the compiler where linking happens, yielding a whole program. Having type information available at link time is important in order to rule out linking with contexts that can attack the component in a way that no source context could have. However, after that point in the compilation pipeline, we can erase types and do whole-program compiler correctness—i.e., there is no further need for types or for proving full abstraction for the remaining passes of the compiler.

There are currently significant performance costs to existing dynamically-enforced secure compilers, and reducing these costs requires reliance on specialized hardware support—for instance, in the form of protected memory architectures (PMAs) as shown by Patrignani et al. [30]. Our work focuses on type-enforced secure compilation to ensure that dynamic checks *can* be avoided, without reliance on any specialized hardware, when low-level code *is* verified and to do this we need to ensure the compiler is secure when linking against a certain class of low-level programs.

Moreover, even if one were only to link against unverified low-level code, in order to minimize dynamic enforcement overhead we *should* use statically sound *intermediate* compiler passes. For example, recent work by Devriese et al. [12] showed full abstraction for a dynamically-protected type erasure pass for simply typed lambda calculus. The protection mechanism introduces deep checks to ensure safety for the compiler which would have to be compounded by further checks when compiling to lower-level languages. If multiple passes of the compiler are introducing deep semantic checks then the compiler will produce highly inefficient code. Thus, to minimize the cost of dynamic checks one should preserve high-level information (e.g., via types) as long as possible in a secure compiler.

Correct and Secure Compilation We start with specifications of correctness and security of our compiler. We give these specifications in terms of relations that we explain intuitively here, and formally define later.

We write $\Gamma \vdash e : \sigma \rightsquigarrow_e e$ when a source component e of type σ (under environment Γ) compiles to the target component e of translation type $\sigma^\dagger$ (under the translated environment Γ^+). We use the relations $\approx_S^{ctx}$ and $\approx_T^{ctx}$ to denote source and target contextual equivalence, respectively.

Theorem 1.1 (Semantics Preservation)

If $\Gamma \vdash e : \sigma \rightsquigarrow_e e$, then $\Gamma \vdash e \stackrel{S \simeq T}{\sim} e : \sigma$.

Semantics preservation, Theorem 1.1, states that the translation of a source *component* correctly preserves the behavior of the original source component, for a given specification $S \simeq T$ of related behavior. We specify related behavior by a cross-language relation that can be seen as a specification for the calling convention and encoding of source values. In other words, it specifies when a source component $e : \sigma$ is behaviorally equivalent to a target component $e : \sigma^\dagger$ (where e may or may not be an output of the compiler).

We define this cross-language relation in §3 by defining interoperability between source and target as a single “multi-language” as in [7, 31]. Then the cross-language relation is defined using contextual equivalence in the multi-language.

Theorem 1.1 gives us compositional compiler correctness, *i.e.*, correctness for components rather than for only whole programs [7, 31, 39, 27, 41]. In the special case that $\Gamma = \cdot$, we recover whole-program correctness.

Theorem 1.2 (Equivalence Reflection)

If $\Gamma \vdash e_1 : \sigma \rightsquigarrow_e e_1$, $\Gamma \vdash e_2 : \sigma \rightsquigarrow_e e_2$ and $\cdot; \Gamma^+ \vdash e_1 \approx_T^{ctx} e_2 : \sigma^\dagger$, then $\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma$.

Equivalence reflection, Theorem 1.2, states that if compiling two different source components results in two target components that are contextually equivalent in the target language, then the original source components must be contextually equivalent in the source

language. This is perhaps clearer if viewed as its contrapositive: if two source components *are distinguishable*, then so are their translations. Viewed in this way it can be seen as a weak form of semantics preservation that can be stated independent of the specification $S \simeq T$. We prove this as a corollary of the stronger semantics preservation theorem in §4.

Theorem 1.3 (Equivalence Preservation)

If $\Gamma \vdash e_1 : \sigma \rightsquigarrow_e e_1$, $\Gamma \vdash e_2 : \sigma \rightsquigarrow_e e_2$ and $\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma$, then $\cdot; \Gamma^+ \vdash e_1 \approx_T^{ctx} e_2 : \sigma^\dagger$.

Equivalence preservation, Theorem 1.3, ensures a *secure* compiler. It states that two observationally equivalent source components will be compiled to observationally equivalent target components. Viewed as a modularity property, this says that the abstraction barriers of the source language are preserved by compilation. Viewed as a security property, it says that attackers in the target language can be emulated by attackers in the source, so any reasoning about security of source programs can be done entirely in the source language. We prove equivalence preservation by back-translating target contexts to the source to show that they can make no more observations than source contexts in §5.

While equivalence preservation and reflection may at first glance seem to subsume the semantics-preservation theorem, note that neither mentions any notion of cross-language equivalence. Just proving full abstraction cannot guarantee that you use the correct calling convention or encoding. For instance, a compiler that compiled **true** to **false**, **false** to **true** and swapped branches of **ifs** preserves and reflects equivalences, but would not be semantics-preserving if the specification says that **true** should be encoded as **true**. Essentially, the abstract properties of equivalence preservation and reflection can’t tell if you agree with a concrete, “arbitrary”, choice of encoding.

Contributions The key to proving Theorem 1.3 is showing that for every well-behaved target context—i.e., context of translation type—there exists a behaviorally equivalent source context. This essentially requires *back-translating* a more expressive target language to a less expressive source language. Prior full abstraction results targeting typed languages performed *back-translation by partial evaluation* [7, 11, 37], a technique that does not seem to scale to non-terminating languages, which we discuss in §7.

We demonstrate a new back-translation technique, *universal embedding*. This technique is to first translate the target language into the source at a universal—i.e., dynamic—source type and then mediate between the universal type and each source type with functions that perform the encoding at run-time. Whereas previous work back-translates to types that precisely encode only target programs, we instead back-translate to a type that is overly large: it contains all the target behaviors but many other undesirable behaviors. This enables us to back-translate when the target is capable of finer type distinctions than the source, exemplified in our paper by the polymorphism of the target language.

This universal embedding demonstrates equivalence between an embedded interpreter and a multi-language semantics, two apparently similar techniques that have never been studied together before.

We apply this technique to prove full abstraction of typed closure conversion from a simply-typed λ -calculus with recursive types to a polymorphic λ -calculus with recursive types and exceptions. Our technique allows back-translating non-terminating programs, target-language features that are untypeable in the source, and well-bracketed/delimited effects. Ours is the first full-abstraction result that targets a *typed, non-terminating* target language with features unavailable in the source.

We provide complete definitions and proofs in our online technical appendix [28].

2. Closure Conversion

Our source and target languages are both call-by-value. They are also in monadic normal form—constructors and eliminators are only applied to syntactic values [10]—meant to represent compiler intermediate languages.

Types	$\sigma ::= \alpha \mid 1 \mid \sigma_1 + \sigma_2 \mid \sigma_1 \times \sigma_2 \mid \sigma_1 \rightarrow \sigma_2 \mid \mu\alpha. \sigma$
Values	$v ::= x \mid \langle \rangle \mid \text{inj}_i v \mid \langle v_1, v_2 \rangle \mid \lambda(x : \sigma). e \mid \text{fold}_{\mu\alpha. \sigma} v$
Expressions	$e ::= v \mid \text{case } v \text{ of } x_1. e_1 \mid x_2. e_2 \mid \pi_i v \mid v_1 v_2 \mid \text{unfold } v \mid \text{let } x = e_1 \text{ in } e_2$
Eval. Contexts	$K ::= [\cdot] \mid \text{let } x = K \text{ in } e_2$
General Contexts	$C ::= [\cdot] \mid \text{case } C \text{ of } x_1. e_1 \mid x_2. e_2 \mid \text{case } e \text{ of } x_1. C \mid x_2. C \mid \pi_i C \mid \lambda(x : \sigma). C \mid \dots$

$e \mapsto e'$	$\dots$
$K[(\lambda(x : \sigma). e) v] \mapsto K[e[v/x]]$	
$K[\text{unfold}(\text{fold}_{\mu\alpha. \sigma} v)] \mapsto K[v]$	

$\Gamma \vdash e : \sigma$	$\dots$
$\frac{x : \sigma \in \Gamma \quad \cdot \vdash \Gamma}{\Gamma \vdash x : \sigma}$	$\frac{\Gamma, x : \sigma_1 \vdash e : \sigma_2}{\Gamma \vdash \lambda(x : \sigma_1). e : \sigma_1 \rightarrow \sigma_2}$
$\frac{\Gamma \vdash v : \sigma[\mu\alpha. \sigma/\alpha]}{\Gamma \vdash \text{fold}_{\mu\alpha. \sigma} v : \mu\alpha. \sigma}$	$\frac{\Gamma \vdash v : \mu\alpha. \sigma}{\Gamma \vdash \text{unfold } v : \sigma[\mu\alpha. \sigma/\alpha]}$

Figure 1. λ^S : Syntax + Semantics (excerpts)

Source Language Our source language λ^S is a simply-typed lambda calculus with unit, sums, pairs, and recursive types. Figure 1 presents the syntax and excerpts of the semantics. We present the dynamic semantics ($e \mapsto e'$) using evaluation contexts K [15] to define a standard left-to-right call-by-value semantics. Since our language is in a normal form, the only non-trivial evaluation contexts are let-bindings. We elide most of the reduction rules and typing rules ($\Gamma \vdash e : \sigma$) as they are completely standard. The typing environment Γ maps term variables x to their types σ .

Figure 1 also presents an excerpt of the syntax for general contexts which are expressions with a single hole in them. We omit some of the details caused by the monadic syntax; for instance, some contexts can only be plugged with values. Context typing ($\Gamma \vdash C : (\Gamma \vdash \sigma) \Rightarrow (\Gamma' \vdash \sigma')$), ensures that for any expression e such that $\Gamma \vdash e : \sigma$, we can conclude that $\Gamma' \vdash C[e] : \sigma'$.

We define contextual equivalence ($\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma$) for λ^S as follows. Informally, two components e_1 and e_2 are contextually equivalent if either can be replaced by the other in any appropriately typed program context C without affecting the program's observable behavior. As it is a simple, functional language, we take termination (written $e \Downarrow$) as our notion of observable behavior. We write $e_1 \Downarrow$ if and only if $e_2 \Downarrow$.

Definition 2.1 (λ^S Contextual Equivalence)

$$\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma \stackrel{\text{def}}{=} \Gamma \vdash e_1 : \sigma \wedge \Gamma \vdash e_2 : \sigma \wedge \forall \sigma', C. \vdash C : (\Gamma \vdash \sigma) \Rightarrow (\cdot \vdash \sigma') \implies (C[e_1] \Downarrow \iff C[e_2] \Downarrow)$$

Target Language Our target language λ^T is a polymorphic λ -calculus with the empty type, sums, n -ary tuples, existential types, recursive types, and exceptions tracked by a modal type system. Figure 2 presents the syntax and excerpts of the dynamic and static semantics.

The target language has three syntactic categories for terms: v is a value, e is a computation that may have effects and r is a result, i.e. a normalized computation: either a returned value **return** v or a raised exception **raise** v .

The let-form of the λ^S is subsumed in λ^T by a combined let and try-catch form called **handle** in the style of Benton and Kennedy [9]. On a successful computation, i.e., a **return**, it continues with the left branch:

$$\text{handle return } v \text{ with } (x. e_1) (y. e_2) \mapsto e_1[v/x]$$

On an exception it continues with the right branch:

$$\text{handle raise } v \text{ with } (x. e_1) (y. e_2) \mapsto e_2[v/y]$$

We define let-forms as syntactic sugar for a handle that immediately re-raises any exception it encounters. We similarly define a more traditional try-catch by doing the opposite:

$$\begin{aligned} \text{let } x = e \text{ in } e' &\stackrel{\text{def}}{=} \text{handle } e \text{ with } (x. e') (y. \text{raise } y) \\ \text{catch } y = e \text{ in } e' &\stackrel{\text{def}}{=} \text{handle } e \text{ with } (x. \text{return } x) (y. e') \end{aligned}$$

We use a modal type system to track exceptions: τ is a value type (for values v) and θ is a computation type (for computations e). If e has type $\theta = E \tau_{\text{exn}} \tau$ then type soundness for this language means that if e reduces to a normal form it will either be a **return** v where v has type τ , or a **raise** v' where v' has type τ_{exn} . Crucially for our compiler, we can use the empty type 0 as the exception type to enforce that a computation does not throw an exception.

Context typing and contextual equivalence are defined analogously to λ^S .

Definition 2.2 (λ^T Contextual Equivalence)

$$\Delta; \Gamma \vdash e_1 \approx_T^{ctx} e_2 : \theta \stackrel{\text{def}}{=} \Delta; \Gamma \vdash e_1 : \theta \wedge \Delta; \Gamma \vdash e_2 : \theta \wedge \forall \theta', C. \vdash C : (\Delta; \Gamma \vdash \theta) \Rightarrow (\cdot \vdash \theta') \implies (C[e_1] \Downarrow \iff C[e_2] \Downarrow)$$

Closure Conversion Closure conversion is a standard internal compiler pass translating a substitution-based language into one that can be implemented with all values being passed by specified registers or memory locations. The pass translates functions with references to free variables, i.e. variables from the local environment, to be closed so that all variables references are bound by the functions parameters. We collect the values of free variables used in a function definition into a *closure environment* that is stored with the function, where the function itself is modified to take the environment as an additional input. We face two challenges in typing this translation to ensure full abstraction.

First, different terms of the same function type may have differently typed closure environments. To see why, consider two functions $e_1 = \lambda x. x$ and $e_2 = \lambda x. z$ of type $\text{bool} \rightarrow \text{bool}$, where z is a free variable of type bool . The function part of the translation of e_1 would have type $\langle \langle \rangle, \text{bool} \rangle \rightarrow \text{bool}$ while the translation of e_2 would have type $\langle \langle \text{bool} \rangle, \text{bool} \rangle \rightarrow \text{bool}$. Furthermore, there is no way to access the closure environment in the source language, so if the interface to the environment is too liberal then equivalence can't possibly be preserved. Both of these problems are solved in Minamide et al. [26] by using *existential types* to hide the type of the environment, as described below.

Second, to preserve equivalence when compiling to a target language with exceptions, we must ensure target contexts cannot use exceptions to make additional observations of translated terms. Consider $e_1 = \lambda f. (f \text{ true}; f \text{ false}; \langle \rangle)$ and $e_2 = \lambda f. (f \text{ false}; f \text{ true}; \langle \rangle)$. In a language with just non-termination, these terms are contextually equivalent. However, if the argument f can raise an exception, a context can distinguish these terms. The context $\text{catch } y = ([\cdot] (\lambda x. \text{raise } x))$ in y returns true when given e_1 and false when given e_2 . We use our checked exception type to ensure an exception cannot propagate into source code.

Our closure conversion pass extends the typed closure conversion in Minamide et al. [26] to accommodate our modal type system. Figure 4 presents the type translation which is split into the value type translation σ^+ and computation type translation $\sigma^\div$. A value of type σ is translated to a value of some value type $\tau = \sigma^+$. Non-trivial expressions of type σ are translated to some computation type

Value Types	$\tau ::= \alpha \mid \tau_1 + \tau_2 \mid \langle \tau \rangle \mid \forall[\alpha]. \tau \rightarrow \theta \mid \mu\alpha. \tau \mid 0 \mid \exists\alpha. \tau$
Computation Types	$\theta ::= E\tau_1 \tau_2$
Values	$v ::= x \mid \text{inj}_1 v_1 \mid \text{inj}_2 v_2 \mid \langle \bar{v} \rangle \mid \lambda[\alpha](x:\tau).e \mid \text{fold}_{\mu\alpha.\tau} v \mid \text{pack}(\tau, v)$
Results	$r ::= \text{return } v \mid \text{raise } v$
Computations	$e ::= r \mid \text{case } v \text{ of } x_1.e_1 \mid x_2.e_2 \mid v.i \mid v_1[\tau] v_2 \mid \text{unfold } v \mid \text{unpack}(\alpha, x) = v \text{ in } e \mid \text{handle } e \text{ with } (x.e_1)(y.e_2)$
Evaluation Contexts	$K ::= [\cdot] \mid \text{handle } K \text{ with } (x.e_1)(y.e_2)$

$e \mapsto e'$

$$\begin{array}{lcl}
 & K[(\lambda[\alpha](x:\tau).e) [\tau'] v] & \mapsto \dots \\
 & K[\text{unpack}(\alpha, x) = (\text{pack}(\tau, v)) \text{ in } e] & \mapsto K[e[\tau'/\alpha][v/x]] \\
 & K[\text{handle}(\text{return } v) \text{ with } (x.e_1)(y.e_2)] & \mapsto K[e[\tau/\alpha][v/x]] \\
 & K[\text{handle}(\text{raise } v) \text{ with } (x.e_1)(y.e_2)] & \mapsto K[e_1[v/x]] \\
 & & \mapsto K[e_2[v/y]]
 \end{array}$$

$\Delta \vdash \tau$

$$\begin{array}{cccc}
 \dots & \frac{\alpha \in \Delta}{\Delta \vdash \alpha} & \frac{\Delta, \alpha \vdash \tau \quad \Delta, \alpha \vdash \theta}{\Delta \vdash \forall[\alpha]. \tau \rightarrow \theta} & \frac{\Delta, \alpha \vdash \tau}{\Delta \vdash \mu\alpha. \tau} & \frac{\Delta, \alpha \vdash \tau}{\Delta \vdash \exists\alpha. \tau}
 \end{array}$$

$\Delta; \Gamma \vdash v : \tau$

$$\begin{array}{ccc}
 \dots & \frac{\Delta; \Gamma \vdash v_i : \tau_i}{\Delta; \Gamma \vdash \langle \bar{v} \rangle : \langle \bar{\tau} \rangle} & \frac{\Delta \vdash \Gamma \quad \alpha; x:\tau \vdash e : \theta}{\Delta; \Gamma \vdash \lambda[\alpha](x:\tau).e : \forall[\alpha]. \tau \rightarrow \theta} & \frac{\Delta; \Gamma \vdash v : \tau[\tau'/\alpha] \quad \Delta \vdash \tau'}{\Delta; \Gamma \vdash \text{pack}(\tau', v) : \exists\alpha. \tau}
 \end{array}$$

$\Delta; \Gamma \vdash e : \theta$

$$\begin{array}{lcl}
 \dots & \frac{\Delta; \Gamma \vdash v : \tau \quad \Delta \vdash \tau_{\text{exn}}}{\Delta; \Gamma \vdash \text{return } v : E\tau_{\text{exn}} \tau} & \frac{\Delta; \Gamma \vdash v : \tau_{\text{exn}} \quad \Delta \vdash \tau}{\Delta; \Gamma \vdash \text{raise } v : E\tau_{\text{exn}} \tau} & \frac{\Delta; \Gamma \vdash v : \exists\alpha. \tau \quad \Delta, \alpha; \Gamma, x:\tau \vdash e : E\tau_{\text{exn}} \tau' \quad \Delta \vdash \tau'}{\Delta; \Gamma \vdash \text{unpack}(\alpha, x) = v \text{ in } e : E\tau_{\text{exn}} \tau'} \\
 & \frac{\Delta; \Gamma \vdash e : E\tau'_{\text{exn}} \tau' \quad \Delta; \Gamma, x:\tau' \vdash e_1 : E\tau_{\text{exn}} \tau \quad \Delta; \Gamma, y:\tau'_{\text{exn}} \vdash e_2 : E\tau_{\text{exn}} \tau}{\Delta; \Gamma \vdash \text{handle } e \text{ with } (x.e_1)(y.e_2) : E\tau_{\text{exn}} \tau} & &
 \end{array}$$

Figure 2. λ^T : Syntax + Semantics (excerpts)

$\boxed{\Gamma \vdash v : \sigma \rightsquigarrow_v v}$	...	$\frac{x : \sigma \in \Gamma}{\Gamma \vdash x : \sigma \rightsquigarrow_v x}$	$\frac{(y_1, \dots, y_n) = \text{fv}(\lambda(x : \sigma').e) \quad \Gamma \vdash y_i : \sigma_i \quad \tau_{\text{env}} = \langle \sigma_1^+, \dots, \sigma_n^+ \rangle \quad \Gamma, x : \sigma \vdash e : \sigma' \rightsquigarrow_e e}{\Gamma \vdash \lambda(x : \sigma).e : \sigma \rightarrow \sigma' \rightsquigarrow_v \text{pack}(\tau_{\text{env}}, \langle \lambda(z : \langle \tau_{\text{env}}, \sigma^+ \rangle). \text{let } x_{\text{env}} = \text{return } z.1 \text{ in} \\ \text{let } y_i = \text{return } x_{\text{env}}.i \text{ in} \\ \dots \\ \text{let } x = \text{return } z.2 \text{ in } e \rangle)}$
$\boxed{\Gamma \vdash e : \sigma \rightsquigarrow_e e}$	...	$\frac{\Gamma \vdash v : \sigma \rightsquigarrow_v v}{\Gamma \vdash v : \sigma \rightsquigarrow_e \text{return } v}$	$\frac{\Gamma \vdash v_1 : \sigma_1 \rightarrow \sigma_2 \rightsquigarrow_v v_1 \quad \Gamma \vdash v_2 : \sigma_1 \rightsquigarrow_v v_2}{\Gamma \vdash v_1 v_2 : \sigma_2 \rightsquigarrow_e \text{unpack}(\alpha, z) = v_1 \text{ in} \\ \text{let } y_1 = \text{return } z.1 \text{ in} \\ \text{let } y_2 = \text{return } z.2 \text{ in} \\ y_1 \langle y_2, v_2 \rangle}$ $\frac{\Gamma \vdash e_1 : \sigma_1 \rightsquigarrow_e e_1 \quad \Gamma, x : \sigma_1 \vdash e_2 : \sigma_2 \rightsquigarrow_e e_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \sigma_2 \rightsquigarrow_e \text{let } x = e_1 \text{ in } e_2}$

Figure 3. Closure Conversion: Term Translation (excerpts)

$\sigma^\div = E0\sigma^+$
$\alpha^+ = \alpha$
$1^+ = \langle \rangle$
$(\sigma_1 + \sigma_2)^+ = \sigma_1^+ + \sigma_2^+ \quad (\cdot)^+ = \cdot$
$(\sigma_1 \times \sigma_2)^+ = \langle \sigma_1^+, \sigma_2^+ \rangle \quad (\Gamma, x:\sigma)^+ = \Gamma^+, x:\sigma^+$
$(\sigma_1 \rightarrow \sigma_2)^+ = \exists\alpha. \langle \langle \alpha, \sigma_1^+ \rangle \rightarrow \sigma_2^\div, \alpha \rangle$
$(\mu\alpha. \sigma)^+ = \mu\alpha. \sigma^+$

Figure 4. Closure Conversion: Type Translation

$\theta = \sigma^\div$, where $\sigma^\div = E0\sigma^+$, indicating that if this computation terminates it will result in a value of type σ^+ . The value type translation σ^+ is defined by structural recursion in all cases except for functions. A function of type $\sigma_1 \rightarrow \sigma_2$ is compiled to a closure, i.e., a pair of the function and its environment: $\exists\alpha. \langle \langle \alpha, \sigma_1^+ \rangle \rightarrow \sigma_2^\div, \alpha \rangle$. The type of the environment is existentially quantified so that functions of the same type but with different environments

are translated to functions of the same type. Parametricity of the language ensures that (standard) typed closure conversion is equivalence preserving. The existential types ensures that the function component of a closure can only ever be called with the environment it is packaged with, and ensures the environment can only be used as an argument to the function it is packaged with. Furthermore the output type of the function is $\sigma_2^\div = E0\sigma_2^+$, guaranteeing that when the function is called, it does not raise an exception. Viewed instead as a restriction on target programs, this means a target context cannot pass a closure that raises uncaught exceptions to compiled source code. Thus, target contexts cannot use exceptions to make additional observations.

The term translation is given in Figure 3. We define a value translation $\Gamma \vdash v : \sigma \rightsquigarrow_v v$ and an expression translation $\Gamma \vdash e : \sigma \rightsquigarrow_e e$. Note that since we translate open terms, we translate a free variable x to x —the same variable name but in the target language. In the expression translation, we translate values by first translating the

value according to the value translation, then wrapping the translated value in a **return**. We translate functions to an existential package, where all free variables referenced in the body of the function are put in the environment part of the package, and the function part of the package first introduces the free variables via **let** before executing the translated function body. We translate function application by first unpacking the closure and then applying the underlying function to the pair of the argument and the environment. We verify that this translation is type preserving.

Theorem 2.3 (Closure Conversion is Type Preserving)

1. If $\Gamma \vdash v : \sigma$ and $\Gamma \vdash v : \sigma \rightsquigarrow_v v$, then $\vdash \Gamma^+ \vdash v : \sigma^+$.
2. If $\Gamma \vdash e : \sigma$ and $\Gamma \vdash e : \sigma \rightsquigarrow_e e$, then $\vdash \Gamma^+ \vdash e : \sigma^+$.

3. Multi-Language Semantics

We use multi-language semantics [24, 7, 31] to define interoperability between source and target components. The multi-language λ^{ST} provides a natural, operational, specification for semantics preservation, given in §3.1.

<i>Value Types</i>	$\tau ::= \sigma \mid \tau$	$v ::= v \mid v$
<i>Comp. Types</i>	$\theta ::= \sigma \mid \theta$	$r ::= v \mid r$
<i>All Types</i>	$\varphi ::= \tau \mid \theta$	$e ::= e \mid e$
	$\Delta ::= \Delta$	$e ::= \dots \mid \sigma ST e$
	$\Gamma ::= \cdot \mid \Gamma, x : \sigma \mid \Gamma, x : \tau$	$e ::= \dots \mid TS^\sigma e$
	$K ::= \dots \mid \sigma ST K$	$C ::= \dots \mid \sigma ST C$
	$K ::= \dots \mid TS^\sigma K$	$C ::= \dots \mid TS^\sigma C$
	$K ::= K \mid K$	$C ::= C \mid C$

$\Delta; \Gamma \vdash e : \theta$		
$\dots$	$\frac{\Delta; \Gamma \vdash e : \sigma^+}{\Delta; \Gamma \vdash \sigma ST e : \sigma}$	$\frac{\Delta; \Gamma \vdash e : \sigma}{\Delta; \Gamma \vdash TS^\sigma e : \sigma^+}$

Figure 5. λ^{ST} : Syntax and Static Semantics (excerpts)

In Figure 5, we define the syntax for λ^{ST} by extending the syntax of λ^S and λ^T . For instance, we extend the source and target terms to include boundary terms, and extend expression contexts to include the boundary contexts.

Boundary terms allow embedding a term from one language into the other language, and their operational semantics define how the source language should interact with the target when compiled. The $\sigma ST e$ boundary embeds the target expression e of type σ^+ into the source at type σ , and the $TS^\sigma e$ embeds a source expression e of type σ into the target at type σ^+ . We define syntactic categories e for all terms, φ for all types, τ for value types, and so on.

We also define typing rules for boundaries in Figure 5. The two boundaries mediate between the source types σ and their translation types σ^+ . This ensures that source (resp. target) components can only be plugged into target (resp. source) contexts if their types are compatible according to the type translation. Other typing rules for the multi-language are inherited from the source and target languages with the environments for all the rules changed to $\Delta; \Gamma$.

Henceforth, when we refer to a “source term” (e) in λ^{ST} , we mean a multi-language term that has type σ . Analogously, a “target term” (e) in λ^{ST} , is a multi-language term of type θ .

The operational semantics for the multi-language, defined in Figure 6, inherits reductions from the source and target languages. Note that we evaluate under a boundary TS until we reach a value v , and under a boundary ST until we reach a result r .

The reduction rules for boundary terms perform a runtime encoding of values that corresponds to the closure conversion pass, directed by the type annotation σ on the boundary. Unlike our compiler which translates based on program syntax, this boundary translation is entirely extensional and only wraps values. For instance, $TS^\sigma \langle \rangle$ evaluates to **return** $\langle \rangle$. Since $TS^\sigma v$ boundaries are computations,

the encoding of v is injected into the exception monad in the target language. For either boundary around a fold, the value is unfolded in one language, then that value is translated and refolded in the other.

The function cases are most interesting because they correspond to the actual work of closure conversion.

To translate a source function, we have to construct a closure in the target. Since at runtime a source function will be closed by substitution, the closure has the empty environment. Additionally, the input and output of the function have to first be translated in order to call the original function. Translating a target closure to a source function is easier: the source function just unpacks the closure, translates the input, calls the target function with the translated input and its packed environment, and translates the output.

Contextual equivalence in the multi-language is defined in the same way as in the source and target languages.

Definition 3.1 (λ^{ST} Contextual Equivalence)

$$\Delta; \Gamma \vdash e_1 \approx_{\text{ST}}^{ctx} e_2 : \theta \stackrel{\text{def}}{=} \Delta; \Gamma \vdash e_1 : \theta \wedge \Delta; \Gamma \vdash e_2 : \theta \wedge \forall \theta', C. \vdash C : (\Delta; \Gamma \vdash \theta) \Rightarrow (\vdash \theta') \Rightarrow C[e_1] \Downarrow C[e_2]$$

Finally, we state Lemma 3.2, which provides a simple correctness property of our multi-language semantics, and implies Corollary 3.3. The latter is important as it justifies stating cross-language theorems with an arbitrary choice of placing the boundary on the source or target term.

Lemma 3.2 (Boundary Cancellation)

1. If $\Delta; \Gamma \vdash e : \sigma$, then $\Delta; \Gamma \vdash e \approx_{\text{ST}}^{ctx} \sigma ST TS^\sigma e : \sigma$.
2. If $\Delta; \Gamma \vdash e : \sigma^+$, then $\Delta; \Gamma \vdash e \approx_{\text{ST}}^{ctx} TS^\sigma \sigma ST e : \sigma^+$.

Corollary 3.3

$$\Delta; \Gamma \vdash e \approx_{\text{ST}}^{ctx} \sigma ST e : \sigma \text{ iff } \Delta; \Gamma \vdash TS^\sigma e \approx_{\text{ST}}^{ctx} e : \sigma^+.$$

3.1 Cross-Language Equivalence

Now that we have defined an interoperability semantics for the two languages, we have a natural way to define when a target term accurately simulates the behavior of a source term. This gives us a clear and concise definition of compiler correctness.

First we consider closed terms. Since our multi-language semantics defines embedding of a source component in a target context, we will consider a source component $e : \sigma$ and a target component $e : \sigma^+$ to be “equivalent” when $TS^\sigma e$ and e are contextually equivalent in λ^{ST} . We define cross-language equivalence on closed terms as follows.

Definition 3.4 (Cross-Lang. Equiv. $\simeq^{\text{ST}} : \text{Closed Terms}$)

$$\vdash e \simeq^{\text{ST}} e : \sigma \stackrel{\text{def}}{=} \vdash TS^\sigma e \approx_{\text{ST}}^{ctx} e : \sigma^+$$

Recall that we wish to prove *compositional* semantics preservation for closure conversion—that is, that compilation of *components*, not just whole programs, is semantics preserving. Since our notion of component is an open term, next we define cross-language equivalence for open terms, which in the special case of closed terms yields the equivalence above.

For cross-language equivalence for open terms, we have an open source term $\Gamma \vdash e : \sigma$ that we want to relate to an open target term $\vdash \Gamma^+ \vdash e : \sigma^+$ —in particular, for compiler correctness we will be interested in the case where e is the compilation of e . The multi-language semantics again gives us a natural definition in terms of boundaries. With closed terms we used boundaries to translate the *output* of the programs. With open terms we will use boundaries to additionally translate the *inputs*, that is, the free variables.

To do this we introduce new syntactic sugar that acts like the boundary on free variables.

$e \mapsto e'$	
$K[\mathcal{TS}^1 v]$	$\mapsto K[\text{return } \langle \rangle]$
$K[\mathcal{TS}^1 \text{return } v]$	$\mapsto K[\langle \rangle]$
$K[\sigma_1 + \sigma_2 \mathcal{TS} \text{return } (\text{inj}_i v)]$	$\mapsto K[\text{let } x = \sigma_i \mathcal{TS} \text{return } v \text{ in } \text{inj}_i x]$
$K[\mathcal{TS}^{\sigma_1 + \sigma_2} \text{inj}_i v]$	$\mapsto K[\text{let } x = \mathcal{TS}^{\sigma_i} v \text{ in } \text{return } \text{inj}_i x]$
$K[\sigma_1 \times \sigma_2 \mathcal{TS} \text{return } v]$	$\mapsto K[\text{let } x_1 = \sigma_1 \mathcal{TS} v.1 \text{ in } \text{let } x_2 = \sigma_2 \mathcal{TS} v.2 \text{ in } \langle x_1, x_2 \rangle]$
$K[\mathcal{TS}^{\sigma_1 \times \sigma_2} v]$	$\mapsto K[\text{let } x_1 = \mathcal{TS}^{\sigma_1} \pi_1 v \text{ in } \text{let } x_2 = \mathcal{TS}^{\sigma_2} \pi_2 v \text{ in } \text{return } \langle x_1, x_2 \rangle]$
$K[\sigma_1 \rightarrow \sigma_2 \mathcal{TS} \text{return } v]$	$\mapsto K[\lambda(x : \sigma_1). \sigma_2 \mathcal{TS} \left(\begin{array}{l} \text{unpack } (\alpha, z) = v \text{ in } \text{let } x_f = \text{return } z.1 \text{ in} \\ \text{let } x_{\text{env}} = \text{return } z.2 \text{ in} \\ \text{let } x = \mathcal{TS}^{\sigma_1} x \text{ in } x_f [\alpha] \langle x_{\text{env}}, x \rangle \end{array} \right)]]$
$K[\mathcal{TS}^{\sigma_1 \rightarrow \sigma_2} v]$	$\mapsto K[\text{return } (\text{pack } (\langle \rangle, \langle \lambda(z : \langle \rangle, \sigma_1^+) \rangle). \mathcal{TS}^{\sigma_2} (\text{let } x = \sigma_1 \mathcal{TS} \text{return } z.2 \text{ in } v x), \langle \rangle))]]]$
$K[\mu\alpha. \sigma \mathcal{TS} \text{return } v]$	$\mapsto K[\text{let } x = \sigma[\mu\alpha. \sigma / \alpha] \mathcal{TS} \text{return } (\text{unfold } v) \text{ in } \text{fold}_{\mu\alpha. \sigma} x]$
$K[\mathcal{TS}^{\mu\alpha. \sigma} v]$	$\mapsto K[\text{let } v = \mathcal{TS}^{\sigma[\mu\alpha. \sigma / \alpha]} \text{unfold } v \text{ in } \text{return } (\text{fold}_{(\mu\alpha. \sigma)^+} v)]$

Figure 6. λ^{ST} : Dynamic Semantics

Definition 3.5 (Boundaries on Free Variables)

For a term $\Gamma \vdash e : \sigma$, where $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$, we define $e \vdash^{\text{ST}} \text{let } x_1 = \sigma_1 \mathcal{TS} \text{return } x_1 \text{ in}$

$\vdots$
 $\text{let } x_n = \sigma_n \mathcal{TS} \text{return } x_n \text{ in } e.$

Note that $\vdash; \Gamma^+ \vdash e \vdash^{\text{ST}} : \sigma$. We define $e \vdash^{\text{ST}} \sigma$ analogously.

We now define a source term e to be equivalent to a target term e when the latter is contextually equivalent to the source term with its input and output suitably translated:

Definition 3.6 (Cross-Lang. Equiv. $\simeq^{\text{ST}}$: Closed Terms)

$\Gamma \vdash e \simeq^{\text{ST}} e : \sigma \stackrel{\text{def}}{=} \vdash; \Gamma^+ \vdash \mathcal{TS}^{\sigma} e \vdash^{\text{ST}} \text{let } x = \sigma \mathcal{TS} \text{return } x : \sigma^+$

4. Correctness and Equivalence Reflection

Direct proofs of contextual equivalence are intractable due to the universal quantification over contexts in the definition of $\approx_{\text{ST}}^{\text{ctx}}$. As is standard, we define a logical relation that is sound and complete with respect to contextual equivalence. We conclude this section with proofs of semantics preservation (Theorem 1.1) and equivalence reflection (Theorem 1.2) using the logical relation.

In Figure 7 we define a logical relation for λ^{ST} . We use a *step-indexed*, *biorthogonal* logical relation [13]. Biorthogonality is a standard technique [22, 32] to ensure that a logical relation is complete with respect to contextual equivalence. Step-indexing is a standard technique to provide an induction metric in the presence of unrestricted recursive types [4]—the logical relation is defined by nested induction on the step-index and types φ . As step-indices are not critical to understanding the definitions, we largely ignore them in our explanation.

The structure of the λ^{ST} logical relation is as follows: $\mathcal{O}$ relates programs that yield the same observations; $\mathcal{V}[\tau]$ relates values at type τ ; $\mathcal{R}[\tau]$ related *results* at type τ ; $\mathcal{K}[\theta]$ relates evaluation contexts that, when given related results of type θ , yield related observations; $\mathcal{E}[\theta]$ relates expressions of type θ that, when plugged into related evaluation contexts, yield related observations. Each relation is restricted to contain only well-typed members (elided for brevity). Each of these relations is indexed by a type φ and a relational interpretation ρ that provides mappings for the free type variables in φ .

The observation relation $\mathcal{O}$ defines when two terms yield the same observation. As is standard for non-total languages, we take termination to be a sufficient notion of observation. Accounting for the index k , $\mathcal{O}$ says that two terms yield the same observations if they both take k steps or if they both terminate.

The value relation $\mathcal{V}[\sigma]$ is the standard inductive definition for an effectful call-by-value language. At type 1 the unit value $\langle \rangle$ is

related to itself. Functions are related at type $\sigma_1 \rightarrow \sigma_2$ when, given arguments related in $\mathcal{V}[\sigma_1]$, they produce related results in $\mathcal{E}[\sigma_2]$.

The relation for polymorphic functions is standard for parametric polymorphism: for arbitrary types and admissible relations on those types, the bodies of the functions must be related under an extended relational interpretation. The dual relation for existential types is also standard: existential packages are related if some relational interpretation of the existential type variable relates them. The parametricity these rules ensure is crucial to the full abstraction proof. We write $\rho[\alpha \mapsto (\tau_1, \tau_2, R)]$ to be the relational interpretation ρ extended with types τ_1 and τ_2 and relation R for the type variable α . At type α two terms are related if they are in the relation for α in ρ , written $\rho_R(\alpha)$.

The relation $\mathcal{K}[\theta]$ is slightly non-standard to account for our modal type system. Specifically, two evaluation contexts are related in $\mathcal{K}[\theta]$ if, when plugged with related *results*—rather than the standard related *values*—yield related observations. Results, related by $\mathcal{R}[\tau]$ are just values in the source language, but are either **return** or **raise** of a value in the target language. In other words, results are what terms evaluate to.

In Figure 8, we lift our logical relation to open terms.

To close a term we first choose relational interpretations for all of the types ($\mathcal{D}[\Delta]$). Then related substitutions are just substitutions whose values are all related ($\mathcal{G}[\Gamma] \rho$). We write $\gamma_1(e)$ to close the term variables with the first component of the relational substitution γ , and similarly write $\gamma_2(e)$ to use the second component. We write $\rho_1(e)$ to substitute all type variables in e with the first component type in ρ and similarly $\rho_2(e)$ to use the second component type.

We prove the fundamental property of this logical relation, Theorem 4.1: every well-typed term is related to itself. The proof is by induction on typing derivations.

Theorem 4.1 (Fundamental Property)

If $\Delta; \Gamma \vdash e : \theta$, then $\Delta; \Gamma \vdash e \approx_{\mathcal{E}}^{\text{log}} e : \theta$.

The cases for boundary terms are most complex and require a proof by mutual induction on source types. We will use this case later (§5) and so specify it as a separate lemma Lemma 4.2.

Lemma 4.2 (Bridge Lemma)

- If $\Delta; \Gamma \vdash e_1 \approx_{\mathcal{E}}^{\text{log}} e_2 : \sigma^+$ then $\Delta; \Gamma \vdash \sigma \mathcal{TS} e_1 \approx_{\mathcal{E}}^{\text{log}} \sigma \mathcal{TS} e_2 : \sigma$.
- If $\Delta; \Gamma \vdash e_1 \approx_{\mathcal{E}}^{\text{log}} e_2 : \sigma$ then $\Delta; \Gamma \vdash \mathcal{TS}^{\sigma} e_1 \approx_{\mathcal{E}}^{\text{log}} \mathcal{TS}^{\sigma} e_2 : \sigma^+$.

Finally, we prove that the logical relation and contextual equivalence for the multi-language coincide. The proof is completely standard [4], using CIU equivalence as an intermediate step.

Theorem 4.3 ($\text{ctx} \equiv \text{log}$)

$\Delta; \Gamma \vdash e_1 \approx_{\text{ST}}^{\text{ctx}} e_2 : \theta$ if and only if $\Delta; \Gamma \vdash e_1 \approx_{\mathcal{E}}^{\text{log}} e_2 : \theta$.

Note: all relations are restricted to well-typed members.

$\text{Rel}[\tau_1, \tau_2]$	$\stackrel{\text{def}}{=} \{ R \mid \forall (k, \mathbf{v}_1, \mathbf{v}_2) \in R. \cdot \vdash \mathbf{v}_1 : \tau_1 \wedge \cdot \vdash \mathbf{v}_2 : \tau_2 \wedge \forall j < k. (j, \mathbf{v}_1, \mathbf{v}_2) \in R \}$
$\text{running}(k, e)$	$\stackrel{\text{def}}{=} \exists e'. e \mapsto^k e'$
$\mathcal{O}$	$\stackrel{\text{def}}{=} \{ (k, e_1, e_2) \mid (e_1 \Downarrow \wedge e_2 \Downarrow) \vee (\text{running}(k, e_1) \wedge \text{running}(k, e_2)) \}$
$\mathcal{V}[\mathbf{1}] \rho$	$\stackrel{\text{def}}{=} \{ (k, \langle \rangle, \langle \rangle) \}$
$\mathcal{V}[\sigma_1 + \sigma_2] \rho$	$\stackrel{\text{def}}{=} \{ (k, \text{inj}_i \mathbf{v}_1, \text{inj}_i \mathbf{v}_2) \mid i \in \{1, 2\} \wedge (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma_i] \rho \}$
$\mathcal{V}[\sigma \times \sigma'] \rho$	$\stackrel{\text{def}}{=} \{ (k, \langle \mathbf{v}_1, \mathbf{v}'_1 \rangle, \langle \mathbf{v}_2, \mathbf{v}'_2 \rangle) \mid (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma] \rho \wedge (k, \mathbf{v}'_1, \mathbf{v}'_2) \in \mathcal{V}[\sigma'] \rho \}$
$\mathcal{V}[\sigma \rightarrow \sigma'] \rho$	$\stackrel{\text{def}}{=} \{ (k, \lambda(x : \rho_1(\sigma)). e_1, \lambda(x : \rho_2(\sigma')). e_2) \mid \forall j \leq k. \forall \mathbf{v}_1, \mathbf{v}_2. (j, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma] \rho \implies (j, e_1[\mathbf{v}_1/x], e_2[\mathbf{v}_2/x]) \in \mathcal{E}[\sigma'] \rho \}$
$\mathcal{V}[\mu\alpha. \sigma] \rho$	$\stackrel{\text{def}}{=} \{ (0, \mathbf{v}_1, \mathbf{v}_2) \} \cup \{ (k+1, \text{fold}_{\rho_1(\mu\alpha.\sigma)} \mathbf{v}_1, \text{fold}_{\rho_2(\mu\alpha.\sigma)} \mathbf{v}_2) \mid \forall j < k+1. (j, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma[\mu\alpha. \sigma/\alpha]] \rho \}$
$\mathcal{V}[\alpha] \rho$	$\stackrel{\text{def}}{=} \rho_R(\alpha)$
$\mathcal{V}[\tau_1 + \tau_2] \rho$	$\stackrel{\text{def}}{=} \{ (k, \text{inj}_i \mathbf{v}_1, \text{inj}_i \mathbf{v}_2) \mid i \in \{1, 2\} \wedge (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau_i] \rho \}$
$\mathcal{V}[\langle \tau_1, \dots, \tau_n \rangle] \rho$	$\stackrel{\text{def}}{=} \{ (k, \langle \mathbf{v}_1, \dots, \mathbf{v}_n \rangle, \langle \mathbf{v}'_1, \dots, \mathbf{v}'_n \rangle) \mid \forall i \in \{1 \dots n\}. (k, \mathbf{v}_i, \mathbf{v}'_i) \in \mathcal{V}[\tau_i] \rho \}$
$\mathcal{V}[\forall[\alpha]. \tau \rightarrow \mathbf{E} \tau_{\text{exn}} \tau'] \rho$	$\stackrel{\text{def}}{=} \{ (k, \lambda[\alpha](x : \rho_1(\tau)). e_1, \lambda[\alpha](x : \rho_2(\tau')). e_2) \mid \forall \tau_1, \tau_2, R \in \text{Rel}[\tau_1, \tau_2]. \forall j \leq k. \forall (j, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau] \rho[\alpha \mapsto (\tau_1, \tau_2, R)]. (j, e_1[\tau_1/\alpha][\mathbf{v}_1/x], e_2[\tau_2/\alpha][\mathbf{v}_2/x]) \in \mathcal{E}[\mathbf{E} \tau_{\text{exn}} \tau'] \rho[\alpha \mapsto (\tau_1, \tau_2, R)] \}$
$\mathcal{V}[\mu\alpha. \tau] \rho$	$\stackrel{\text{def}}{=} \{ (0, \mathbf{v}_1, \mathbf{v}_2) \} \cup \{ (k+1, \text{fold}_{\rho_1(\mu\alpha.\tau)} \mathbf{v}_1, \text{fold}_{\rho_2(\mu\alpha.\tau)} \mathbf{v}_2) \mid \forall j < k+1. (j, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau[\mu\alpha. \tau/\alpha]] \rho \}$
$\mathcal{V}[\mathbf{0}] \rho$	$\stackrel{\text{def}}{=} \emptyset$
$\mathcal{V}[\exists\alpha. \tau] \rho$	$\stackrel{\text{def}}{=} \{ (k, \text{pack}(\tau_1, \mathbf{v}_1), \text{pack}(\tau_2, \mathbf{v}_2)) \mid \exists R \in \text{Rel}[\tau_1, \tau_2]. (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau] \rho[\alpha \mapsto (\tau_1, \tau_2, R)] \}$
$\mathcal{R}[\sigma] \rho$	$\stackrel{\text{def}}{=} \mathcal{V}[\sigma] \rho$
$\mathcal{R}[\mathbf{E} \tau_{\text{exn}} \tau] \rho$	$\stackrel{\text{def}}{=} \{ (k, \text{return } \mathbf{v}_1, \text{return } \mathbf{v}_2) \mid (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau] \rho \} \cup \{ (k, \text{raise } \mathbf{v}_1, \text{raise } \mathbf{v}_2) \mid (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\tau_{\text{exn}}] \rho \}$
$\mathcal{E}[\theta] \rho$	$\stackrel{\text{def}}{=} \{ (k, e_1, e_2) \mid \forall K_1, K_2. (k, K_1, K_2) \in \mathcal{K}[\theta] \rho \implies (k, K_1[e_1], K_2[e_2]) \in \mathcal{O} \}$
$\mathcal{K}[\theta] \rho$	$\stackrel{\text{def}}{=} \{ (k, K_1, K_2) \mid \forall j \leq k, r_1, r_2. (j, r_1, r_2) \in \mathcal{R}[\theta] \rho \implies (j, K_1[r_1], K_2[r_2]) \in \mathcal{O} \}$
$\mathcal{D}[\cdot]$	$\stackrel{\text{def}}{=} \{ \emptyset \}$
$\mathcal{D}[\Delta, \alpha]$	$\stackrel{\text{def}}{=} \{ \rho[\alpha \mapsto (\tau_1, \tau_2, R)] \mid \cdot \vdash \tau_1 \wedge \cdot \vdash \tau_2 \wedge \rho \in \mathcal{D}[\Delta] \wedge R \in \text{Rel}[\tau_1, \tau_2] \}$
$\mathcal{G}[\cdot] \rho$	$\stackrel{\text{def}}{=} \{ (k, \emptyset) \mid k \in \mathbb{N} \}$
$\mathcal{G}[\Gamma, x : \tau] \rho$	$\stackrel{\text{def}}{=} \{ (k, \gamma[x \mapsto (v_1, v_2)]) \mid (k, \gamma) \in \mathcal{G}[\Gamma] \rho \wedge (k, v_1, v_2) \in \mathcal{V}[\tau] \rho \}$

Figure 7. λ^{ST} : Logical Relation for Closed Terms

$\Delta; \Gamma \vdash e_1 \approx_{\mathcal{E}}^{\text{log}} e_2 : \theta \stackrel{\text{def}}{=} \Delta; \Gamma \vdash e_1 : \theta \wedge \Delta; \Gamma \vdash e_2 : \theta \wedge \forall k \geq 0. \forall \rho, \gamma. \rho \in \mathcal{D}[\Delta] \wedge (k, \gamma) \in \mathcal{G}[\Gamma] \rho \implies (k, \rho_1(\gamma_1(e_1)), \rho_2(\gamma_2(e_2))) \in \mathcal{E}[\theta] \rho$
$\Delta; \Gamma \vdash v_1 \approx_{\mathcal{V}}^{\text{log}} v_2 : \tau \stackrel{\text{def}}{=} \dots \implies (k, \rho_1(\gamma_1(e_1)), \rho_2(\gamma_2(e_2))) \in \mathcal{V}[\tau] \rho$

Figure 8. λ^{ST} : Logical Relation for Open Terms

Semantics Preservation To model the cross-language contextual equivalence we defined in §3.1, we define a “cross-language” logical relation using our multi-language logical relation. The cross-language relation, defined in Figure 9, relates source terms of type σ to target terms of type $\sigma^\div$, and similarly source values of type σ to target values of type σ^+ . The relation is defined from the perspective of the source language and places boundaries on the target language; again, this arbitrary choice is justified by Corollary 3.3.

We then prove the semantics preservation theorem, restated below, by using the logical relation instead of contextual equivalence. Unwrapping the definition, the theorem says that the syntactic closure conversion translation produces equivalent results as placing the multi-language boundaries.

Theorem 1.1 (Semantics Preservation)

If $\Gamma \vdash e : \sigma \rightsquigarrow_e e$, then $\Gamma \vdash e \stackrel{\text{S}}{\simeq} e : \sigma$.

$\mathcal{V}^+[\sigma] \stackrel{\text{def}}{=} \{ (k, \mathbf{v}_1, \mathbf{v}_2) \mid \exists \mathbf{v}_2. \sigma \text{ST } \mathbf{v}_2 \mapsto^* \mathbf{v}_2 \wedge (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma] \emptyset \}$
$\mathcal{E}^\div[\sigma] \stackrel{\text{def}}{=} \{ (k, \mathbf{e}, \mathbf{e}) \mid (k, \mathbf{e}, \sigma \text{ST } \mathbf{e}) \in \mathcal{E}[\sigma] \emptyset \}$
$\mathcal{G}^+[\cdot] \stackrel{\text{def}}{=} \{ (k, \emptyset, \emptyset) \mid k \in \mathbb{N} \}$
$\mathcal{G}^+[\Gamma, x : \sigma] \stackrel{\text{def}}{=} \{ (k, \gamma[x \mapsto \mathbf{v}], \gamma[x \mapsto \mathbf{v}]) \mid (k, \gamma, \gamma) \in \mathcal{G}^+[\Gamma] \wedge (k, \mathbf{v}, \mathbf{v}) \in \mathcal{V}^+[\sigma] \}$
$\Gamma \vdash v \approx_+ v : \sigma \stackrel{\text{def}}{=} v \in \lambda^{\text{S}} \wedge v \in \lambda^{\text{T}} \wedge \Gamma \vdash v : \sigma \wedge \cdot; \Gamma^+ \vdash v : \sigma^+ \wedge \forall (k, \gamma, \gamma) \in \mathcal{G}^+[\Gamma]. (k, \gamma(\mathbf{v}), \gamma(\mathbf{v})) \in \mathcal{V}^+[\sigma]$
$\Gamma \vdash e \approx_\div e : \sigma \stackrel{\text{def}}{=} e \in \lambda^{\text{S}} \wedge e \in \lambda^{\text{T}} \wedge \Gamma \vdash e : \sigma \wedge \cdot; \Gamma^+ \vdash e : \sigma^\div \wedge \forall (k, \gamma, \gamma) \in \mathcal{G}^+[\Gamma]. (k, \gamma(\mathbf{e}), \gamma(\mathbf{e})) \in \mathcal{E}^\div[\sigma]$

Figure 9. Cross-Language Logical Relation for λ^{ST}

Proof Sketch It suffices to show $\Gamma \vdash e \approx_\div e : \sigma$ by Theorem 4.3. The proof is by induction on the translation judgment. Most cases follow trivially by the induction hypothesis. The case for application requires some expected parametric reasoning about the closure environment. $\square$

Equivalence Reflection Equivalence reflection is a simple corollary of semantics preservation, after noting that the compositional nature of the translation allows us to easily define closure-conversion for contexts:

Lemma 4.5 (Context Translation)

If $\vdash C : (\Gamma \vdash \sigma) \Rightarrow (\Gamma' \vdash \sigma')$, $\Gamma \vdash e : \sigma$ and $\Gamma \vdash e : \sigma \rightsquigarrow_e e$,
then there exists C such that $\Gamma' \vdash C[e] : \sigma \rightsquigarrow_e C[e]$.
Furthermore, if $\Gamma \vdash e' : \sigma$ and $\Gamma \vdash e' : \sigma \rightsquigarrow_e e'$,
then $\Gamma' \vdash C[e'] : \sigma \rightsquigarrow_e C[e']$.

Theorem 1.2 (Equivalence Reflection)

If $\Gamma \vdash e_1 : \sigma \rightsquigarrow_e e_1$, $\Gamma \vdash e_2 : \sigma \rightsquigarrow_e e_2$ and $\vdash \Gamma^+ \vdash e_1 \approx_T^{ctx} e_2 : \sigma^\dagger$,
then $\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma$.

Proof Let $C \in \lambda^S$ be an appropriately typed context. We need to show that $C[e_1] \Downarrow C[e_2]$. First by Lemma 4.5, we compile $\cdot \vdash C[e_1] : \sigma' \rightsquigarrow_e C[e_1]$ and $\cdot \vdash C[e_2] : \sigma' \rightsquigarrow_e C[e_2]$. Then by Theorem 1.1, $C[e_1] \Downarrow C[e_1]$ and $C[e_2] \Downarrow C[e_2]$. Finally since $\vdash \Gamma^+ \vdash e_1 \approx_T^{ctx} e_2 : \sigma^\dagger$, $C[e_1] \Downarrow C[e_2]$ giving:

$$C[e_1] \Downarrow C[e_1] \Downarrow C[e_2] \Downarrow C[e_2]. \quad \square$$

5. Back-Translation and Equivalence

Preservation

To prove equivalence preservation, it is sufficient to be able to “back-translate” a target context to an equivalent source context. To see why, suppose that $\Gamma \vdash e_1 : \sigma \rightsquigarrow_e e_1$, $\Gamma \vdash e_2 : \sigma \rightsquigarrow_e e_2$ and $\Gamma \vdash e_1 \approx_S^{ctx} e_2 : \sigma$ and that we want to show $\vdash \Gamma^+ \vdash e_1 \approx_T^{ctx} e_2 : \sigma^\dagger$. Given a target context C such that $C[e_1] \Downarrow$, we want to come up with a source context C such that $C[e_1] \Downarrow C[e_1]$ and $C[e_2] \Downarrow C[e_2]$. Since e_1 and e_2 are equivalent we know that $C[e_1] \Downarrow C[e_2]$, so we can conclude that $C[e_2] \Downarrow$.

However, there are several differences between the source and target languages that make this difficult. First, the target has exceptions while the source does not. Second, the target has polymorphic functions, while the source language has a simple type system.

While the target language has features that are untypeable in the source language, we can still encode the target language into the source “loosely”.

In fact, we first back-translate to a single type of “untyped” terms that encodes all values of the source and target languages: unit, sum, product, function and recursive type:

$$\begin{aligned} U &\stackrel{\text{def}}{=} \mu\alpha. 1 + (\alpha + \alpha) + (\alpha \times \alpha) + (\alpha \rightarrow R(\alpha)) + \alpha \\ R(\alpha) &\stackrel{\text{def}}{=} \alpha + \alpha \end{aligned}$$

The type $R(\alpha)$ is the parameterized type of *results* in the target language, i.e., either a successful or exceptional value. We write R to mean the closed version of this type, $R(\alpha)[U/\alpha]$.

This encoding allows us to back-translate target terms into the source language, but it is not enough to ensure full abstraction. Note that back-translating a term of type $\sigma^\dagger$ produces a term of type R , not a term of type σ .

However, R is not just a type of untyped terms, but a *universal* type for the source language as in Scott [36] (see Longley [23] for a more accessible introduction). By this we mean every type σ can be embedded in R via a function we denote $\text{EMBED}(\sigma)$ and there is a partial function that selects the elements of σ from R which we denote $\text{PROJECT}(\sigma)$. These two form a *retract* which means that $\text{PROJECT}(\sigma)(\text{EMBED}(\sigma)x) \approx_S^{ctx} x$. This suffices to make our encoding at translation type “tight” though our translation at other types was “loose” (i.e., not even equivalence preserving).

This property is perfectly analogous with one case of Lemma 3.2, that is $\sigma ST TS^\sigma x \approx_S^{ctx} x$, and indeed our full abstraction proof relies on a formal correspondence between the embedding-projection pairs and the multi-language boundaries.

More precisely, when we want to prove that equivalence is preserved, we start with a target context C interacting with the translation of a source term e , which is equivalent under the multi-language semantics to a target context interacting with the source term, mediated by the language boundaries (ST, TS) .

$$\begin{aligned} &\cdot \xrightarrow{\text{def}} \cdot \\ (\Gamma, x : \sigma) &\xrightarrow{\text{def}} \Gamma \rightarrow, x : \sigma \\ (\Gamma, y : \tau) &\xrightarrow{\text{def}} \Gamma \rightarrow, y : U \\ \text{If } \Delta; \Gamma \vdash e : \sigma, \text{ then } \Gamma \rightarrow \vdash e \rightarrow : \sigma & \\ x &\xrightarrow{\text{def}} x \\ (\text{let } x = e \text{ in } e') &\xrightarrow{\text{def}} \text{let } x = e \rightarrow \text{ in } e' \rightarrow \\ &\vdots \\ (\sigma ST e) &\xrightarrow{\text{def}} \text{let } x = e \rightarrow \text{ in } \text{PROJECT}(\sigma) x \\ \text{If } \Delta; \Gamma \vdash v : \tau, \text{ then } \Gamma \rightarrow \vdash v \rightarrow : U & \\ y &\xrightarrow{\text{def}} y \\ \text{inj}_i v &\xrightarrow{\text{def}} \text{IN}(i, v \rightarrow) \\ \langle \rangle &\xrightarrow{\text{def}} \text{UNIT} \\ \langle v_1, \dots, v_n \rangle &\xrightarrow{\text{def}} \text{CONS}(v_1 \rightarrow, \langle \dots, v_n \rangle \rightarrow) \\ (\lambda[\alpha](y : \tau). e) &\xrightarrow{\text{def}} \text{LAMBDA}(\lambda(y : U). e \rightarrow) \\ (\text{fold}_{\mu\alpha. \tau} v) &\xrightarrow{\text{def}} \text{FOLD}(v \rightarrow) \\ (\text{pack}(\tau', v)) &\xrightarrow{\text{def}} v \rightarrow \\ \text{If } \Delta; \Gamma \vdash e : \theta, \text{ then } \Gamma \rightarrow \vdash e \rightarrow : R & \\ (\text{return } v) &\xrightarrow{\text{def}} \text{RETURN}(v \rightarrow) \\ (\text{raise } v) &\xrightarrow{\text{def}} \text{RAISE}(v \rightarrow) \\ (v.i) &\xrightarrow{\text{def}} \text{let } x = \text{PRJ}(i, v \rightarrow) \text{ in } \text{RETURN}(x) \\ (v_1 v_2) &\xrightarrow{\text{def}} \text{let } x = \text{TOFUN}(v_1 \rightarrow) \text{ in } x v_2 \rightarrow \\ (\text{unfold } v) &\xrightarrow{\text{def}} \text{let } x = \text{TOFOLD}(v \rightarrow) \text{ in } \text{RETURN}(x) \\ (\text{unpack } (\alpha, y) = v \text{ in } e) &\xrightarrow{\text{def}} \text{let } y = v \rightarrow \text{ in } e \\ (\text{handle } e \text{ with } (x_1.e_1) \dots (x_2.e_2)) &\xrightarrow{\text{def}} \text{let } x_r = e \rightarrow \text{ in case } x_r \text{ of } \\ &\quad x_1.e_1 \rightarrow \dots x_2.e_2 \rightarrow \\ &\quad \vdots \\ (TS^\sigma e) &\xrightarrow{\text{def}} \text{let } x = e \rightarrow \text{ in } \text{EMBED}(\sigma) x \end{aligned}$$

Figure 10. Back-Translation (excerpts)

We want a source context C whose interaction with e is the same as that of C . To this end we backtranslate the target context to a source context at universal type and replace the language boundaries (TS, ST) with the universal type boundaries $(\text{EMBED}(\sigma), \text{PROJECT}(\sigma))$.

The crucial detail is then a proof of correspondence between these two types of boundaries: Formally stated in Lemma 5.4 and Lemma 5.5, we essentially prove that if a context C with hole of type $\sigma^\dagger$ and return type $\sigma'^\dagger$ is back translated to C with hole of type R and return type $\sigma'^\dagger$, then

$$\text{PROJECT}(\sigma') (C[\text{EMBED}(\sigma) []]) \approx \sigma' ST (C[TS^\sigma []]).$$

5.1 Back-Translation, Formally

We present a portion of the back-translation in Figure 10. To enhance readability, we define metafunctions that construct terms of universal type via the appropriate application of fold_U and inj_i . For instance, $\text{UNIT} = \text{fold}_U(\text{inj}_1 \langle \rangle)$. The elimination metafunctions and $\text{PROJECT}()$ also handle failure since the sub-language is untyped. The metafunction $\text{TOLHS}()$ is one such metafunction which returns the expected value if back-translation succeeds, and fails otherwise. For example, in the metafunction $\text{TOFUN}(e')$, if the term e' represents a function, it is applied, otherwise it results in failure:

$$\text{let } x = \text{TOFUN}(\text{LAMBDA}(\lambda(y : U). e)) \text{ in } x v \mapsto^* e[v/y]$$

$\text{let } x = \text{TOFUN}(\text{UNIT}) \text{ in } x \vdash \mapsto^* \mathcal{U}$

where $\mathcal{U}$ represents undefined behavior. Since our back-translation only produces well-formed terms, our development is independent of what $\mathcal{U}$ means.¹

We back-translate source terms to themselves, target terms to constructors for the universal type, and boundaries to $\text{EMBED}()$ and $\text{PROJECT}()$. While some types, like sums, can be directly simulated in the source, others must be encoded. For instance, tuples are back-translated to cons-lists. Crucially, all polymorphism is erased, using dynamic typing to encode the behavior of the polymorphic terms of the target language. Finally, we preserve the **fold** constructor for formal convenience.

In Figure 11, we define $\text{EMBED}()$ and $\text{PROJECT}()$ for each type. The code is fairly ugly since we are programming in a bare-bones calculus in monadic normal form, but its meaning is straightforward and follows the multi-language boundaries.

At the top level we have $\text{EMBED}(\sigma)$ and $\text{PROJECT}(\sigma)$ which translate between $\mathbf{R}$ and σ . These deal with the exception mechanisms of the target language. As in the multi-language boundary, no exception should be thrown from target to source or vice-versa. Thus, $\text{EMBED}(\sigma)$ uses the $\text{RETURN}()$ function to return a successful value and $\text{PROJECT}(\sigma)$ uses the $\text{TOLHS}()$ function to handle any possible exceptions. If it is a successful value of type $\mathbf{R}$ then this is further projected but if it is an exception then $\text{TOLHS}()$ diverges.

These functions are defined in terms of $\text{EMBED}(\delta, \sigma)$ and $\text{PROJECT}(\delta, \sigma)$, which mediate between σ and $\mathbf{U}$, the type of values in the target language.

Since we have higher-order recursive types, we must define $\text{EMBED}(\delta, \sigma)$ and $\text{PROJECT}(\delta, \sigma)$ mutually recursively. This is the purpose of the δ parameter. It maps free type variables to the recursive type that they represent and the recursive binding to a thunk of the embedding-projection pair for that type. We use the notation δ_σ to denote the associated closing type substitution and δ_x to denote the map from type variables to term variables of the associated closed type. In the definition, σ 's free type variables are always in δ 's domain, and the definition of $\text{EMBED}()$, $\text{PROJECT}()$ is structurally decreasing in σ . The key recursive step is in Figure 12 which uses a standard call-by-value fixed point combinator (definition elided) to simultaneously define embedding and projection at recursive type.

For function types, $\text{EMBED}(\delta, \sigma_1 \rightarrow \sigma_2)$ packs a function with an empty environment and $\text{PROJECT}(\delta, \sigma_1 \rightarrow \sigma_2)$ unpacks a closure and applies the function to its environment.

Pairs are straightforward, handling the fact that a two-tuple in the target is a cons list of length two. Sums and the unit type are even more straightforward so we elide them.

5.2 Correctness of Back-Translation

The correctness theorem for the back-translation mirrors Theorem 1.1 (semantics preservation)—that is, back-translation of $\mathbf{e}$ is equivalent to placing appropriate boundaries around $\mathbf{e}$. However the back-translation thus far uses the universal type, since we back-translate all target language terms and not just those that are of translation type. So we get a more refined theorem, one that says that a target term and its back-translation are indistinguishable in their *interactions with source programs*.

Programs of universal type interact with programs of source type mediated by $\text{EMBED}()$ and $\text{PROJECT}()$ just as multi-language interactions are mediated by boundaries. For closed programs, we prove that they satisfy the following equivalence:

¹ In our proofs, we model failure via divergence since it can be defined uniformly for any type.

δ	$::= \emptyset \mid \delta[\alpha \mapsto \sigma, x]$
$\emptyset_\Gamma$	$\stackrel{\text{def}}{=} \cdot$
$(\delta[\alpha \mapsto \sigma, x])_\Gamma$	$\stackrel{\text{def}}{=} \delta_\Gamma, x : 1 \rightarrow ((\sigma \rightarrow \mathbf{U}) \times (\mathbf{U} \rightarrow \sigma))$
$\cdot \vdash \text{EMBED}(\sigma) : \sigma \rightarrow \mathbf{R}$	$\delta_\Gamma \vdash \text{EMBED}(\delta, \sigma) : \delta_\sigma(\sigma) \rightarrow \mathbf{U}$
$\text{EMBED}(\sigma)$	$\stackrel{\text{def}}{=} \lambda(x : \sigma). \text{let } x_u = \text{EMBED}(\emptyset, \sigma) \times \text{in } \text{RETURN}(x_u)$
$\text{PROJECT}(\sigma)$	$\stackrel{\text{def}}{=} \lambda(x_r : \mathbf{R}). \text{let } x_u = \text{TOLHS}(x_r) \text{ in } \text{PROJECT}(\emptyset, \sigma) \times_u$
$\text{EMBED}(\delta, \alpha)$	$\stackrel{\text{def}}{=} \lambda(x : \delta_\sigma(\alpha)). \text{let } x_{ep} = \delta_x(\alpha) \langle \rangle \text{ in } \text{let } x_{embed} = \pi_1 x_{ep} \text{ in } x_{embed} \times$
$\text{EMBED}(\delta, \sigma_1 \times \sigma_2)$	$\stackrel{\text{def}}{=} \lambda(x : \delta_\sigma(\sigma_1 \times \sigma_2)). \text{let } x_1 = \pi_1 x \text{ in } \text{let } x_2 = \pi_2 x \text{ in } \text{let } x'_1 = \text{EMBED}(\delta, \sigma_1) x_1 \text{ in } \text{let } x'_2 = \text{EMBED}(\delta, \sigma_2) x_2 \text{ in } \text{CONS}(x'_1, \text{CONS}(x'_2, \text{UNIT}))$
$\text{EMBED}(\delta, \sigma_1 \rightarrow \sigma_2)$	$\stackrel{\text{def}}{=} \lambda(x_f : \delta_\sigma(\sigma_1 \rightarrow \sigma_2)). \text{let } x'_f = \lambda(x_u : \mathbf{U}). \text{let } x'_u = \text{PRJ}(2, x_u) \text{ in } \text{let } x = \text{PROJECT}(\delta, \sigma_1) x'_u \text{ in } \text{let } y = x_f \times \text{in } \text{let } x''_u = \text{EMBED}(\delta, \sigma_2) y \text{ in } \text{RETURN}(x''_u) \text{ in } \text{CONS}(x'_f, \text{CONS}(\text{UNIT}, \text{UNIT}))$
$\text{EMBED}(\delta, \mu\alpha. \sigma)$	$\stackrel{\text{def}}{=} \lambda(x : \delta_\sigma(\mu\alpha. \sigma)). \text{let } x_{ep} = \text{EP}(\delta, \mu\alpha. \sigma) \langle \rangle \text{ in } \text{let } x_{embed} = \pi_1 x_{ep} \text{ in } x_{embed} \times$
$\cdot \vdash \text{PROJECT}(\sigma) : \mathbf{R} \rightarrow \sigma$	$\delta_\Gamma \vdash \text{PROJECT}(\delta, \sigma) : \mathbf{U} \rightarrow \delta_\sigma(\sigma)$
$\text{PROJECT}(\delta, \alpha)$	$\stackrel{\text{def}}{=} \lambda(x_u : \mathbf{U}). \text{let } x = \delta_x(\alpha) \langle \rangle \text{ in } \text{let } x_f = \pi_2 x \text{ in } x_f \times$
$\text{PROJECT}(\delta, \sigma_1 \times \sigma_2)$	$\stackrel{\text{def}}{=} \lambda(x_u : \mathbf{U}). \text{let } x = \text{TOPAIR}(x_u) \text{ in } \text{let } x_1 = \pi_1 x \text{ in } \text{let } x'_1 = \text{PROJECT}(\delta, \sigma_1) x_1 \text{ in } \text{let } y = \pi_2 x \text{ in } \text{let } y' = \text{TOPAIR}(y) \text{ in } \text{let } x_2 = \pi_1 y' \text{ in } \text{let } x'_2 = \text{PROJECT}(\delta, \sigma_2) x_2 \text{ in } \langle x'_1, x'_2 \rangle$
$\text{PROJECT}(\delta, \sigma_1 \rightarrow \sigma_2)$	$\stackrel{\text{def}}{=} \lambda(x_u : \mathbf{U}). \text{let } x'_u = \text{TOPAIR}(x_u) \text{ in } \text{let } x_f = \text{PRJ}(1, x'_u) \text{ in } \text{let } x_{env} = \text{PRJ}(2, x'_u) \text{ in } \lambda(y : \delta_\sigma(\sigma_1)). \text{let } y_u = \text{EMBED}(\delta, \sigma_1) y \text{ in } \text{let } x = \text{CONS}(x_{env}, \text{CONS}(y_u, \text{UNIT})) \text{ in } \text{let } x_r = x_f \times \text{in } \text{let } x''_u = \text{TOLHS}(x_r) \text{ in } \text{PROJECT}(\delta, \sigma_2) x''_u$
$\text{PROJECT}(\delta, \mu\alpha. \sigma)$	$\stackrel{\text{def}}{=} \lambda(x_u : \mathbf{U}). \text{let } x = \text{EP}(\delta, \mu\alpha. \sigma) \langle \rangle \text{ in } \text{let } x_f = \pi_2 x \text{ in } x_f \times_u$

Figure 11. Embedding and Projection Functions

Definition 5.1 (Universal Type Equivalence for Closed Programs)

If $\cdot \vdash \mathbf{e} : \sigma^\div$ and $\cdot \vdash \mathbf{e}_u : \mathbf{R}$, then

$\cdot \vdash \mathbf{e}_u \stackrel{\text{def}}{\sim} \mathbf{e} : \sigma \stackrel{\text{def}}{=} \cdot \vdash \text{PROJECT}(\sigma) \circ \mathbf{e}_u \approx_{\text{ST}}^{\text{ctx}} \sigma \mathcal{T} \mathbf{e} : \sigma$

where $\text{PROJECT}(\sigma) \circ \mathbf{e}_u \stackrel{\text{def}}{=} \text{let } x_u = \mathbf{e}_u \text{ in } \text{PROJECT}(\sigma) x_u$.

$$\boxed{\delta_\Gamma \vdash \underline{\text{EP}}(\delta, \mu\alpha. \sigma) : 1 \rightarrow ((\delta_\sigma(\mu\alpha. \sigma) \rightarrow \mathbf{U}) \times (\mathbf{U} \rightarrow \delta_\sigma(\mu\alpha. \sigma)))}$$

$$\begin{aligned}
&\underline{\text{EP}}(\delta, \mu\alpha. \sigma) \stackrel{\text{def}}{=} \\
&\quad \underline{\text{FIX}}_1 \rightarrow ((\delta_\sigma(\mu\alpha. \sigma) \rightarrow \mathbf{U}) \times (\mathbf{U} \rightarrow \delta_\sigma(\mu\alpha. \sigma))) \\
&\quad \lambda(x_{\mu\alpha. \sigma} : 1 \rightarrow ((\delta_\sigma(\mu\alpha. \sigma) \rightarrow \mathbf{U}) \times (\mathbf{U} \rightarrow \delta_\sigma(\mu\alpha. \sigma)))) \\
&\quad \lambda(x_{\text{unit}} : 1). \\
&\quad \text{let } x_{\text{embed}} = \\
&\quad \quad \lambda(x : \delta_\sigma(\mu\alpha. \sigma)). \\
&\quad \quad \text{let } y = \text{unfold } x \text{ in} \\
&\quad \quad \text{let } y_u = \underline{\text{EMBED}}(\delta[\alpha \mapsto \mu\alpha. \sigma, x_{\mu\alpha. \sigma}], \sigma) y \text{ in} \\
&\quad \quad \underline{\text{FOLD}}(y_u) \\
&\quad \text{in let } x_{\text{project}} = \\
&\quad \quad \lambda(x_u : \mathbf{U}). \\
&\quad \quad \text{let } y_u = \underline{\text{TOFOLD}}(x_u) \text{ in} \\
&\quad \quad \text{let } y = \underline{\text{PROJECT}}(\delta[\alpha \mapsto \mu\alpha. \sigma, x_{\mu\alpha. \sigma}], \sigma) y_u \text{ in} \\
&\quad \quad \text{fold}_{\mu\alpha. \sigma} y \\
&\quad \text{in } \langle x_{\text{embed}}, x_{\text{project}} \rangle
\end{aligned}$$

Figure 12. Embedding-Projection Pair at Recursive Type

To define equivalence for open programs, we need to define how to *embed* all the free variables from source to universal type just as we had to extend the boundary to act on the free variables. Note that here we use the version of `EMBED()` whose return type is $\mathbf{U}$. The asymmetry here is because we are in a call-by-value language. We then arrive at the following definition for open programs

Definition 5.2 (Universal Type Equivalence)

If $\vdash \Gamma^+ \vdash e : \sigma^\div$ and $\Gamma^+ \vdash e_u : R$, and $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$, we define

$$\begin{aligned}
&\Gamma \vdash e_u \stackrel{\mathcal{U}}{\simeq} e : \sigma \stackrel{\text{def}}{=} \\
&\quad \vdash \Gamma \vdash \underline{\text{PROJECT}}(\sigma) \circ e_u \circ \underline{\text{EMBED}}(\Gamma) \approx_{\text{ST}}^{\text{ctx}} \sigma \text{ST} e \text{ST} \Gamma \\
&\quad \text{where } e_u \circ \underline{\text{EMBED}}(\Gamma) \stackrel{\text{def}}{=} \text{let } x_{u1} = \underline{\text{EMBED}}(\emptyset, \sigma_1) x_1 \text{ in} \\
&\quad \quad \vdots \\
&\quad \quad \text{let } x_n = \underline{\text{EMBED}}(\emptyset, \sigma_n) x_n \text{ in } e_u.
\end{aligned}$$

In order to have an inductive invariant strong enough to prove back-translation is correct, we define a cross-language logical relation to formalize when a term of universal type represents an *arbitrary* target term, rather than just a target term of translation type.

Crucially, we cannot reuse the original logical relation’s case for the universal type because each target type corresponds to a different “subset” of the universal type. Furthermore, since these subsets are not syntactic types but purely logical, we have to restrict the continuations linked against to be well-behaved.

In Figure 13 we define a step-indexed biorthogonal cross-language logical relation that relates the typed target language to the untyped embedding in the source language. As usual, we lift the relation on closed terms to a relation on open terms via closing substitution. However, recall that we define the back-translation for the entire multi-language and not just the target language. While the logical relation on closed terms only relates source terms to target terms, we must extend the relation on open terms to relate source terms to source terms. The definition for source terms merely appeals to the multi-language logical relation defined in §4.

The fundamental property of this logical relation, Theorem 5.3, gives us that every term is related to its back-translation. As in §4, we require a form of bridge lemma, Lemma 5.4, which also connects our two notions of logical equivalence, namely the embedding interpretation and the multi-language interoperability semantics.

Theorem 5.3 (Fundamental Property)

1. If $\Delta; \Gamma \vdash e : \sigma$ then $\Delta; \Gamma \vdash e^\rightarrow \approx_{\mathcal{U}}^{\log} e : \sigma$.
2. If $\Delta; \Gamma \vdash e : \theta$ then $\Delta; \Gamma \vdash e^\rightarrow \approx_{\mathcal{U}}^{\log} e : \theta$.

Lemma 5.4 (Interpret = Interoperate)

For brevity, some cases are elided

1. If $\Delta; \Gamma \vdash v_u \approx_{\mathcal{U}}^{\log} v : \sigma^+$, then for any appropriate k, ρ^U, γ^U
 $\exists v_1, v_2, \underline{\text{PROJECT}}(\sigma) \circ \rho^U(\gamma^U(v_u)) \mapsto^{n_1} v_1$,
 $\sigma \text{ST} \text{return } \rho^U(\gamma^U(v)) \mapsto^{n_2} v_2$ and either $n_1, n_2 > k$ or
 $(k, v_1, v_2) \in \mathcal{V}_{\llbracket \sigma \rrbracket}^U \rho^U$.
2. If $\Delta; \Gamma \vdash e_u \approx_{\mathcal{U}}^{\log} e : \sigma^\div$, then
 $\Delta; \Gamma \vdash \underline{\text{PROJECT}}(\sigma) \circ e_u \approx_{\mathcal{U}}^{\log} \sigma \text{ST} e : \sigma$.
3. If $\Delta; \Gamma \vdash e \approx_{\mathcal{U}}^{\log} e' : \sigma$, then
 $\Delta; \Gamma \vdash \underline{\text{EMBED}}(\sigma) \circ e \approx_{\mathcal{U}}^{\log} \text{TS } \sigma e' : \sigma^\div$.

As a corollary, we have the correctness theorem for our back-translation.

Corollary 5.5 (Back-Translation Correctness)

If $\vdash \Gamma^+ \vdash e : \sigma^\div$, then $\Gamma \vdash e^\rightarrow \stackrel{\mathcal{U}}{\simeq} e : \sigma$.

Proof Let $(k, \gamma) \in \mathcal{G}[\llbracket \Gamma \rrbracket] \emptyset$ and $v_i = \gamma(x_i)$.

Using Lemma 5.4, for each $x_i : \sigma_i \in \Gamma$, we define v_i, v_{ui} to be the unique values satisfying:

$$\begin{aligned}
&\text{TS } \sigma_i v_i \mapsto^* \text{return } v_i \\
&\underline{\text{EMBED}}(\emptyset, \sigma_i) v_i \mapsto^* v_{ui}
\end{aligned}$$

If these take more than k steps we’re done, otherwise by Lemma 5.4, we define γ^U such that $\gamma^U(x_{ui}) = v_{ui}$ and $\gamma^U(x_i) = v_i$, so $(k, \gamma^U) \in \mathcal{G}[\llbracket \Gamma^+ \rrbracket]$. Then we have

$$\gamma(\sigma \text{ST} e \text{ST} \Gamma) \mapsto^* \sigma \text{ST} \gamma^U(e)$$

$$\gamma(\underline{\text{PROJECT}}(\sigma) \circ e^\rightarrow \circ \underline{\text{EMBED}}(\Gamma)) \mapsto^* \underline{\text{PROJECT}}(\sigma) \circ \gamma^U(e^\rightarrow)$$

and the result follows from Theorem 5.3, Lemma 5.4 and closure of relatedness under anti-reduction. $\square$

5.3 Equivalence Preservation

We prove equivalence preservation as a consequence of the correctness of the back-translation.

In fact, we present two proofs, one *syntactic* based on translation of contexts and the other *algebraic* based on rewriting. The latter proof especially highlights the benefits of multi-language semantics, crucially using “mixed” equations to derive single-language equivalences.

For our first proof, we extend the back-translation to *contexts*. We exploit the fact that our back-translation is compositional, *i.e.*, we have that $(C[e])^\rightarrow = C^\rightarrow[e^\rightarrow]$, which simplifies our first equivalence preservation proof.

Theorem 1.3 (Equivalence Preservation)

If $\Gamma \vdash e_1 : \sigma \rightsquigarrow_e e_1$, $\Gamma \vdash e_2 : \sigma \rightsquigarrow_e e_2$ and $\Gamma \vdash e_1 \approx_5^{\text{ctx}} e_2 : \sigma$, then $\vdash \Gamma^+ \vdash e_1 \approx_{\text{ST}}^{\text{ctx}} e_2 : \sigma^\div$.

Proof Suppose $\mathbf{C} \in \lambda^T$. We proceed in two steps. First we use the compiler semantics preservation theorem to construct a multi-language context for the source terms:

$$\vdash \Gamma^+ \vdash e_1 \approx_{\text{ST}}^{\text{ctx}} \text{TS } \sigma e_1 \text{ST} : \sigma^\div.$$

Let $C = \mathbf{C}[\text{TS } \sigma [\cdot] \text{ST}]$. Then we have $\mathbf{C}[e_1] \Downarrow C[e_1]$ and similarly $\mathbf{C}[e_2] \Downarrow C[e_2]$.

We now have a context for the source terms, however the source terms are equivalent in the *source language* but C is a *multi-language* context. To arrive at a fully source context, we back-translate:

$$(C[e_1])^\rightarrow = C^\rightarrow[e_1^\rightarrow] = C^\rightarrow[e_1]$$

and

$$(C[e_2])^\rightarrow = C^\rightarrow[e_2].$$

Then by Theorem 5.3 we get $C[e_1] \Downarrow C^\rightarrow[e_1]$ and $C[e_2] \Downarrow C^\rightarrow[e_2]$. Finally, since $\Gamma \vdash e_1 \approx_5^{\text{ctx}} e_2 : \sigma$ and $C^\rightarrow \in \lambda^S$, $C^\rightarrow[e_1] \Downarrow C^\rightarrow[e_2]$. This completes the chain of ifs giving us $\mathbf{C}[e_1] \Downarrow \mathbf{C}[e_2]$:
 $\mathbf{C}[e_1] \Downarrow C[e_1] \Downarrow C^\rightarrow[e_1] \Downarrow C^\rightarrow[e_2] \Downarrow C[e_2] \Downarrow \mathbf{C}[e_2]$. $\square$

Note: all relations are restricted to well-typed members.

$$\begin{aligned}
\text{Rel}^U[\tau] &\stackrel{\text{def}}{=} \{R \mid \forall (k, \mathbf{v}, \mathbf{v}') \in R. \cdot \vdash \mathbf{v} : \mathbf{U} \wedge \cdot \vdash \mathbf{v}' : \tau \wedge \forall j < k. (j, \mathbf{v}, \mathbf{v}') \in R\} \\
\mathcal{V}^U[\alpha] \rho^U &\stackrel{\text{def}}{=} \rho_R^U(\alpha) \\
\mathcal{V}^U[\tau_1 + \tau_2] \rho^U &\stackrel{\text{def}}{=} \{(k, \text{IN}(i, \mathbf{v}_u), \text{inj}_i \mathbf{v}) \mid (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau_i] \rho^U\} \\
\mathcal{V}^U[\langle \rangle] \rho^U &\stackrel{\text{def}}{=} \{(k, \text{UNIT}, \langle \rangle)\} \\
\mathcal{V}^U[\langle \tau_1, \tau_2, \dots, \tau_n \rangle] \rho^U &\stackrel{\text{def}}{=} \{(k, \text{CONS}(\mathbf{v}_u, \mathbf{v}'_u), \langle \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n \rangle) \mid (k, \mathbf{v}_u, \mathbf{v}_1) \in \mathcal{V}^U[\tau_1] \rho^U \wedge (k, \mathbf{v}'_u, \langle \mathbf{v}_2, \dots, \mathbf{v}_n \rangle) \in \mathcal{V}^U[\langle \tau_2, \dots, \tau_n \rangle] \rho^U\} \\
\mathcal{V}^U[\forall[\alpha]. \tau \rightarrow \theta] \rho^U &\stackrel{\text{def}}{=} \left\{ (k, \text{LAMBDA}(\lambda(x_u : \mathbf{U}). \mathbf{e}_u), \lambda[\alpha](\mathbf{x} : \rho^U(\tau)). \mathbf{e}) \mid \forall \tau', R \in \text{Rel}^U[\tau'], j \leq k, (j, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau] \rho^{U'} \right. \\
&\quad \left. (j, \mathbf{e}_u[\mathbf{v}_u/x_u], \mathbf{e}[\tau'/\alpha][\mathbf{v}/\mathbf{x}]) \in \mathcal{E}^U[\theta] \rho^{U'} \text{ where } \rho^{U'} = \rho^U[\alpha \mapsto \tau', R] \right\} \\
\mathcal{V}^U[\mu\alpha. \tau] \rho^U &\stackrel{\text{def}}{=} \{(0, \mathbf{v}_u, \mathbf{v})\} \cup \{(k+1, \text{FOLD}(\mathbf{v}_u), \text{fold}_{\rho^U(\mu\alpha. \tau)} \mathbf{v}) \mid (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau[\mu\alpha. \tau/\alpha]] \rho^U\} \\
\mathcal{V}^U[0] \rho^U &\stackrel{\text{def}}{=} \emptyset \\
\mathcal{V}^U[\exists\alpha. \tau] \rho^U &\stackrel{\text{def}}{=} \{(k, \mathbf{v}_u, \text{pack}(\tau', \mathbf{v})) \mid \exists R \in \text{Rel}^U[\tau']. (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau] \rho^U[\alpha \mapsto \tau', R]\} \\
\mathcal{R}^U[\text{E } \tau_{\text{exn}} \tau] \rho^U &\stackrel{\text{def}}{=} \{(k, \text{RETURN}(\mathbf{v}_u), \text{return } \mathbf{v}) \mid (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau] \rho^U\} \cup \{(k, \text{RAISE}(\mathbf{v}_u), \text{raise } \mathbf{v}) \mid (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau_{\text{exn}}] \rho^U\} \\
\mathcal{E}^U[\theta] \rho^U &\stackrel{\text{def}}{=} \{(k, \mathbf{e}_u, \mathbf{e}) \mid \forall j \leq k, K_1, K_2. (j, K_1, K_2) \in \mathcal{K}^U[\theta] \rho^U \implies (j, K_1[\mathbf{e}_u], K_2[\mathbf{e}]) \in \mathcal{O}\} \\
\mathcal{K}^U[\theta] \rho^U &\stackrel{\text{def}}{=} \{(k, K_1, K_2) \mid \forall j \leq k, \mathbf{v}_u, \mathbf{r}. (j, \mathbf{v}_u, \mathbf{r}) \in \mathcal{R}^U[\theta] \rho^U \implies (j, K_1[\mathbf{v}_u], K_2[\mathbf{r}]) \in \mathcal{O}\} \\
\mathcal{D}^U[\cdot] &\stackrel{\text{def}}{=} \{\emptyset\} \\
\mathcal{D}^U[\Delta, \alpha] &\stackrel{\text{def}}{=} \{\rho^U[\alpha \mapsto \tau, R] \mid \cdot \vdash \tau \wedge \rho^U \in \mathcal{D}[\Delta] \wedge R \in \text{Rel}^U[\tau]\} \\
\mathcal{G}^U[\cdot] \rho^U &\stackrel{\text{def}}{=} \{(k, \emptyset) \mid k \in \mathbb{N}\} \\
\mathcal{G}^U[\Gamma, \mathbf{x} : \sigma] \rho^U &\stackrel{\text{def}}{=} \{(k, \gamma^U[\mathbf{x} \mapsto \mathbf{v}_1, \mathbf{v}_2]) \mid (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U \wedge (k, \mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[\sigma] \emptyset\} \\
\mathcal{G}^U[\Gamma, \mathbf{x} : \tau] \rho^U &\stackrel{\text{def}}{=} \{(k, \gamma^U[\mathbf{x} \mapsto \mathbf{v}][x_u \mapsto \mathbf{v}_u]) \mid (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U \wedge (k, \mathbf{v}_u, \mathbf{v}) \in \mathcal{V}^U[\tau] \rho^U\} \\
\Delta; \Gamma \vdash \mathbf{v}' \stackrel{\log}{\sim}_{\mathcal{V}^U} \mathbf{v} : \sigma &\stackrel{\text{def}}{=} \mathbf{v}' \in \lambda^S \wedge \forall \rho^U \in \mathcal{D}^U[\Delta], (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U. (k, \gamma^U(\mathbf{v}'), \rho^U(\gamma^U(\mathbf{v}))) \in \mathcal{V}[\sigma] \emptyset \\
\Delta; \Gamma \vdash \mathbf{e}' \stackrel{\log}{\sim}_{\mathcal{E}^U} \mathbf{e} : \sigma &\stackrel{\text{def}}{=} \mathbf{e}' \in \lambda^S \wedge \forall \rho^U \in \mathcal{D}^U[\Delta], (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U. (k, \gamma^U(\mathbf{e}'), \rho^U(\gamma^U(\mathbf{e}))) \in \mathcal{E}[\sigma] \emptyset \\
\Delta; \Gamma \vdash \mathbf{v}_u \stackrel{\log}{\sim}_{\mathcal{R}^U} \mathbf{r} : \theta &\stackrel{\text{def}}{=} \mathbf{v}_u \in \lambda^S \wedge \forall \rho^U \in \mathcal{D}^U[\Delta], (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U. (k, \gamma^U(\mathbf{v}_u), \rho^U(\gamma^U(\mathbf{r}))) \in \mathcal{R}^U[\theta] \rho^U \\
\Delta; \Gamma \vdash \mathbf{e}_u \stackrel{\log}{\sim}_{\mathcal{E}^U} \mathbf{e} : \theta &\stackrel{\text{def}}{=} \mathbf{e}_u \in \lambda^S \wedge \forall \rho^U \in \mathcal{D}^U[\Delta], (k, \gamma^U) \in \mathcal{G}^U[\Gamma] \rho^U. (k, \gamma^U(\mathbf{e}_u), \rho^U(\gamma^U(\mathbf{e}))) \in \mathcal{E}^U[\theta] \rho^U
\end{aligned}$$

Figure 13. Universal Type Logical Relation

Our second proof is more abstract, but also succinct. First, we define some simplifying notation. For $\Gamma \vdash \mathbf{e} : \sigma$ we define

$$[\mathbf{e}] \stackrel{\text{def}}{=} \mathbf{e} \text{ where } \Gamma \vdash \mathbf{e} : \sigma \rightsquigarrow_e \mathbf{e}$$

For $\cdot; \Gamma^+ \vdash \mathbf{e} : \sigma^\dagger$,

$$(\mathbf{e}) \stackrel{\text{def}}{=} \text{PROJECT}(\sigma) \circ \mathbf{e}^{\rightarrow} \circ \text{EMBED}(\Gamma).$$

These theorems give us that, up to equivalence, the back translation is a post-inverse to the compiler.

Lemma 5.7 (There and Back is Identity)

For all $\Gamma \vdash \mathbf{e} : \sigma$,

$$\Gamma \vdash \mathbf{e} \approx_S^{ctx} \llbracket \mathbf{e} \rrbracket : \sigma$$

Proof By Theorem 1.1, Lemma 5.5, and Lemma 3.2 (eliding obvious types):

$$\begin{aligned}
\llbracket \mathbf{e} \rrbracket &\approx_{ST}^{ctx} ST \llbracket \mathbf{e} \rrbracket TS \\
&\approx_{ST}^{ctx} ST TS \mathbf{e} ST TS \\
&\approx_{ST}^{ctx} \mathbf{e}
\end{aligned}$$

Now we have $\llbracket \mathbf{e} \rrbracket \approx_{ST}^{ctx} \mathbf{e}$, which is sufficient to prove that $\llbracket \mathbf{e} \rrbracket \approx_S^{ctx} \mathbf{e}$, since $\approx_{ST}^{ctx} \subset \approx_S^{ctx}$. $\square$

Note the simplicity of this proof using multi-language semantics. By providing a framework in which source and target programs can be freely mixed, we enable direct equational reasoning between programs and their compilation. Since multi-language equivalence is trivially sound with respect to source or target equivalence, a chain of equations using intermediate multi-language terms can be used to prove purely source or purely target programs are equivalent.

As an obvious corollary, the composition of compiler and back-translation is fully abstract. However we only use that it preserves equivalence, since we can already easily prove that it reflects equivalence.

Lemma 5.8 (There and Back is Fully Abstract)

$$\Gamma \vdash \mathbf{e} \approx_S^{ctx} \mathbf{e}' : \sigma \text{ iff } \Gamma \vdash \llbracket \mathbf{e} \rrbracket \approx_S^{ctx} \llbracket \mathbf{e}' \rrbracket : \sigma$$

Finally, we have that the back-translation is equivalence reflecting.

Lemma 5.9 (Back Translation is Equivalence Reflecting)

If $\Gamma \vdash (\mathbf{e}) \approx_S^{ctx} \llbracket \mathbf{e}' \rrbracket : \sigma$ then $\cdot; \Gamma^+ \vdash \mathbf{e} \approx_T^{ctx} \mathbf{e}' : \sigma^\dagger$.

Proof Analogous to proof of Theorem 1.2, using Lemma 5.5. $\square$

We conclude with the proof of equivalence preservation as an obvious consequence. This part of the proof is highly general and completely elementary. Its content is that if the composition of two translations ($\llbracket \cdot \rrbracket$ and $\llbracket \cdot \rrbracket$) is equivalence preserving and the second ($\llbracket \cdot \rrbracket$) is equivalence reflecting, then the first ($\llbracket \cdot \rrbracket$) is equivalence preserving.

Proof Assume $\Gamma \vdash \mathbf{e} \approx_S^{ctx} \mathbf{e}' : \sigma$.

By Lemma 5.8, $\Gamma \vdash \llbracket \mathbf{e} \rrbracket \approx_S^{ctx} \llbracket \mathbf{e}' \rrbracket : \sigma$.

By Lemma 5.9, $\cdot; \Gamma^+ \vdash \mathbf{e} \approx_T^{ctx} \llbracket \mathbf{e}' \rrbracket : \sigma^\dagger$. $\square$

6. Related Work

Interest in fully abstract translation has a long history [25, 34, 17, 40, 37, 35, 6, 7, 2, 18, 3, 16, 29, 30], with several authors pointing out that failures of full abstraction can be turned into security exploits [1, 21, 14].

Until recently, most results on fully abstract translations have strategically used source and target languages that are in close correspondence, either (1) picking syntactically identical source and target languages [6, 16] or (2) adding problematic target-language features to the source language [34, 35, 16]. This can be attributed to the lack of back-translation techniques for settings where the

target is more expressive than the source. Our *Universal Embedding* technique offers a solution to this problem, discussed further in §7.

Protection via Types Ahmed and Blume [6] (AB08) were first to prove that typed closure conversion of System F with recursive types is fully abstract. Since their source and target languages coincide, they define type-directed wrappers in the same language to coerce terms from type τ to τ^+ , and vice versa, instead of a back-translation. Our work, with different source and target languages (even without exceptions), provides one significant advantage over theirs: our target language’s type system guarantees that function bodies are closed, as they should be after closure conversion.

Later, Ahmed and Blume [7] (AB11) developed a typed CPS translation from simply typed λ -calculus to System F and proved it fully abstract using a multi-language semantics and *Back-Translation by Partial Evaluation*. To back-translate a term of translation type, the technique uses partial evaluation to eliminate all subterms of non-translation type.

Bowman and Ahmed [11] define a noninterference-preserving translation from DCC to F_ω , both terminating languages. Like contextual equivalence, noninterference is a relational property. Their proof, like AB11, uses back-translation by partial evaluation, but also requires a complex well-foundedness argument as their back-translation is not inductively defined.

Glew [17] presents closure conversion for an object calculus and proves it fully abstract. He notes, however, that object closure conversion is simpler than functional closure conversion. In particular, the latter can be encoded as the composition of (1) encoding functions as objects, (2) object closure conversion, and (3) an object encoding. Glew only shows that step (2) is fully abstract. Hence, for full abstraction of functional closure conversion, one would also need to prove encodings (1) and (3) fully abstract.

Protection via Dynamic Checks Fournet et al. [16] prove full abstraction of a translation from monomorphic ML-like language to js^* . While their actual compiler implementation produces JavaScript, their proof of full abstraction is for the translation to js^* , which is not actual JavaScript, but rather an encoding of JavaScript in their source language. Their proof technique, like AB08’s, relies on the source and target languages being identical, hence they added exceptions and fatal errors to their source. Our work is the first full abstraction result where the target language contains exceptions but the source does not. Like AB08, Fournet et al. use type-based invariants to prove full abstraction. However, unlike AB08, they use type-directed wrappers to *dynamically* protect translated components from target contexts.

Agten, Patrignani, et al. [3, 29, 30] give a fully abstract translation from an object-based language—Java Jr. [19] extended with exceptions—to a *protected module architecture*—an untyped assembly language extended with a hardware-supported isolation mechanism. Their translation relies on hardware-provided memory access protection to ensure full abstraction. They use fully abstract trace semantics for Java Jr. [19] so they can prove their theorems in terms of trace equivalence.

Universal Type In Scott [36], types are given manifest computational meaning by being defined as idempotent functions on a universal domain, or equivalently, the embedding-projection pair (retract) of the domain that compose to that idempotent. In this paper, we construct a universal type in the source language and define explicit retracts for all *source* types. We never explicitly prove this property, but it is a simple consequence of Lemma 5.4 and Lemma 4.2.

The types of the target language, as it is encoded by back-translation into the source, need *not* be retracts for our technique to work. For example, we do nothing to enforce parametricity conditions on terms of polymorphic type, and it seems unlikely

that it could be enforced in our source language. In this way our technique uses a mix of computational types given by retracts and purely logical types as encoded by our back-translation logical relation.

Embedded Interpreters and Multi-Language Semantics The concept of type-indexed embedding-projection pairs for mediating between a typed source language and an embedded untyped language draws directly from Benton [8] and Ramsey [33], where they use this as an *implementation technique* for embedding a scripting language communicating with a host language. Here we use their implementation technique as a formal tool to prove secure compilation.

Multi-language semantics was first developed in Matthews and Findler [24] as a general technique to describe interacting languages operationally. While the connection to embedded interpreters was noted from the start, this paper is the first to prove a precise correspondence between the two techniques.

Approximate Back-Translation In recent work, Devriese et al. [12] showed how to prove full abstraction of a translation from simply typed λ -calculus with recursive functions (but not recursive types) to the untyped λ -calculus. Since their target is untyped, they back-translate to a universal type in the source language. However, the difficulty in that work is in the fact that the appropriate universal type cannot be constructed in the source (due to a lack of recursive types). Thus they show that the full universal type is not actually necessary, as long as arbitrarily large approximations of it exist. We discuss a potential use of this in §7.

In comparison, our work—done independently of theirs—tackles a nontrivial program transformation, and demonstrates universal embedding for a *more richly typed* target language with exceptions. Also, since the universal type can be constructed in our source language, we are able to separate the core of the technique—universal embedding—from the use of approximate back-translation, making our presentation more elementary. We view their independent discovery and use of (approximate) back-translation to universal type in a fairly different setting to be evidence for the generality of the technique.

Compositional Compiler Correctness While the central focus of this paper has been on using *full abstraction* to ensure security of compiled components with respect to source-language reasoning, we have also proved *compositional compiler correctness* (Theorem 1.1), a topic that has been the focus of several recent papers. Note that none of the results discussed below tackle fully abstract compilation.

Perconti and Ahmed [31] showed how to compile components correctly while allowing interoperability with target components that have no equivalent in the source. This allows linking with components that are inexpressible in the source, but requires extra-language reasoning to achieve security. Like us, Perconti and Ahmed [31] use multi-language semantics to specify interoperability between source and target components, and prove compositional compiler correctness of a multi-pass compiler from System F to a low-level language with mutable references. They allow linking with any well-typed target component of translation type, but unlike our work, their type translation does not restrict the target language to make only the same observations as the source language.

Stewart et al. [39] use *interaction semantics*, which provide an abstract specification of interoperability between source and target components, to prove compositional correctness of the CompCert C compiler. This work allows linking with any target component that respects restrictions imposed by CompCert’s memory model via a specified interaction semantics.

Neis et al. [27] prove compositional correctness of a compiler for an ML-like language to an assembly-like language. They allow linking with only those target components that are related to some source

component via a *parametric inter-language simulation* (PILS), a relation that specifies equivalence between source-language and target-language code. In practical terms, this means that code produced by a PILS-verified compiler can only be linked with target code produced either by the same compiler, or by a different verified compiler from the same source language to the same target language using the same PILS specification.

Recent work by Kang et al. [20] demonstrates that if one makes the pragmatic choice to restrict linking to only those components produced by the same verified compiler—i.e., supporting separate compilation—that significantly eases the proof effort. In particular, they prove correctness of SepCompCert which permits linking only with other components produced by their modified version of CompCert.

7. Discussion

With much recent work, we can begin to classify different techniques for proving full abstraction. Some techniques are subsumed by others, but offer a trade-off of power vs. simplicity. Here we focus on comparing back-translation-based techniques, though there are other techniques for proving full abstraction, such as those that make use fully abstract trace semantics [3, 29, 30] as mentioned in §6.

The first technique may be called *Back Translation by Partial Evaluation* as used in Ahmed and Blume [7], Bowman and Ahmed [11], Shikuma and Igarashi [38]. Back-translation by partial evaluation relies on the fact that in some languages, if a term has translation type, then any uses of non-translation type in a subterm are *inessential* and can therefore be eliminated by partial evaluation. However, this is not a realistic assumption for non-terminating languages or languages with information hiding. For instance, a diverging program may do arbitrary computation using non-translation type that cannot be eliminated by partial evaluation. Furthermore, this technique alone will not work when both languages have state or existential types since programs of target type can use values of non-translation type in their closure or as the existential witness. For this reason, back translation by partial evaluation has only been applied to purely functional, normalizing languages. However, it has the benefit of being fairly systematic when it is applicable.

The second technique may be called *Back Translation by Embedding* as used in Ahmed and Blume [6], Fournet et al. [16], Devriese et al. [12], and this paper.

This technique translates all target code, not just those at translation type. The translation type code is not translated directly to the source type, but instead wrapped with a semantic boundary.

In Ahmed and Blume [6] and Fournet et al. [16], the target and source are the same so the back-translation is the identity and the boundaries are called “wrappers”. These wrappers are witness to a type *isomorphism* and consequently the translation is fully abstract. We say that this back-translation is *precise* because the embedding of the target language is into isomorphic types.

In this paper, we back-translate to a universal type and the boundaries are *retractions*, that is $\text{PROJECT}(\sigma) \text{ EMBED}(\sigma) e \approx_S^{ctx} e$. We say our back-translation is *over-approximating* in that it embeds the target language at types which include many more behaviors (full untyped lambda calculus), but due to the boundaries the code is only run on “good” values, i.e., those that represent source values.

Finally, in Devriese et al. [12], they back-translate an untyped language to a simply typed language *without* recursive types, so they cannot construct the universal type as in this work. However, they can construct arbitrarily large approximations to the universal type, and a family of increasingly precise approximations to it. They show that due to the finitary nature of observation, i.e. that observing termination only takes a finite number of steps, they can find a *large enough* approximation to show that equivalence is preserved for any particular program and context. We say their back-translation is

under-approximating since it embeds the target language at types which include only a subset of the behaviors of the target.

It seems likely to be fruitful to combine the last two embedding approaches, to construct a sequence of *under-approximations* to an *over-approximating* type (like a universal type).

Candidates for such a technique would include embedding languages into intermediate languages with richer type systems where a precise universal type for the source cannot be constructed. For instance, proving that embedding a non-terminating System F-like language into an $F\omega$ -like or dependently-typed language would require simulating the more flexible typed terms of the target language.

What should a universal type for an F-like language be? In addition to being able to encode all of the base types it should also encode all terms of abstract types that are in scope, so we need not one U , but an indexed family:

$$U_\Delta \cong (U_\Delta \rightarrow U_\Delta) + (\forall \alpha. U_{\Delta, \alpha}) + \Sigma(\Delta)$$

where $\Sigma(\alpha_1, \dots, \alpha_n) = \alpha_1 + \dots + \alpha_n$. Since the universal type includes polymorphic types, it depends on the definition of the universal type with one type parameter. Clearly recursive types alone are not enough to construct such a type so it seems that a sequence of under-approximations is necessary.

This source-target combination may seem strange but this scenario would arise if one were to construct a reusable compiler backend for fully abstract compilation. Since the backend would have to have typing features of the most richly typed source language it supports, then using it as a fully abstract backend would require proving that embedding less richly typed languages into more richly typed languages is secure.

We now have a spectrum of embedding techniques, and as a practical matter one should use the simplest one necessary for the job. For instance, a precise encoding is fairly simple to prove fully abstract using wrappers. However, if the target has programs untypeable *precisely* in the source then an over-approximation seems necessary. Finally, if the source is too weak to express the intended back-translation type, then a sequence of approximations may be used.

Effects Our technique should be effective in back-translating other well-bracketed/delimited effects. For instance, if compiling a purely functional language to a stateful language that ensures that the compiled terms only link with pure code, target terms can be back-translated to store-passing style, changing the result type to something like $R = S \rightarrow (S \times U)$. Similarly, if the target contains delimited continuations (protecting source code with delimiters), it can be back-translated to continuation-passing style, with $R = (U \rightarrow U) \rightarrow U$.

Information Hiding The technique should easily handle information hiding, such as state and existential types. As an example if the source and target have existential types, then a term of type $\exists \alpha. \sigma^+$ must be back-translated to a source term of existential type. An existential package may instantiate α with a target type not in the image of the back-translation, meaning we must be able to simulate the behavior of a target term of non-translation type using a source term. With our technique this is simple, because the witness type can be back-translated to dynamic type. On the other hand, back-translation by partial evaluation, as explained above, could not be used.

Acknowledgments

This research was supported in part by the National Science Foundation (grants CCF-1422133 and CCF-1453796).

References

- [1] M. Abadi. Protection in programming-language translations. In *International Colloquium on Automata, Languages and Programming*

- (ICALP), pages 868–883, 1998.
- [2] M. Abadi and G. Plotkin. On protection by layout randomization. *ACM Transactions on Information and Systems Security*, 15(2), July 2012.
 - [3] P. Agten, R. Strackx, B. Jacobs, and F. Piessens. Secure compilation to modern processors. In *Computer Security Foundations Symposium (CSF)*, pages 171–185, 2012.
 - [4] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *European Symposium on Programming (ESOP)*, pages 69–83, Mar. 2006.
 - [5] A. Ahmed. Verified Compilers for a Multi-Language World. In T. Ball, R. Bodik, S. Krishnamurthi, B. S. Lerner, and G. Morrisett, editors, *1st Summit on Advances in Programming Languages (SNAPL 2015)*, volume 32 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15–31, 2015. URL <http://drops.dagstuhl.de/opus/volltexte/2015/5013>.
 - [6] A. Ahmed and M. Blume. Typed closure conversion preserves observational equivalence. In *International Conference on Functional Programming (ICFP)*, Victoria, British Columbia, Canada, pages 157–168, Sept. 2008.
 - [7] A. Ahmed and M. Blume. An equivalence-preserving CPS translation via multi-language semantics. In *International Conference on Functional Programming (ICFP)*, Tokyo, Japan, pages 431–444, Sept. 2011.
 - [8] N. Benton. Embedded interpreters. *Journal of Functional Programming*, 15(04):503–542, 2005.
 - [9] N. Benton and A. Kennedy. Exceptional syntax. *Journal of Functional Programming*, 11(04):395–410, 2001.
 - [10] N. Benton, A. Kennedy, and G. Russell. Compiling Standard ML to Java bytecodes. In *International Conference on Functional Programming (ICFP)*, Baltimore, Maryland, USA, pages 129–140, 1998. URL <http://doi.acm.org/10.1145/289423.289435>.
 - [11] W. J. Bowman and A. Ahmed. Noninterference for free. In *International Conference on Functional Programming (ICFP)*, Vancouver, British Columbia, Canada, 2015. URL <http://dx.doi.org/10.1145/2784731.2784733>.
 - [12] D. Devriese, M. Patrignani, and F. Piessens. Fully-abstract compilation by approximate back-translation. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg, Florida, page To appear, 2016.
 - [13] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming*, 22(4&5):477–528, 2012.
 - [14] V. D’Silva, M. Payer, and D. Song. The correctness-security gap in compiler optimization. In *Language-theoretic Security IEEE Security and Privacy Workshop (LangSec)*, 2015.
 - [15] M. Felleisen and R. Hieb. A revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.*, 103(2):235–271, 1992.
 - [16] C. Fournet, N. Swamy, J. Chen, P.-E. Dagand, P.-Y. Strub, and B. Livshits. Fully abstract compilation to JavaScript. In *ACM Symposium on Principles of Programming Languages (POPL)*, Rome, Italy, pages 371–384, 2013.
 - [17] N. Glew. Object closure conversion. In *Higher-Order Operational Techniques in Semantics (HOOTS ’99)*, Sept. 1999.
 - [18] R. Jagadeesan, C. Pitcher, J. Rathke, and J. Riely. Local memory via layout randomization. In *Computer Security Foundations Symposium (CSF)*, pages 161–174, 2011.
 - [19] A. Jeffrey. A fully abstract semantics for a concurrent functional language with monadic types. In *IEEE Symposium on Logic in Computer Science (LICS)*, San Diego, California, 1995.
 - [20] J. Kang, Y. Kim, C.-K. Hur, D. Dreyer, and V. Vafeiadis. Lightweight verification of separate compilation. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg, Florida, pages 178–190. ACM, 2016.
 - [21] A. Kennedy. Securing the .NET programming model. *Theoretical Computer Science*, 364(3):311–317, 2006.
 - [22] J.-L. Krivine. Classical logic, storage operators and second-order lambda-calculus. *Annals of Pure and Applied Logic*, 68(1):53–78, 1994.
 - [23] J. R. Longley. Universal types and what they are good for. In *Domain theory, logic and computation*, pages 25–63. Springer, 2003.
 - [24] J. Matthews and R. B. Findler. Operational semantics for multi-language programs. In *ACM Symposium on Principles of Programming Languages (POPL)*, Nice, France, pages 3–10, Jan. 2007.
 - [25] A. Meyer and J. G. Riecke. Continuations may be unreasonable. In *Conf. on LISP and functional programming, LFP ’88*, pages 63–71, 1988.
 - [26] Y. Minamide, G. Morrisett, and R. Harper. Typed closure conversion. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, pages 271–283, Jan. 1996.
 - [27] G. Neis, C.-K. Hur, J.-O. Kaiser, C. McLaughlin, D. Dreyer, and V. Vafeiadis. Pilsner: A compositionally verified compiler for a higher-order imperative language. In *International Conference on Functional Programming (ICFP)*, Vancouver, British Columbia, Canada, Aug. 2015.
 - [28] M. New, W. J. Bowman, and A. Ahmed. Fully abstract compilation via universal embedding (technical appendix). July 2016. URL <https://williamjbowman.com/resources/fabcc-techrpt.pdf>.
 - [29] M. Patrignani, D. Clarke, and F. Piessens. Secure compilation of object-oriented components to protected module architectures. In *Proceedings of the 11th Asian Symposium on Programming Languages and Systems (APLAS)*, Melbourne, Australia, Dec. 2013.
 - [30] M. Patrignani, P. Agten, R. Strackx, B. Jacobs, D. Clarke, and F. Piessens. Secure compilation to protected module architectures. *ACM Transactions on Programming Languages and Systems*, 37(2):6:1–6:50, Apr. 2015.
 - [31] J. T. Perconti and A. Ahmed. Verifying an open compiler using multi-language semantics. In *European Symposium on Programming (ESOP)*, Apr. 2014.
 - [32] A. M. Pitts and I. D. Stark. Operational reasoning for functions with local state. *Higher order operational techniques in semantics*, pages 227–273, 1998.
 - [33] N. Ramsey. Embedding an interpreted language using higher-order functions and types. *Journal of Functional Programming*, 21(06):585–615, 2011.
 - [34] J. G. Riecke. Fully abstract translations between functional languages. In *ACM Symposium on Principles of Programming Languages (POPL)*, Orlando, Florida, pages 245–254, 1991.
 - [35] S. B. Sanjabi and C.-H. L. Ong. Fully abstract semantics of additive aspects by translation. In *Proceedings of the 6th international conference on Aspect-oriented software development (AOSD)*, pages 135–148, 2007.
 - [36] D. Scott. Data types as lattices. *Siam Journal on computing*, 5(3):522–587, 1976.
 - [37] N. Shikuma and A. Igarashi. Proving noninterference by a fully complete translation to the simply typed λ -calculus. In *Proceedings of the 11th Asian computing science conference on Advances in computer science: secure software and related issues*, pages 301–315. Springer-Verlag, 2007.
 - [38] N. Shikuma and A. Igarashi. Proving noninterference by a fully complete translation to the simply typed λ -calculus. *Logical Methods in Computer Science*, 4(3:10):1–31, 2008.
 - [39] G. Stewart, L. Beringer, S. Cuellar, and A. W. Appel. Compositional compcert. In *ACM Symposium on Principles of Programming Languages (POPL)*, Mumbai, India, 2015.
 - [40] S. Tse and S. Zdancewic. Translating dependency into parametricity. In *International Conference on Functional Programming (ICFP)*, Snowbird, Utah, pages 115–125. ACM, 2004.
 - [41] P. Wang, S. Cuellar, and A. Chlipala. Compiler verification meets cross-language linking via data abstraction. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, Oct. 2014.