

Operational Semantics for Multi-Language Programs

Jacob Matthews Robert Bruce Findler

University of Chicago

{jacobm, robbey}@cs.uchicago.edu

Abstract

Inter-language interoperability is big business, as the success of Microsoft's .NET and COM and Sun's JVM show. Programming language designers are designing programming languages that reflect that fact — SML#, Mondrian, and Scala, to name just a few examples, all treat interoperability with other languages as a central design feature. Still, current multi-language research tends not to focus on the semantics of interoperation features, but only on how to implement them efficiently. In this paper, we take first steps toward higher-level models of interoperating systems. Our technique abstracts away the low-level details of interoperability like garbage collection and representation coherence, and lets us focus on semantic properties like type-safety and observable equivalence.

Beyond giving simple expressive models that are natural compositions of single-language models, our studies have uncovered several interesting facts about interoperability. For example, higher-order contracts naturally emerge as the glue to ensure that inter-operating languages respect each other's type systems. While we present our results in an abstract setting, they shed light on real multi-language systems and tools such as the JNI, SWIG, and Haskell's stable pointers.

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

General Terms Languages, theory

Keywords Interoperability, multi-language systems, operational semantics

1. Introduction

A modern large-scale software system is likely written in a variety of languages: its core might be written in Java, while it has specialized system interaction routines written in C and a web-based user interface written in PHP. And even academic languages have caught multi-language programming fever, perhaps due to temptingly large numbers of pre-existing libraries written in other languages. This has prompted language implementors to target COM [19, 45], Java Virtual Machine bytecode [7, 28, 36], and most recently Microsoft's Common Language Runtime [8, 34, 39]. Furthermore, where foreign function interfaces have historically been used in practice to allow high-level safe languages to call libraries written in low-level unsafe languages like C (as was the motivation for the popular wrapper generator SWIG [5]), these new foreign

function interfaces are built to allow high-level, safe languages to interoperate with other high-level, safe languages, such as Python with Scheme [35] and Lua with OCaml [42].

Since these embeddings are driven by practical concerns, the research that accompanies them rightly focuses on the bits and bytes of interoperability — how to represent data in memory, how to call a foreign function efficiently, and so on. But an important theoretical problem arises, independent of these implementation-level concerns: how can we reason formally about multi-language programs? This is a particularly important question for systems that involve typed languages, because we have to show that the embeddings respect their constituents' type systems.

In this paper we present a simple method for giving operational semantics to multi-language systems. Our models are rich enough to support a wide variety of multi-language embedding strategies, and powerful enough that we have been able to use them for type soundness and contextual equivalence proofs. Our technique is based on simple constructs we call *boundaries*, cross-language casts that regulate both control flow and value conversion between languages. We introduce boundaries through series of operational semantics in which we combine a simple ML-like language with a simple Scheme-like language.

In section 2, we introduce those two constituent languages formally and connect them using a primitive embedding where values in one language are opaque to the other. In section 3, we enrich that embedding so that boundaries use type information to convert one language's values into their counterparts in the other, and we show that this embedding naturally leads to higher-order contracts. Section 4 shows a surprising relationship between the expressive power of these two embeddings, and section 5 shows how the system scales beyond purely type-directed conversion.

2. The lump embedding

To begin, we pick two languages, give them formal models, and then tie those formal models together. In the interest of focusing on interoperation rather than the special features of particular languages, we have chosen two simple calculi: an extended model of the untyped call-by-value lambda calculus, which we use as a stand-in for Scheme, and an extended model of the simply-typed lambda calculus, which we use as a stand-in for ML (though it more closely resembles Plotkin's PCF without fixpoint operators [40]). Figure 1 presents these languages in an abstract manner that we instantiate multiple ways to model different forms of interoperability. One goal of this section is to explain that figure's peculiarities, but for now notice that aside from unusual subscripts and font choices, the two language models look pretty much as they would in a normal Felleisen-and-Hieb-style presentation [16].

To make the preparation more concrete, as we explain our presentation of the core models we also simultaneously develop our first interoperation model, which we call the lump embedding. In the lump embedding, ML values can appear in Scheme and Scheme values can appear in ML. However, ML treats Scheme val-

This technical report accompanies the paper of the same title to appear at the Principles of Programming Languages (POPL) '07 conference. The two papers are identical except that this version includes a technical appendix of proofs left out of the conference version.

value across the boundary (which will be a lump in ML); if $\tau \neq \mathbf{L}$ we are sending an ML value across the boundary (which will be a lump in Scheme).

In these typing judgments, Scheme rules use the same type environment that ML does. We do that to allow ML expressions to refer to variables bound by ML (or vice versa) even if they are separated by a sequence of boundaries. This is necessary to give types to functions that use foreign arguments, as we shall see.

2.3 Operational semantics

We use Felleisen and Hieb-style reduction semantics to specify operational semantics for our systems. In figure 1, we define an evaluation context for ML (\mathbf{E}) and one for Scheme (\mathbf{E}), and we use a third, unspecified evaluation context (\mathcal{E}) to represent the top-level context. In examples in this paper, we will assume $\mathcal{E} = \mathbf{E}$ and thus that the top-level program is written in ML; we could set $\mathcal{E} = \mathbf{E}$ to reverse that assumption. To allow Scheme expressions to evaluate inside ML expressions and vice versa, we define evaluation contexts mutually recursively at boundaries: \mathbf{E} produces $({}^\tau MS \mathbf{E})$ and \mathbf{E} produces $(SM^\tau \mathbf{E})$.

The reduction rules in the core model are all reasonably standard, with a few peculiarities. On the ML side, we allow subtraction of two numbers but floor all results at zero. The Scheme side has a bit more going on dynamically. Since Scheme’s type system does not protect it, we add dynamic checks to every appropriate form that reduce to `wrong` if they receive an illegal value. The reduction rule for `wrong` itself discards the entire program context and aborts the program with an error message.

To combine the languages, we might hope to just merge their sets of reductions together. That does not quite work. For instance, the ML term $(\lambda \mathbf{x} : \iota. \mathbf{x}) (\bar{1} \bar{1})$ would reduce to (`wrong` “Non-procedure”), rather than getting stuck, since Scheme’s reduction for a non-procedure application would apply to $(\bar{1} \bar{1})$. To remedy this, we extend Felleisen and Hieb’s context-sensitive rewriting framework by differentiating $[\]_M$ holes generated by ML evaluation contexts from $[\]_S$ holes generated by Scheme evaluation contexts. Scheme’s rewriting rules only apply to evaluation contexts with $[\]_S$ holes, (and similarly for ML); we indicate that restriction by writing an S (or M) subscript next to the bracket on each evaluation context on the left-hand side of Scheme’s (or ML’s) rewriting rules.¹

With this extension, the example above decomposes into the context $(\lambda \mathbf{x} : \iota. \mathbf{x})[\]_M$ but not $(\lambda \mathbf{x} : \iota. \mathbf{x})[\]_S$. Because the hole is named $[\]_M$ and the Scheme application rule only applies to holes named $[\]_S$, the term remains stuck. However, if instead we had written $(\lambda \mathbf{x} : \iota. \mathbf{x}) ({}^L MS (\bar{1} \bar{1}))$ it would decompose into an ML context with a Scheme hole: $(\lambda \mathbf{x} : \iota. \mathbf{x}) ({}^L MS [\]_S)$ with the erroneous application $(\bar{1} \bar{1})$ inside. Since that redex is in a Scheme hole it would reduce to an error.

To finish the lump embedding, all that remains is to specify the reduction rules and values for the boundaries between languages. If an MS boundary of type \mathbf{L} has a Scheme value inside, then the boundary is an ML value. Similarly, when an SM boundary of a non-lump type has an ML value inside, then it is a Scheme value. However, MS boundaries with a non-lump type that contain Scheme values and SM boundaries of type \mathbf{L} that contain ML values should reduce, since they represent foreign values returning to a native context. We do that by cancelling matching boundaries, as the reduction rules in figure 2 show.

¹We could also have introduced two different notations for Scheme and ML application, but we find it inelegant; it suggests that a multi-language implementation would decide how to evaluate each term by inspecting it, when real systems decide how to evaluate a term based on the language in which the term is being evaluated — *i.e.*, its context. Also, this is the same extension we made in earlier work to model Scheme’s multiple values [32].

ML’s typing rules guarantee that values that appear inside $(SM^\mathbf{L} \mathbf{v})$ expressions will in fact be lump values, so the $SM^\mathbf{L}$ reduction can safely restrict its attention to values of the correct form. Scheme offers no such guarantee, so the rule for eliminating an ${}^\tau MS$ boundary must apply whenever the Scheme expression is a value at all.

These additions give us a precise account of the behavior for lump embedding we described at the beginning of this section. To get a sense of how the calculus works, consider this example:

$$\begin{aligned} & ((\lambda \mathbf{fa} : \mathbf{L} \rightarrow \mathbf{L} \rightarrow \mathbf{L}. ((\mathbf{fa}) ({}^L MS (\lambda \mathbf{x} : \iota. (+ \times \bar{1})))))) \\ & \quad ({}^L MS \bar{3})) \\ & \rightarrow^2 (\lambda \mathbf{f} : \mathbf{L}. \lambda \mathbf{x} : \mathbf{L}. ({}^L MS ((SM^\mathbf{L} \mathbf{f}) (SM^\mathbf{L} \mathbf{x})))) \\ & \rightarrow^2 ({}^L MS ((SM^\mathbf{L} ({}^L MS (\lambda \mathbf{x} : \iota. (+ \times \bar{1})))) (SM^\mathbf{L} ({}^L MS \bar{3})))) \\ & \rightarrow^2 ({}^L MS ((\lambda \mathbf{x} : \iota. (+ \times \bar{1})) \bar{3})) \rightarrow^2 ({}^L MS \bar{4}) \end{aligned}$$

In the initial term of this reduction sequence, we use a “left-left-lambda” encoding of `let` to bind the name `fa` (for “foreign-apply”) to a curried ML function that takes two foreign values, applies the first to the second, and returns the result as another foreign value. The program uses `fa` to apply a Scheme add-one function to the Scheme number $\bar{3}$. In two computation steps, we plug in the Scheme function and its argument into the body of `fa`. In that term there are two instances of $(SM^\mathbf{L} ({}^L MS \mathbf{v}))$ subterms, both of which are cancelled in the next two computation steps. After those cancellations, the term is just a Scheme application of the add-one function to $\bar{3}$, which reduces to the Scheme value $\bar{4}$.

If we try to apply the ML add-one function to the Scheme number $\bar{3}$ instead (and adjust `fa`’s type to make that possible), we will end up with an intermediate term like this:

$$({}^L MS ((SM^{\iota \rightarrow \iota} (\lambda \mathbf{x} : \iota. (+ \times \bar{1}))) \bar{3})) \rightarrow^2 \mathbf{Error}: \text{non-procedure}$$

Here, Scheme tries to apply the ML function directly, which leads to a runtime error since it is illegal for Scheme to apply ML functions. We cannot make the analogous mistake and try to apply a Scheme function in ML, since terms like $({}^L MS (\lambda \mathbf{x} : \iota. (+ \times \bar{1}))) \bar{3}$ are ill-typed.

The formulation of the lump embedding in figure 2 allows us to prove type soundness in the standard way, by establishing preservation and progress lemmas. Notice that because of the way we have combined the two languages, type soundness entails that both languages are type-sound with respect to *their own* type systems — in other words, that both single-language type soundness proofs are special cases of the soundness theorem for the entire system.

Theorem 1 (Lump type soundness). *If $\Gamma \vdash_M e : \tau$, then either $e \rightarrow^* v$, $e \rightarrow^* \mathbf{Error}$: str, or e diverges.*

Proof. (Sketch) A mutually-recursive variation on the standard preservation and progress lemmas. \square

Proofs of this and the other theorems in this paper appear in appendix A.

3. The natural embedding

The lump embedding is a useful starting point, but many multi-language systems offer richer cross-language communication primitives. A more natural way to pass values between our Scheme and ML models, suggested many times in the literature (*e.g.*, [6,37,42]) is to use a type-directed strategy to convert ML numbers to equivalent Scheme numbers and ML functions to equivalent Scheme functions (for some suitable notion of equivalence) and vice versa. We call this the natural embedding.

We can quickly get at the essence of this strategy by extending the core calculi from figure 1, just as we did before to form the lump embedding. Again, we add new syntax and reduction rules to figure 1. In this section we will add ${}^\tau MS_N$ and SM_N^τ boundaries,

adding the subscript N (for “natural”) to distinguish these new boundaries from lump boundaries from section 2.

We will assume we can translate numbers from one language to the other, and give reduction rules for boundary-crossing numbers based on that assumption:

$$\begin{aligned} \mathcal{E}[(SM_N^t \bar{n})_S] &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[(^t MS_N \bar{n})_M] &\rightarrow \mathcal{E}[\bar{n}] \end{aligned}$$

In some multi-language settings, differing byte representations might complicate this task. Worse, some languages may have more expansive notions of numbers than others — for instance, the actual Scheme language treats many different kinds of numbers uniformly (e.g., integers, floating-point numbers, arbitrary precision rationals, and complex numbers are all operated on by the same operators), whereas the actual ML language imposes much more structure on its number representations. More sophisticated versions of the above rules would address these problems straightforwardly.

We must be more careful with procedures, though. We cannot get away with just moving the text of a Scheme procedure into ML or vice versa; aside from the obvious problem that their grammars generate different sets of terms, ML does not even have a reasonable equivalent for every Scheme procedure. Instead, for this embedding we represent a foreign procedure with a proxy. We represent a Scheme procedure in ML at type $\tau_1 \rightarrow \tau_2$ by a new procedure that takes an argument of type τ_1 , converts it to a Scheme equivalent, runs the original Scheme procedure on that value, and then converts the result back to ML at type τ_2 . For example, the scheme function f becomes $(\lambda x : \tau_1. (^t MS_N (f (SM_N^{\tau_1} x))))$ when embedded in ML. Since function arguments flow in the opposite direction from function results, the boundary that converts the argument to the Scheme function must be an $SM_N^{\tau_1}$ boundary, not an $^t MS_N$ boundary.

This would complete the natural embedding, but for one important problem: the system has stuck states, since a boundary might receive a value of an inappropriate shape. Stuck states violate type-soundness, and in an implementation they might correspond to segmentation faults or other undesirable behavior. As it turns out, higher-order contracts [17, 18] arise naturally as the checks required to protect against these stuck states. We show that in the next three sections: first we add dynamic guards directly to boundaries to provide a baseline, then show how to separate them, and finally observe that these separated guards are precisely contracts between ML and Scheme, and that since ML statically guarantees that it always lives up to its contracts, we can eliminate their associated dynamic checks.

3.1 A simple method for adding error transitions

In the lump embedding, we can always make a single, immediate check that would tell us if the value Scheme provided to ML was consistent with the type ML ascribed to it. This is no longer possible, since we cannot know if a Scheme function always produces a value that can be converted to the appropriate type. Still, we can perform an optimistic check that preserves ML’s type safety: when a Scheme value crosses a boundary, we only check its first-order qualities — *i.e.*, whether it is a number or a procedure. If it has the appropriate first-order behavior, we assume the type ascription is correct and perform the conversion, distributing into domain and range conversions as before. If it does not, we immediately signal an error. This method works to catch all errors that would lead to stuck states. Although it only checks first-order properties, the program can only reach a stuck state if a value is used in such a way that it does not have the appropriate first-order properties anyway.

To model this method, rather than adding the SM_N^t and $^t MS_N$ constructs to our core languages from figure 1, we instead add “guarded” versions GSM^t and MSG^t shown in figure 3. These

$$\begin{array}{l} \mathbf{e} = \dots \mid (MSG^t \mathbf{e}) \\ \mathbf{e} = \dots \mid (GSM^t \mathbf{e}) \end{array} \left| \begin{array}{l} \frac{\Gamma \vdash_S \mathbf{e} : \mathbf{TST}}{\Gamma \vdash_M (MSG^t \mathbf{e}) : \tau} \\ \frac{\Gamma \vdash_M \mathbf{e} : \tau}{\Gamma \vdash_S (GSM^t \mathbf{e}) : \mathbf{TST}} \end{array} \right.$$

$$\begin{aligned} \mathcal{E}[MSG^t \bar{n}]_M &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[MSG^t v]_M &\rightarrow \mathcal{E}[MSG^t (\text{wrong “Non-number”})] \\ &\quad v \neq \bar{n} \\ \mathcal{E}[MSG^{\tau_1 \rightarrow \tau_2} \lambda x. \mathbf{e}]_M &\rightarrow \mathcal{E}[\lambda x : \tau_1. MSG^{\tau_2} ((\lambda x. \mathbf{e}) (GSM^{\tau_1} x))] \\ &\quad x \text{ not free in } \mathbf{e} \\ \mathcal{E}[MSG^{\tau_1 \rightarrow \tau_2} v]_M &\rightarrow \mathcal{E}[MSG^{\tau_1 \rightarrow \tau_2} \text{wrong “Non-procedure”}] \\ &\quad v \neq \lambda x. \mathbf{e} \\ \mathcal{E}[(GSM^t \bar{n})_S] &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[(GSM^{\tau_1 \rightarrow \tau_2} v)_S] &\rightarrow \mathcal{E}[(\lambda x. (MSG^{\tau_2} (v (MSG^{\tau_1} x))))] \end{aligned}$$

Figure 3. Extensions to figure 1 to form the simple natural embedding

rules translate values in the same way that SM_N^t and $^t MS_N$ did before, but also detect concrete, first-order witnesses to an invalid type ascription (*i.e.*, numbers for procedures or procedures for numbers) and abort the program if one is found. We call the language formed by these rules the simple natural embedding. We give its rules in figure 3, but it may be easier to understand how it works by reconsidering the examples we gave at the end of section 2. Each of those examples, modified to use the natural embedding rather than the lump embedding, successfully evaluates to the ML number $\bar{4}$. Here is the reduction sequence produced by the last of those examples, which was ill-typed before:

$$\begin{aligned} &((MSG^{\iota \rightarrow \iota} (\lambda x. (+ x 1))) \bar{3}) \\ \rightarrow &((\lambda x' : \iota. (MSG^{\iota} ((\lambda x. (+ x 1)) (GSM^{\iota} x')))) \bar{3}) \\ \rightarrow &(MSG^{\iota} ((\lambda x. (+ x 1)) (GSM^{\iota} \bar{3}))) \\ \rightarrow &(MSG^{\iota} ((\lambda x. (+ x 1)) \bar{3})) \\ \rightarrow^2 &(MSG^{\iota} \bar{4}) \rightarrow \bar{4} \end{aligned}$$

ML converts the Scheme add-one function to an ML function with type $\iota \rightarrow \iota$ by replacing it with a function that converts its argument to a Scheme number, feeds that number to the original Scheme function, and then converts the result back to an ML number. Then it applies this new function to the ML number $\bar{3}$, which gets converted to the Scheme number $\bar{3}$, run through the Scheme function, and finally converted back to the ML number $\bar{4}$, which is the program’s final answer.

The method works at higher-order types because it applies type conversions recursively. Consider this expression:

$$(^{\iota \rightarrow \iota} \rightarrow \iota MS_N (\lambda f. (if 0 (f \bar{1}) \bar{2} f)))$$

Depending on the behavior of its arguments, the Scheme procedure may or may not always produce numbers. ML treats it as though it definitely had type $(\iota \rightarrow \iota) \rightarrow \iota$, and wraps it to the ML value

$$\lambda x : \iota \rightarrow \iota. (^{\iota} MS_N ((\lambda f. if 0 (f \bar{1}) \bar{2} f) (SM_N^{\iota \rightarrow \iota} x)))$$

Whenever this value is applied to a function, that function is converted to a Scheme value at type $\iota \rightarrow \iota$ and the result is converted from Scheme to ML at type ι . Thus, intuitively, conversion in either direction works at a given type if it works in *both* directions at all smaller types.

Theorem 2. *If $\vdash_M \mathbf{e} : \tau$, then either $\mathbf{e} \rightarrow^* v$, $\mathbf{e} \rightarrow^* \mathbf{Error}$: str, or \mathbf{e} diverges.*

$$\begin{aligned} \mathbf{e} &= \dots \mid ({}^\tau MS_N \mathbf{e}) \\ \mathbf{e} &= \dots \mid (\mathcal{G}^\tau \mathbf{e}) \mid (SM_N^\tau \mathbf{e}) \end{aligned}$$

$$\begin{aligned} \mathbf{E} &= \dots \mid ({}^\tau MS_N \mathbf{E}) \\ \mathbf{E} &= \dots \mid (\mathcal{G}^\tau \mathbf{E}) \mid (SM_N^\tau \mathbf{E}) \end{aligned}$$

$$\frac{\Gamma \vdash_S \mathbf{e} : \mathbf{TST}}{\Gamma \vdash_S (\mathcal{G}^\tau \mathbf{e}) : \mathbf{TST}} \quad \frac{\Gamma \vdash_S \mathbf{e} : \mathbf{TST}}{\Gamma \vdash_M ({}^\tau MS_N \mathbf{e}) : \tau} \quad \frac{\Gamma \vdash_M \mathbf{e} : \tau}{\Gamma \vdash_S (SM_N^\tau \mathbf{e}) : \tau}$$

$$\begin{aligned} \mathcal{E}[\iota MS_N \bar{n}]_M &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[SM_N^\iota \bar{n}]_S &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[\tau_1 \rightarrow \tau_2 MS_N \lambda x. \mathbf{e}]_M &\rightarrow \mathcal{E}[\lambda x : \tau_1. SM_N^{\tau_2} ((\lambda x. \mathbf{e}) (SM_N^{\tau_1} \mathbf{x}))] \\ \mathcal{E}[(SM_N^{\tau_1 \rightarrow \tau_2} \mathbf{v})]_S &\rightarrow \mathcal{E}[(\lambda x. (SM_N^{\tau_2} (\mathbf{v} ({}^\tau MS_N \mathbf{x}))))] \\ \mathcal{E}[(\mathcal{G}^\iota \bar{n})]_S &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[(\mathcal{G}^\iota \mathbf{v})]_S &\rightarrow \mathcal{E}[\text{wrong "Non-number"}] \quad (\mathbf{v} \neq \bar{n}) \\ \mathcal{E}[(\mathcal{G}^{\tau_1 \rightarrow \tau_2} (\lambda x. \mathbf{e}))]_S &\rightarrow \mathcal{E}[(\lambda x'. (\mathcal{G}^{\tau_2} ((\lambda x. \mathbf{e}) (\mathcal{G}^{\tau_1} x'))))] \\ \mathcal{E}[(\mathcal{G}^{\tau_1 \rightarrow \tau_2} \mathbf{v})]_S &\rightarrow \mathcal{E}[\text{wrong "Non-procedure"}] \quad (\mathbf{v} \neq \lambda x. \mathbf{e}) \end{aligned}$$

Figure 4. Extensions to figure 1 to form the separated-guards natural embedding

3.2 A refinement: guards

Adding dynamic checks to boundaries themselves is an expedient way to ensure type soundness, but we find it a little troublesome. For one thing, boundaries are necessarily the core of any multi-language system, so they should be as small and simple as possible. For another, coupling the task of value conversion with the conceptually unrelated task of detecting and signalling errors means that changing the method of signalling errors requires modifying the internals of value conversion.

To decouple error-handling from value conversion, we separate the guarded boundaries of the previous subsection into their constituent parts: boundaries and guards. These separated boundaries have the semantics of the ${}^\tau MS_N$ and SM_N^τ boundaries we introduced at the beginning of this section. Guards will be new expressions of the form $(\mathcal{G}^\tau \mathbf{e})$ that perform all dynamic checks necessary to ensure that their arguments behave as τ in the sense of the previous subsection. In all initial terms, we will wrap every boundary with an appropriate guard: $({}^\tau MS_N (\mathcal{G}^\tau \mathbf{e}))$ instead of $(MSG^\tau \mathbf{e})$ and $(\mathcal{G}^\tau (SM_N^\tau \mathbf{e}))$ instead of $(GSM^\tau \mathbf{e})$.

Figure 4 shows the rules for guards. An ι guard applied to a number reduces to that number, and the same guard applied to a procedure aborts the program. A $\tau_1 \rightarrow \tau_2$ guard aborts the program if given a number, and if given a procedure reduces to a new procedure that applies a τ_1 guard to its input, runs the original procedure on that value, and then applies a τ_2 guard to the original procedure’s result. This is just like the strategy we use to convert functions in the first place, but it doesn’t perform any foreign-language translation by itself; it just distributes the guards in preparation for conversion later on.

The guard distribution rules for functions can move a guard arbitrarily far away from the boundary it protects; if this motion ever gave a value the opportunity to get to a boundary without first being blessed by the appropriate guard, the system would get stuck. We can prove this never happens by defining a combined language that has *both* guarded boundaries GSM^τ and MSG^τ and unguarded boundaries with separated guards ${}^\tau MS_N$, SM_N^τ , and \mathcal{G}^τ ; *i.e.* the language formed by combining figure 1 with figures 3 and 4. In this combined language, an argument by induction shows that guarded boundaries are observably equivalent to guards combined with unguarded boundaries.

Theorem 3. *For all Scheme expressions e and ML expressions e , the following propositions hold:*

$$\begin{aligned} (1) \quad &(MSG^\tau e) \simeq ({}^\tau MS_N (\mathcal{G}^\tau e)) \\ (2) \quad &(GSM^\tau e) \simeq (\mathcal{G}^\tau (SM_N^\tau e)) \end{aligned}$$

where \simeq is observable equivalence in the combined language.

In other words, we can replace any or all MSG^τ and GSM^τ boundaries with their separated versions without affecting the program, and therefore a program with *only* MSG^τ and GSM^τ boundaries is equivalent to the same program with *only* separated guards and unchecked boundaries; and therefore the language of figure 3 is the same as the language of figure 4.

3.3 A further refinement: contracts

While the guard strategy of the last subsection works, an implementation based on it would perform many unnecessary dynamic checks. For instance, the term $({}^\iota MS_N (\mathcal{G}^{\iota \rightarrow \iota} (\lambda x. (x))))$ is equivalent to

$$(\lambda x : \iota. ({}^\iota MS_N (\mathcal{G}^\iota ((\lambda x. x) (\mathcal{G}^\iota (SM_N^\iota \mathbf{x}))))))$$

The check performed by the leftmost guard is necessary, but the check performed by the rightmost guard could be omitted — since the value is coming directly from ML, ML’s type system guarantees that the conversion will succeed.

We can refine our guarding strategy to eliminate those unnecessary checks. We split guards into two varieties: positive guards, written \mathcal{G}_+^τ , that apply to values going from Scheme to ML, and negative guards, written \mathcal{G}_-^τ , that apply to values going from ML to Scheme. Their reduction rules are:

$$\begin{aligned} \mathcal{E}[(\mathcal{G}_+^\iota \bar{n})]_S &\rightarrow \mathcal{E}[\bar{n}] \\ \mathcal{E}[(\mathcal{G}_+^\iota \mathbf{v})]_S &\rightarrow \mathcal{E}[\text{wrong "Non-number"}] \quad (\mathbf{v} \neq \bar{n}) \\ \mathcal{E}[(\mathcal{G}_+^{\tau_1 \rightarrow \tau_2} \mathbf{v})]_S &\rightarrow \mathcal{E}[(\mathcal{G}_+^{\tau_2} (\mathbf{v} (\mathcal{G}_+^{\tau_1} \mathbf{x})))] \quad (\mathbf{v} = \lambda x. \mathbf{e}) \\ \mathcal{E}[(\mathcal{G}_+^{\tau_1 \rightarrow \tau_2} \mathbf{v})]_S &\rightarrow \mathcal{E}[\text{wrong "Non-function"}] \quad (\mathbf{v} \neq \lambda x. \mathbf{e}) \\ \mathcal{E}[(\mathcal{G}_-^\iota \mathbf{v})]_S &\rightarrow \mathcal{E}[\mathbf{v}] \\ \mathcal{E}[(\mathcal{G}_-^{\tau_1 \rightarrow \tau_2} \mathbf{v})]_S &\rightarrow \mathcal{E}[(\lambda x. (\mathcal{G}_-^{\tau_2} (\mathbf{v} (\mathcal{G}_-^{\tau_1} \mathbf{x}))))] \end{aligned}$$

The function cases are the interesting rules here. Since functions that result from positive guards are bound for ML, we check the inputs that ML will supply them using a negative guard; since the result of those functions will be Scheme values going to ML, they must be guarded with positive guards. Negative guards never directly signal an error; they exist only to protect ML functions from erroneous Scheme inputs. They put positive guards on the arguments to ML functions but use negative guards on their results because those values will have come from ML.

This new system eliminates half of the first-order checks associated with the first separated-guard system, but maintains equivalence with the simple natural embedding system.

Theorem 4. *For all ML expressions e and Scheme expressions e in the language that combines simple natural boundaries and positive and negative guards, both of the following propositions hold:*

$$\begin{aligned} (1) \quad &(MSG^\tau e) \simeq ({}^\tau MS_N (\mathcal{G}_+^\tau e)) \\ (2) \quad &(GSM^\tau e) \simeq (\mathcal{G}_-^\tau (SM_N^\tau e)) \end{aligned}$$

As it happens, we do not need to add \mathcal{G}_+^τ and \mathcal{G}_-^τ as extensions, because they can be implemented directly in the core Scheme language of figure 1. In fact, they are exactly contracts, or more specifically pairs of projections in exactly the sense of Findler and Blume [17] where the two parties are + (Scheme) and – (ML). Since ML’s type system guarantees that – never breaks its contract, only + will ever be blamed for a violation. Practically, this means we can safely use contracts to protect type invariants in foreign function interfaces. This adds to our confidence in Gray *et al*’s fine-grained interoperability scheme [24], for example. More theoretically, it means we can use the simple system of section 3.1 for our models

$$\begin{aligned}
\mathcal{T}_M^{\tau_1 \rightarrow \tau_2} &\stackrel{\text{def}}{=} \lambda \mathbf{x} : \mathbf{L}. \lambda \mathbf{y} : \tau_1. (\mathcal{T}_M^{\tau_2} (\mathcal{L}MS_L ((SM_L^{\mathbf{L}} \mathbf{x}) \mathcal{T}_S^{\tau_1} (SM_L^{\tau_1} \mathbf{y})))) \\
\mathcal{T}_S^{\tau_1 \rightarrow \tau_2} &\stackrel{\text{def}}{=} \lambda x. \lambda y. \mathcal{T}_S^{\tau_2} (SM_L^{\tau_2} ((\tau_1 \rightarrow \tau_2 MS_L x) (\mathcal{T}_M^{\tau_1} (\mathcal{L}MS_L y)))) \\
\mathcal{T}_M^{\iota} &\stackrel{\text{def}}{=} (\lambda \mathbf{x} : \mathbf{L}. \mathcal{L}MS_L \\
&\quad (\mathbf{Y} (\lambda f. \lambda n. \\
&\quad \quad (\text{if0 } n \\
&\quad \quad \quad (SM_L^{\iota} \bar{0}) \\
&\quad \quad \quad (SM_L^{\iota} (+ \bar{1} (\mathcal{L}MS_L (f (- n \bar{1})))))) \\
&\quad \quad (SM_L^{\iota} \mathbf{x}))) \\
\mathcal{T}_S^{\iota} &\stackrel{\text{def}}{=} (\lambda x. \\
&\quad (\mathbf{Y} (\lambda f. \lambda n. (SM_L^{\iota} \\
&\quad \quad (\text{if0 } (\mathcal{L}MS_L n) \\
&\quad \quad \quad (\mathcal{L}MS_L \bar{0}) \\
&\quad \quad \quad (\mathcal{L}MS_L (+ \bar{1} (f (SM_L^{\iota} (- (\mathcal{L}MS_L n) \bar{1})))))) \\
&\quad \quad x))
\end{aligned}$$

Figure 5. Translation functions for lump values

and be more confident that our soundness results apply to actual multi-language embeddings that we write with contracts. From the contract perspective, it also shows one way of using mechanized reasoning to statically eliminate dynamic assertions from annotated programs. In this light it can be seen as a hybrid type system [21].

4. What have we gained?

The natural embedding fits our intuitive notion of what a useful interoperability system ought to look like much more than the lump embedding does, so it seems like it should give us the power to write more programs. However, that is not the case: τMS_N and SM_N^{τ} boundaries are macro-expressible in the lump embedding (in the sense of Felleisen [14]), meaning that any natural-embedding program can be translated using local transformations into a lump-embedding program.

To show that, we define two type-indexed functions, \mathcal{T}_S^{τ} and \mathcal{T}_M^{τ} , that can be written in the lump embedding. These functions translate values whose ML type is τ from Scheme to ML and from ML to Scheme, respectively; they are shown in figure 5. (for clarity, in that figure and below we use the notation τMS_L and SM_L^{τ} rather than τMS and SM^{τ} to refer to the lump embedding’s boundaries.) The translation functions for higher-order types use essentially the same method we presented in section 3 for converting a procedure value — they translate a procedure by left- and right-composing appropriate translators for its input and output types. That leaves us with nothing but the base types, in our case numbers.

Those require more work. The key insight is that we can use a two-language equivalent of Church-encoded numbers to send a number across a boundary, and both translators do that the same way: the receiver gives the sender the opaque number to be translated, a zero value, and a successor function. The sender then performs a foreign application of the successor function to the zero value n times where n is the number to be sent. The result of that operation is a foreign value representing the native number that was to be sent. (The two translators are superficially different from each other, but only because we restrict the \mathbf{Y} fixed-point combinator to Scheme, where it can be written directly.)

Theorem 5. *In the language formed by extending the language of figure 1 with both figure 2 and the language of section 3.3, both of the following propositions hold:*

$$\begin{aligned}
(\mathcal{G}_-^{\tau} (SM_N^{\tau} e)) &\simeq (\mathcal{G}_-^{\tau} (\mathcal{T}_S^{\tau} (SM_L^{\tau} e))) \\
(\tau MS_N (\mathcal{G}_+^{\tau} e)) &\simeq (\mathcal{T}_M^{\tau} (\tau MS_L (\mathcal{G}_+^{\tau} e)))
\end{aligned}$$

In other words, we can run any natural-embedding program in the lump embedding by converting all of its natural-embedding boundaries into lump boundaries wrapped in the appropriate guard and conversion functions. Since each translation in figure 5 can be applied independently, since each one is a local transformation, and since guards are expressible in the core Scheme language, it follows that the SM_N^{τ} and τMS_N boundaries are macro-expressible in the lump embedding.

Based on these translation rules, we were able to implement a program that communicated numbers between PLT Scheme and C using PLT Scheme’s foreign interface [3] but maintaining a strict lump discipline.

While the translations of figure 5 suffice for our purposes here, it is worth mentioning that they could be made faster. Rather than having the sender transmit “add-one” and “stop” signals, the receiver could give the sender representations of 0 and 1 and let the sender send successive bits of the number to be converted. This approach would run in time proportional to the log of the converted number, assuming efficient bit-shift operations. This is likely to be linear in the size of the representation of integers, though this is still suboptimal since such conversions are likely to be either simple bit operations or even noops at the level of the underlying machine.

5. From type-directed to type-mapped conversion

In the previous sections, we have gotten double duty out of the type annotations on boundaries: statically, they have indicated the type of each boundary, and dynamically they have indicated how to convert any value that appears at the boundary. That overloading is fine (even desirable) for some multi-language systems, but it is insufficient for others. For instance, since C does not have an exception mechanism, many C functions (e.g., `malloc`, `fopen`, `sqrt`) return a normal value on success and a sentinel value on error. A foreign function interface might automatically convert error indicators from such functions into exceptions, while converting non-errors as normal. To model such a conversion, we must generalize boundaries so that instead of containing a *type*, they contain a *conversion strategy* from which a type can be derived.

In the numbers-for-errors example, we can consider a variant of our earlier multi-language systems in which we expect ML to use the zero for error convention for some functions, and Scheme has a very simple exception mechanism in which (`wrong str`) raises an exception, and (`handle e1 e2`) tries to evaluate e_2 unless it raises an exception, in which case it abandons e_2 and instead evaluates the handler expression e_1 . If a Scheme exception reaches ML, it aborts the program unless the ML boundary it reaches expects a value following the zero-for-error convention, in which case it becomes $\bar{0}$. Similarly if $\bar{0}$ flows from ML to Scheme in a context where the zero-for-error convention is in place, it becomes a Scheme exception.

We give a model of this system in figure 6. The Scheme contexts identify two different layers: those in which no boundary or exception handler appears (H), and those in which boundaries and handlers may appear (E). This layering is a simplified variant of the exception model presented by Wright and Felleisen [47].

The core of the system is a conversion strategy κ that replaces the type on all boundaries, and an associated $[\cdot]$ metafunction from κ to τ that takes a conversion strategy “down to” a type. We use the conversion strategy to add a single conversion, ι , indicating places where $\bar{0}$ indicates an error. Typing judgments are as before, except that boundary rules must apply the $[\cdot]$ function as necessary. Reducing a boundary is also as before, with additions corresponding to the ML-to-Scheme and Scheme-to-ML conversions for values at ι boundaries.

Theorem 6. *The language of figure 6 is type-sound.*

$$\begin{aligned}
\mathbf{e} &= \dots \mid (MSG^\kappa \mathbf{e}) \\
\mathbf{e} &= \dots \mid (\mathbf{handle} \mathbf{e} \mathbf{e}) \mid (GSM^\kappa \mathbf{e}) \mid (\mathbf{handle} \mathbf{e} \mathbf{e}) \\
\kappa &= \iota \mid \iota! \mid \kappa \rightarrow \kappa \\
\mathbf{E} &= \dots \mid \mathbf{E}[(MSG^\kappa \mathbf{E})]_M \\
\mathbf{H} &= [\]_S \mid (\mathbf{H} \mathbf{e}) \mid (\vee \mathbf{H}) \mid (op \mathbf{H} \mathbf{e}) \mid (op \vee \mathbf{H}) \mid (\text{if} 0 \mathbf{H} \mathbf{e} \mathbf{e}) \mid (pr \mathbf{H}) \\
\mathbf{E} &= \mathbf{H} \mid \mathbf{H}[(GSM^\kappa \mathbf{E})]_S \mid \mathbf{H}[(\mathbf{handle} \mathbf{e} \mathbf{E})]_S \\
\frac{\Gamma \vdash_S \mathbf{e} : \mathbf{TST}}{\Gamma \vdash_M (MSG^\kappa \mathbf{e}) : [\kappa]} & \quad \frac{\Gamma \vdash_M \mathbf{e} : [\kappa]}{\Gamma \vdash_S (GSM^\kappa \mathbf{e}) : \mathbf{TST}} \\
[\iota] &= \iota \\
[\iota!] &= \iota \\
[\kappa_1 \rightarrow \kappa_2] &= [\kappa_1] \rightarrow [\kappa_2] \\
\mathcal{E}[(MSG^{\iota!} \bar{n})]_M &\rightarrow \mathcal{E}[\bar{n}] \\
\mathcal{E}[(MSG^{\iota!} \mathbf{H}[\mathbf{wrong} \mathit{str}])]_M &\rightarrow \mathcal{E}[\bar{0}] \\
\mathcal{E}[(MSG^\kappa \mathbf{H}[\mathbf{wrong} \mathit{str}])]_M &\rightarrow \mathbf{Error} : \mathit{str} \text{ (if } \kappa \neq \iota!) \\
\mathcal{E}[(GSM^{\iota!} \bar{0})]_S &\rightarrow \mathcal{E}[(\mathbf{wrong} \text{ "zero"})] \\
\mathcal{E}[(GSM^{\iota!} \bar{n})]_S &\rightarrow \mathcal{E}[\bar{n}] \text{ (if } n \neq 0) \\
\mathcal{E}[(\mathbf{handle} \mathbf{e}_1 \mathbf{H}[\mathbf{wrong} \mathit{str}])]_S &\rightarrow \mathcal{E}[\mathbf{e}_1]
\end{aligned}$$

Figure 6. Extensions to figure 3 for mapped embedding

This example demonstrates a larger point: although we have used a boundary’s type as its conversion strategy for most of the systems in this paper, they are separate ideas. Decoupling them has a number of pleasant effects: first, it allows us to use non-type-directed conversions, as we have shown. Second, it illustrates that boundaries depend on conversion strategies, not type information *per se*, which suggests that type erasure (but not conversion strategy erasure) remains possible in our framework. Finally, the separation makes it easier to understand the connection between these formal systems and tools like SWIG, in particular SWIG’s type-map feature [4]: from this perspective, SWIG is a tool that automatically generates boundaries that pull C++ values into Python (or another high-level language), and type-maps allow the user to write a new conversion strategy and specify the circumstances under which it should be used.

6. Related work

The work most directly related to ours is Benton’s “Embedded Interpreters” [6], which lays out a method for embedding interpreters into statically typed languages using type-indexed embeddings and projections (the same method was also discovered independently by Ramsey [42]). His work considers only the asymmetric case where a typed host language embeds an untyped language, and focuses on implementation rather than formal techniques. Still, readers will find that the flavor of his work is quite similar to this work. Zdancewic, Grossman, and Morrisett’s work [48] is also similar to ours in that it introduces two-agent calculi and boundaries; their work, however, focuses on information-hiding properties and does not allow different languages to interoperate.

Other work on the semantics of interoperability tends to focus on the properties of multi-language runtime systems. This includes a pair of formalisms for the semantics of COM, the first by Ibrahim and Szyperski [27] and the second by Pucella [41] and also Gordon and Syme’s formalization of a type-safe intermediate language designed specifically for multi-language interoperation [23]. Kennedy [29] pointed out that in multi-language systems, observations in one language can break equations in the other and that this is a practical problem. Our system is one way to reason about these problems precisely. Trifonov and Shao have developed an ab-

stract intermediate language for multi-language programs that aids reasoning about interactions between effects in the two source languages [46]. While our framework can also address effects, it does not address their implementation. Instead our work focuses on their semantics as seen by the source languages, a topic Trifonov and Shao do not discuss. Finally, Furr and Foster have built a system for verifying certain safety properties of the OCaml foreign-function interface by analyzing C code for problematic uses of OCaml values [22].

On the issue of combining typed and untyped code, Henglein and Rehof [25,26] have done work on translating Scheme into ML, inserting ML equivalents of our guards to simulate Scheme’s dynamic checks. Some languages have introduced ways of mixing typed and untyped code using a dynamic type [1], similar to our boundaries and lumps; for instance Cardelli’s Amber [10], Chambers *et al*’s Cecil [13] or Gray, Findler and Flatt’s ProfessorJ [24].

There has been far too much implementation work connecting high-level languages to list it all here. In addition to the projects we have already mentioned, there are dozens of compilers that target the JVM, the .NET CLR, or COM. There have also been more exotic embeddings; two somewhat recent examples are an embedding of Alice (an SML extension) into the Oz programming language [30] and the LazyScheme educational embedding of a lazy variant of Scheme into strict-evaluation-order Scheme [2]. There has been even more work on connecting high-level languages to low-level languages; in addition to the venerable SWIG [5] there are many more systems that try to sanitize the task [3, 9, 11, 20, 43].

7. Conclusion

We have shown how to give operational semantics to multi-language systems and still use the same formal techniques that apply to single languages. This work has focused on two points in the design space for interoperating languages: the lump and natural embeddings. We see aspects of these two points in many real foreign function interfaces: for instance, SML.NET translates flat values and .NET objects, but forces lump-like behavior for higher-order ML functions. The Java Native Interface provides all foreign values as lumps, but provides a large set of constant functions for manipulating them, similar to our **fa** (“foreign-apply”) function from section 2. The lump embedding’s ease of implementation and natural embedding’s ease of use both pull on language design, and most real multi-language systems lie in between.

Furthermore, we have talked mostly about foreign function interfaces in this paper, but our technique scales to other kinds of multi-language systems. Any situation in which two logically different languages interact is a candidate for this treatment — *e.g.* embedded domain-specific languages, some uses of metaprogramming, and even contract systems, viewing each party to the contract as a separate language. By treating these uniformly as multi-language systems, we might be able to connect them in fruitful ways. Even with the limited scope considered in this paper, we have discovered connections between contracts, foreign function interfaces, and hybrid type systems (as discussed in section 3.3). Widening the net to include a broader interpretation of interoperating systems suggests that this work barely scratches the surface.

Acknowledgments The authors would like to thank Cormac Flanagan for his suggestion of the phrase “cross-language cast”; to Matthias Felleisen, Dave MacQueen, and the anonymous reviewers for their comments on earlier versions of this paper; and to David Walker, for shepherding this paper from its original twelve-page form to the eight-page version presented here. This work was funded in part by the National Science Foundation.

References

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically typed language. *ACM Transactions on Programming Languages and Systems*, 13(2):237–268, April 1991.
- [2] E. Barzilay and J. Clements. Laziness without all the hard work. In *Workshop on Functional and Declarative Programming in Education (FDPE)*, 2005.
- [3] E. Barzilay and D. Orlovsky. Foreign interface for PLT Scheme. In *Workshop on Scheme and Functional Programming*, 2004.
- [4] D. Beazley. Pointers, constraints, and typemaps. In SWIG 1.1 Users Manual. Available online: <http://www.swig.org/Doc1.1/HTML/Typemaps.html>.
- [5] D. Beazley. SWIG: An easy to use tool for integrating scripting languages with C and C++. In *4th Tcl/Tk Workshop*, pages 129–139, 1996. Available online: <http://www.swig.org/papers/Tc196/tc196.html>.
- [6] N. Benton. Embedded interpreters. *Journal of Functional Programming*, 15:503–542, July 2005.
- [7] N. Benton and A. Kennedy. Interlanguage working without tears: Blending SML with Java. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 126–137, 1999.
- [8] N. Benton, A. Kennedy, and C. V. Russo. Adventures in interoperability: the SML.NET experience. In *Proceedings of the 6th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*, pages 215–226, 2004.
- [9] M. Blume. No-longer-foreign: Teaching an ML compiler to speak C “natively”. In *Workshop on Multi-Language Infrastructure and Interoperability (BABEL)*, 2001.
- [10] L. Cardelli. Amber. In G. Cousineau, P.-L. Curien, and B. Robinet, editors, *Combinators and functional programming languages*, volume 242. Springer-Verlag, 1986.
- [11] M. M. T. Chakravarty. C→HASKELL, or yet another interfacing tool. In *International Workshop on Implementation of Functional Languages (IFL)*, 1999.
- [12] M. M. T. Chakravarty. The Haskell 98 foreign function interface 1.0, 2002. Available online: <http://www.cse.unsw.edu.au/~chak/haskell/ffi/>.
- [13] C. Chambers and The Cecil Group. The Cecil language: Specification and rationale, version 3.2. Technical report, Department of Computer Science and Engineering, University of Washington, February 2004. Available online: <http://www.cs.washington.edu/research/projects/cecil/pubs/cecil-spec.html>.
- [14] M. Felleisen. On the expressive power of programming languages. *Science of Computer Programming*, 17:35–75, 1991.
- [15] M. Felleisen, D. Friedman, E. Kohlbecker, and B. Duba. A syntactic theory of sequential control. pages 205–237, 1987.
- [16] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 102:235–271, 1992. Original version in: Technical Report 89-100, Rice University, June 1989.
- [17] R. B. Findler and M. Blume. Contracts as pairs of projections. In *International Symposium on Functional and Logic Programming (FLOPS)*, 2006.
- [18] R. B. Findler and M. Felleisen. Contracts for higher-order functions. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2002.
- [19] S. Finne, D. Leijen, E. Meijer, and S. Peyton Jones. Calling Hell from Heaven and Heaven from Hell. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 114–125, 1999.
- [20] K. Fisher, R. Pucella, and J. Reppy. A framework for interoperability. In *Workshop on Multi-Language Infrastructure and Interoperability (BABEL)*, 2001.
- [21] C. Flanagan. Hybrid type checking. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2006.
- [22] M. Furr and J. S. Foster. Checking type safety of foreign function calls. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 62–72, 2005.
- [23] A. D. Gordon and D. Syme. Typing a multi-language intermediate code. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 248–260, 2001.
- [24] K. E. Gray, R. B. Findler, and M. Flatt. Fine grained interoperability through mirrors and contracts. In *Conference on Object-Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, 2005.
- [25] F. Henglein. Dynamic typing: Syntax and proof theory. *Science of Computer Programming*, 22(3):197–230, 1994.
- [26] F. Henglein and J. Rehof. Safe polymorphic type inference for a dynamically typed language: translating Scheme to ML. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture (FPCA)*, 1995.
- [27] R. Ibrahim and C. Szyperski. The COMEL language. Technical Report FIT-TR-97-06, Faculty of Information Technology, Queensland University of Technology, Brisbane, Australia, 1997.
- [28] Jython. <http://www.jython.org/>.
- [29] A. Kennedy. Securing the .NET programming model. *Theoretical Computer Science*. To appear. <http://research.microsoft.com/~akenn/sec/>.
- [30] L. Kornstaedt. Alice in the land of Oz - an interoperability-based implementation of a functional language on top of a relational language. In *Workshop on Multi-Language Infrastructure and Interoperability (BABEL)*, 2001.
- [31] I. Mason and C. Talcott. Equivalence in functional languages with effects. *Journal of Functional Programming*, 1:287–327, July 1991.
- [32] J. Matthews and R. B. Findler. An operational semantics for R5RS Scheme. In *Workshop on Scheme and Functional Programming*, 2005.
- [33] J. Matthews, R. B. Findler, M. Flatt, and M. Felleisen. A visual environment for developing context-sensitive term rewriting systems. In *Proceedings of the International Conference on Rewriting Techniques and Applications (RTA)*, 2004.
- [34] E. Meijer, N. Perry, and A. van Yzendoorn. Scripting .NET using Mondrian. In *European Conference on Object-Oriented Programming (ECOOP)*, pages 150–164, London, UK, 2001. Springer-Verlag.
- [35] P. Meunier and D. Silva. From Python to PLT Scheme. In *Proceedings of the Fourth Workshop on Scheme and Functional Programming*, pages 24–29, 2003.
- [36] M. Odersky, P. Altherr, V. Cremet, B. Emir, S. Micheloud, N. Mihaylov, M. Schinz, E. Stenman, and M. Zenger. An Introduction to Scala. <http://scala.epfl.ch/docu/files/ScalaIntro.pdf>, 2005.
- [37] A. Ohori and K. Kato. Semantics for communication primitives in a polymorphic language. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 99–112, 1993.
- [38] B. C. Pierce. *Types and Programming Languages*. The MIT Press, 2002.
- [39] P. Pinto. Dot-Scheme: A PLT Scheme FFI for the .NET framework. In *Workshop on Scheme and Functional Programming*, November 2003.
- [40] G. D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, pages 223–255, 1977.
- [41] R. Pucella. Towards a formalization for COM, part I: The primitive calculus. In *Conference on Object-Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, 2002.
- [42] N. Ramsey. Embedding an interpreted language using higher-order functions and types. In *Interpreters, Virtual Machines and Emulators (IVME '03)*, pages 6–14, 2003.
- [43] J. Reppy and C. Song. Application-specific foreign-interface generation. In *International Conference on Generative Programming and Component Engineering (GPCE)*, 2006.
- [44] A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. *Lisp and Symbolic Computation*, 1993.
- [45] P. Steckler. MysterX: A Scheme toolkit for building interactive applications with COM. In *Technology of Object-Oriented Languages and Systems (TOOL)*, pages 364–373, 1999. Available online: citeseer.ist.psu.edu/steckler99mysterx.html.
- [46] V. Trifonov and Z. Shao. Safe and principled language interoperation. In *European Symposium on Programming (ESOP)*, pages 128–146, 1999.
- [47] A. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, pages 38–94, 1994. First appeared as Technical Report TR160, Rice University, 1991.
- [48] S. Zdanczewicz, D. Grossman, and G. Morrisett. Principals in programming languages. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 1999.

PLT Redex [33] implementations of all the formal systems presented here are available online at <http://www.cs.uchicago.edu/~jacobm/papers/multilang/>.

A. Proofs

A.1 Proof of theorem 1

Lemma 7 (Preservation).

If $\vdash_M e : \tau$ and $e \rightarrow e'$, then $\vdash_M e' : \tau$.

Proof. By cases on the reduction $e \rightarrow e'$. We will make use of three standard lemmas: a uniqueness of types lemma, an inversion lemma, and a replacement lemma. The first two are the straightforward adaptations of uniqueness and inversion in, e.g., Pierce [38]. We take the replacement lemma from Wright and Felleisen [47]. It states that if $\Gamma \vdash_M e : \tau$, e' is a Scheme or ML subterm of e , and $\Gamma' \vdash e' : \tau$ (where \vdash is \vdash_M or \vdash_S as appropriate, and τ may be τ' or **TST**), then for all terms e'' where $\Gamma' \vdash e'' : \tau'$, $\Gamma \vdash_M e[e'/e''] : \tau$.

With those lemmas, the cases are entirely standard except for the boundary reductions. We present only those.

Case $\mathcal{E}[\tau' MS (SM^{\tau'} \mathbf{v})] \rightarrow \mathcal{E}[\mathbf{v}]$

By premise and uniqueness of types, we have that $\Gamma \vdash_M \tau' MS (SM^{\tau'} \mathbf{v}) : \tau'$. By inversion we have that $\Gamma \vdash_S SM^{\tau'} \mathbf{v} : \mathbf{TST}$, and by inversion again we have that $\Gamma \vdash_M \mathbf{v} : \tau'$. Thus by replacement, $e' = \mathcal{E}[\mathbf{v}]$ has type τ .

Case $\mathcal{E}[\tau' MS \mathbf{v}]_M \rightarrow \mathcal{E}[\tau' MS (\mathbf{wrong} \text{ “Bad value”})]$ where $\mathbf{v} \neq (SM^{\tau'} \mathbf{v})$

By premise and uniqueness of types, $\Gamma \vdash_M \tau' MS \mathbf{v} : \tau'$. A calculation shows that $(\tau' MS (\mathbf{wrong} \text{ “Bad value”}))$ also has type τ' in Γ , so by replacement $e' = \mathcal{E}[\tau' MS (\mathbf{wrong} \text{ “Bad value”})]$ has type τ .

Case $\mathcal{E}[SM^L ({}^L MS \mathbf{v})]_S \rightarrow \mathcal{E}[\mathbf{v}]$

By premise and uniqueness of types, we have that $\Gamma \vdash_S SM^L ({}^L MS \mathbf{v}) :$

TST. By inversion we have that $\Gamma \vdash_M ({}^L MS \mathbf{v}) : \mathbf{L}$, and by inversion again we have that $\Gamma \vdash_S \mathbf{v} : \mathbf{TST}$. Thus by replacement, $e' = \mathcal{E}[\mathbf{v}]$ has type τ . \square

For progress, we have to strengthen the statement of the lemma; we need to be able to use induction on both ML's and Scheme's typing judgments, and we need to prove the statement for all evaluation contexts because we have no notion of evaluating Scheme programs in empty contexts. Nevertheless, the proof is straightforward.

Lemma 8 (Progress). *For all ML expressions e , Scheme expressions e , both of the following hold:*

- (1) If $\vdash_M e : \tau$, then either e is an ML value or for all top-level evaluation contexts $\mathcal{E}[\]_M$ either there exists an e' such that $\mathcal{E}[e] \rightarrow e'$ or $\mathcal{E}[e] \rightarrow \mathbf{Error}$: str for some error message str.
- (2) If $\vdash_S e : \mathbf{TST}$, then either e is a Scheme value or for all top-level evaluation contexts $\mathcal{E}[\]_S$ either there exists an e' such that $\mathcal{E}[e] \rightarrow e'$ or $\mathcal{E}[e] \rightarrow \mathbf{Error}$: str for some error message str.

Proof. By simultaneous induction on the size of the typing derivation. Cases will generally make use of the fact that we can compose contexts where the hole in the outer context corresponds to the outermost language of the inner context, but are otherwise straightforward. We show the hardest case.

Case $\frac{\vdash_S e : \mathbf{TST}}{\vdash_M (\tau' MS e) : \tau}$:

We must show that either $(\tau' MS e)$ is a value or that for any given top-level evaluation context $\mathcal{E}[\]_M$, $\mathcal{E}[\tau' MS e]$ reduces. Either e is a Scheme value or e is not a Scheme value. If it is, then depending on τ either the entire expression is a value or one of the two reduction rules for $\tau' MS$ boundaries directly applies. If e is not a Scheme value, then by induction on (2) we have that for all ML evaluation contexts with Scheme holes, e can reduce. The composition $\mathcal{E}[SM^{\tau} [\]_S]$ is an ML evaluation context with a Scheme hole; thus $\mathcal{E}[SM^{\tau} e]$ reduces as required. \square

Theorem 1. *If $\Gamma \vdash_M e : \tau$, then either $e \rightarrow^* \mathbf{v}$, $e \rightarrow^* \mathbf{Error}$: str, or e diverges.*

Proof. Combination of lemmas 7 and 8. \square

The proofs of theorems 2 and 6 are analogous.

A.2 Proof of theorem 3

Lemma 9 (Context equivalence). *If for all \mathcal{E} , $eval_N(\mathcal{E}[e_1]) = eval_N(\mathcal{E}[e_2])$ then $e_1 \simeq e_2$.*

Proof. As in Felleisen *et al* [15, theorem 4.11]. Note that this is also Mason and Talcott's *ciu*-equivalence notion [31]. \square

Lemma 10 (Closure under reduction). *If for all evaluation contexts $\mathcal{E}[\]$, $\mathcal{E}[e_1] \rightarrow^* \mathcal{E}[e'_1]$ (uniquely) and $\mathcal{E}[e_2] \rightarrow^* \mathcal{E}[e'_2]$ (uniquely) and $e'_1 \simeq e'_2$, then $e_1 \simeq e_2$.*

Proof. Assume $e_1 \not\simeq e_2$. Then there is some context $\mathcal{E}'[\]$ such that $eval_N(\mathcal{E}'[e_1])$ produces a different result from $eval_N(\mathcal{E}'[e_2])$. But by premises, $\mathcal{E}'[e_1] \rightarrow^* \mathcal{E}'[e'_1]$ and $\mathcal{E}'[e_2] \rightarrow^* \mathcal{E}'[e'_2]$, and $e'_1 \simeq e'_2$. By definition the evaluations of these two terms are the same. This is a contradiction. \square

Theorem 3. *For all Scheme expressions e and ML expressions e , the following propositions hold:*

- (1) $MSG^{\tau} e \simeq \tau MS_N(\mathcal{G}^{\tau} e)$
- (2) $GSM^{\tau} e \simeq \mathcal{G}^{\tau}(SM_N^{\tau} e)$

Proof. By lemma 9, for each proposition it suffices to show that the two forms are equivalent in any evaluation context $\mathcal{E}[\]$. If e in (1) or e in (2) diverges or signals an error when evaluated, then both sides of each equivalence do so as well. Thus it suffices to show that the evaluation function produces identical results when e (or e) is a value \mathbf{v} (or \mathbf{v}) and the entire expression appears in an evaluation context. We prove that by induction on τ .

Case $\tau = \iota$: For (1): if \mathbf{v} is the number n :

$$\begin{aligned} \mathcal{E}[MSG^{\iota} n] &\rightarrow \mathcal{E}[n] \\ \mathcal{E}[\iota MS_N(\mathcal{G}^{\iota} n)] &\rightarrow \mathcal{E}[\iota MS_N n] \rightarrow \mathcal{E}[n] \end{aligned}$$

If \mathbf{v} is $(\lambda x.e')$, then

$$\begin{aligned} \mathcal{E}[MSG^{\iota}(\lambda x.e')] &\rightarrow \mathcal{E}[MSG^{\iota}(\mathbf{wrong} \text{ “Non-number”})] \\ \mathcal{E}[\iota MS_N(\mathcal{G}^{\iota}(\lambda x.e'))] &\rightarrow \mathcal{E}[\iota MS_N(\mathbf{wrong} \text{ “Non-number”})] \end{aligned}$$

For (2): the analogous argument applies, but by inversion we know that \mathbf{v} must be a natural number so the error case is impossible.

Case $\tau = \tau_1 \rightarrow \tau_2$: For (1): if \mathbf{v} is $(\lambda x.e')$, then the left-hand side expression reduces as follows:

$$\mathcal{E}[MSG^{\tau_1 \rightarrow \tau_2}(\lambda x.e')] \rightarrow \mathcal{E}[(\lambda \mathbf{x} : \tau. MSG^{\tau_2}((\lambda x.e') (GSM^{\tau_1} \mathbf{x})))]$$

For the right-hand side:

$$\begin{aligned}
& \mathcal{E}[\tau_1 \rightarrow \tau_2 MS_N(\mathcal{G}^{\tau_1 \rightarrow \tau_2}(\lambda x.e'))] \\
& \rightarrow \hspace{15em} \text{(by guard reduction)} \\
& \mathcal{E}[\tau_1 \rightarrow \tau_2 MS_N(\lambda x'. \mathcal{G}^{\tau_2}((\lambda x.e')(\mathcal{G}^{\tau_1} x')))] \\
& \rightarrow \hspace{15em} \text{(by boundary reduction)} \\
& \mathcal{E}[\lambda x'' : \tau_1. \tau_2 MS_N((\lambda x'. \mathcal{G}^{\tau_2}((\lambda x.e')(\mathcal{G}^{\tau_1} x')))(SM_N^{\tau_1} x''))] \\
& \simeq \hspace{15em} \text{(by } \beta_\omega \text{ [44])} \\
& \mathcal{E}[\lambda x'' : \tau_1. \tau_2 MS_N(\mathcal{G}^{\tau_2}((\lambda x.e')(\mathcal{G}^{\tau_1}(SM_N^{\tau_1} x''))))] \\
& \simeq \hspace{15em} \text{(by induction hypothesis 1)} \\
& \mathcal{E}[\lambda x'' : \tau_1. \tau_2 MS_N(\mathcal{G}^{\tau_2}((\lambda x.e')(\mathcal{GSM}^{\tau_1} x'')))] \\
& \simeq \hspace{15em} \text{(by induction hypothesis 2)} \\
& \mathcal{E}[\lambda x'' : \tau_1. MSG^{\tau_2}((\lambda x.e')(\mathcal{GSM}^{\tau_1} x''))]
\end{aligned}$$

If v is a number n , then $\mathcal{E}[MSG^{\tau_1 \rightarrow \tau_2} n] \rightarrow \text{wrong}$ “Non-procedure” and $\mathcal{E}[\tau_1 \rightarrow \tau_2 MS_N(\mathcal{G}^{\tau_1 \rightarrow \tau_2} n)] \rightarrow \text{wrong}$ “Non-procedure”.

For (2): the analogous argument applies, with boundaries swapped in the final steps. Also, again by inversion, we can rule out the case in which an error occurs. \square

The proof of theorem 4 is as above, *mutatis mutandis*.

A.3 Proof of theorem 5

Theorem 5. *In the language formed by extending the language of figure 1 with both figure 2 and the language of section 3.3, both of the following propositions hold:*

$$\begin{aligned}
(\mathcal{G}_N^\tau(SM_N^\tau e)) & \simeq (\mathcal{G}_S^\tau(\mathcal{T}_S^\tau(SM_L^\tau e))) \\
(\tau MS_N(\mathcal{G}_N^\tau e)) & \simeq (\mathcal{T}_M^\tau(\tau MS_L(\mathcal{G}_L^\tau e)))
\end{aligned}$$

Proof. As before, it is sufficient to consider the case where $e = v$ (for the first proposition) or $e = v$ (for the second) and the overall expression is in an evaluation context.

Case $\tau = \iota$: By lemma 11.

Case $\tau = \tau_1 \rightarrow \tau_2$: By induction as in the proof of theorem 3. \square

Lemma 11. *Both of the following propositions hold:*

- For any Scheme value $(SM_L^\tau v)$ (where $\tau \neq L$) and for any top-level evaluation context $\mathcal{E}[]_S$:
 - If $v = \bar{n}$, then $\mathcal{E}[\mathcal{T}_S^\tau(SM_L^\tau v)]_M \rightarrow^* \mathcal{E}[v]_M$.
 - If $v \neq \bar{n}$, then $\mathcal{E}[\mathcal{T}_S^\tau(SM_L^\tau v)]_M \rightarrow^* \text{Error}$: s for some s .
- For any ML value $({}^L MS_L v)$ and for any top-level evaluation context $\mathcal{E}[]_M$:
 - If $v = \bar{n}$, then $\mathcal{E}[\mathcal{T}_M^\iota({}^L MS_L v)] \rightarrow^* \mathcal{E}[v]$.
 - If $v \neq \bar{n}$, then $\mathcal{E}[\mathcal{T}_M^\iota({}^L MS_L v)] \rightarrow^* \text{Error}$: s for some s .

Proof. The error cases hold by calculation of the reductions; the non-error cases require induction on the number n . \square