



HAL
open science

FabULous Interoperability for ML and a Linear Language

Gabriel Scherer, Max New, Nick Rioux, Amal Ahmed

► **To cite this version:**

Gabriel Scherer, Max New, Nick Rioux, Amal Ahmed. FabULous Interoperability for ML and a Linear Language. International Conference on Foundations of Software Science and Computation Structures (FoSSaCS), Apr 2018, Thessaloniki, Greece. 10.1007/978-3-319-89366-2_8 . hal-01929158

HAL Id: hal-01929158

<https://inria.hal.science/hal-01929158>

Submitted on 21 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FabULous Interoperability for ML and a Linear Language

Gabriel Scherer¹², Max New¹, Nick Rioux¹, and Amal Ahmed¹³

¹ Northeastern University, US

² INRIA Saclay, France

³ INRIA Paris, France

Abstract. Instead of a monolithic programming language trying to cover all features of interest, some programming systems are designed by combining together simpler languages that cooperate to cover the same feature space. This can improve usability by making each part simpler than the whole, but there is a risk of *abstraction leaks* from one language to another that would break expectations of the users familiar with only one or some of the involved languages.

We propose a formal specification for what it means for a given language in a multi-language system to be usable without leaks: it should embed into the multi-language in a *fully abstract* way, that is, its contextual equivalence should be unchanged in the larger system.

To demonstrate our proposed design principle and formal specification criterion, we design a multi-language programming system that combines an ML-like statically typed functional language and another language with linear types and linear state. Our goal is to cover a good part of the expressiveness of languages that mix functional programming and linear state (ownership), at only a fraction of the complexity. We prove that the embedding of ML into the multi-language system is fully abstract: functional programmers should not fear abstraction leaks. We show examples of combined programs demonstrating in-place memory updates and safe resource handling, and an implementation extending OCaml with our linear language.

1 Introduction

Feature accretion is a common trend among mature but actively evolving programming languages, including C++, Haskell, Java, OCaml, Python, and Scala. Each new feature strives for generality and expressiveness, and may provide a large usability improvement to users of the particular problem domain or programming style it was designed to empower (e.g., XML documents, asynchronous communication, staged evaluation). But feature creep in general-purpose languages may also make it harder for programmers to master the language as a whole, degrade the user experience (e.g., leading to more cryptic error messages), require additional work on the part of tooling providers, and lead to fragility in language implementations.

A natural response to increased language complexity is to define subsets of the language designed for a better programming experience. For instance, a subset can be easier to teach (e.g., “Core” ML⁴, Haskell 98 as opposed to GHC Haskell, Scala mastery levels⁵); it can facilitate static analysis or decrease the risk of programming errors, while remaining sufficiently expressive for the target users’ needs (e.g., MISRA C, Spark/Ada); it can enforce a common style within a company; or it can be designed to encourage a transition to deprecate some ill-behaved language features (e.g., strict Javascript).

Once a subset has been selected, it may be the case that users write whole programs purely in the subset (possibly using tooling to enforce that property), but programs will commonly rely on other libraries that are not themselves implemented in the same subset of the language. If users stay in the subset while using these libraries, they will only interact with the part of the library whose interface is expressible in the subset. But does the behavior of the library respect the expectations of users who only know the subset? When calling a function from within the subset breaks subset expectations, it is a sign of *leaky abstraction*.

How should we design languages with useful subsets that manage complexity and avoid abstraction leaks?

We propose to look at this question from a different, but equivalent, angle: instead of designing a single big monolithic language with some nicer subsets, we propose to consider *multi-language* programming systems where several smaller programming languages interact together to cover the same feature space. Each language or sub-combination of languages is a subset, in the above sense, of the multi-language, and there is a clear definition of *abstraction leaks* in terms of user experience: a user who only knows some of the languages of the system should be able to use the multi-language system, interacting with code written in the other languages, without have their expectations violated. If we write a program in Java and call a function that, internally, is implemented in Scala, there should be no surprises—our experience should be the same as when calling a pure Java function. Similarly, consider the subset of Haskell that does not contain IO (input-output as a type-tracked effect): the expectations of a user of this language, for instance in terms of valid equational reasoning, should not be violated by adding IO back to the language—in the absence of the abstraction-leaking `unsafePerformIO`.

We propose a *formal specification* for a “no abstraction leaks” guarantee that can be used as a design criterion to design new multi-language systems, with graceful interoperation properties. It is based on the formal notion of *full abstraction* which has previously been used to study the denotational semantics of programming languages [Meyer and Sieber, 1988, Milner, 1977, Cartwright and Felleisen, 1992, Jeffrey and Rathke, 2005, Abramsky, Jagadeesan, and Malacaria, 2000], and the formal property of compilers [Ahmed and Blume, 2008, 2011, Devriese, Patrignani, and Piessens, 2016, New, Bowman, and Ahmed, 2016, Patrignani, Agten, Strackx, Jacobs, Clarke, and Piessens, 2015], but not for user-facing languages. A compiler C from a source language S to a target language T

⁴ <https://caml.inria.fr/pub/docs/u3-ocaml/ocaml-ml.html>

⁵ <http://www.scala-lang.org/old/node/8610>

is *fully abstract* if, whenever two source terms s_1 and s_2 are indistinguishable in S , their translations $C(s_1)$ and $C(s_2)$ are indistinguishable in T . In a multi-language $G + E$ formed of a general-purpose, user-friendly language G and a more advanced language E —one that provides an *escape hatch* for *experts* to write code that can't be implemented in G —we say that E does not *leak* into G if the embedding of G into the multi-language $G + E$ is fully abstract.

To demonstrate that our formal specification is reasonable, we design a novel multi-language programming system that satisfies it. Our multi-language λ^{UL} combines a general-purpose functional programming language λ^{U} (unrestricted) of the ML family with an advanced language λ^{L} (linear) with *linear types* and linear state. It is less convenient to program in λ^{L} 's restrictive type system, but users can write programs in λ^{L} that could not be written in λ^{U} : they can use linear types, locally, to enforce resource usage protocols (typestate), and they can use linear state and the linear ownership discipline to write programs that do in-place update to allocate less memory, yet remain observationally pure.

Consider for example the following mixed-language program. The blue fragments are written in the general-purpose, user-friendly functional language, while the red fragments are written in the linear language. The boundaries `UL` and `LU` allow switching between languages. The program reads all lines from a file, accumulating them in a list, and concatenating it into a single string when the end-of-file (EOF) is reached.

```
let concat_lines path : String = UL(
  loop (open LU(path)) LU(Nil)
  where rec loop handle LU(acc : List String) =
    match line handle with
    | Next line LU(handle) -> loop handle LU(Cons line acc)
    | EOF handle -> close handle; LU(rev_concat "\n" acc))
```

The linear type system ensures that the file handle is properly closed: removing the `close handle` call would give a type error. On the other hand, only the parts concerned with the resource-handling logic need to be written in the red linear language; the user can keep all general-purpose logic (here, how to accumulate lines and what to do with them at the end) in the more convenient general-purpose blue language—and call this function from a blue-language program. Fine-grained boundaries allow users to rely on each language's strength and to use the advanced features only when necessary.

In this example, the file-handle API specifies that the call to `line`, which reads a line, returns the data at type `![String]`. The latter represents how `U` values of type `String` can be put into a *lump* type to be passed to the linear world where they are treated as opaque blackboxes that must be passed back to the ML world for consumption. For other examples, such as in-place list manipulation or transient operations on an persistent data structure, we will need a deeper form of interoperability where the linear world creates, dissects or manipulates `U` values. To enable this, our multi-language supports translation of types from one language to the other, using a *type compatibility* relation $\sigma \simeq \sigma$ between λ^{U} types σ and λ^{L} types σ .

We claim the following contributions:

1. We propose a formal specification of what it means for advanced language features to be introduced in a (multi-)language system without introducing a class of abstraction leaks that break equational reasoning. This specification captures a useful *usability* property, and we hope it will help us and others design more usable programming languages, much like the formal notion of *principal types* served to better understand and design type inference systems.
2. We design a simple linear language, λ^L , that supports linear state (Section 2). This simple design for linear state is a contribution of its own. A nice property of the language (shared by some other linear languages) is that the code has both an imperative interpretation—with in-place memory update, which provides resource guarantees—and a functional interpretation—which aids program reasoning. The imperative and functional interpretations have different resource usage, but the same input/output behavior.
3. We present a multi-language programming system λ^{UL} combining a core ML language, λ^U (U for Unrestricted, as opposed to Linear) with λ^L and prove that the embedding of the ML language λ^U in λ^{UL} is fully abstract (Section 3). Moreover, the multi-language is designed to ensure that our full abstraction result is stable under extension of the embedded ML language λ^U .
4. We define a logical relation and prove parametricity for λ^{UL} . The logical relation illustrates, semantically, why one can reason functionally about programs in λ^{UL} despite the presence of state and strong updates (Section 4).
5. We evaluate the resulting language design by providing examples of hybrid λ^{UL} programs that exhibit programming patterns inaccessible to ML alone, such that safe in-place updates and typestate-like static protocol enforcement (Section 5).

2 The λ^U and λ^L languages

Types $\sigma ::= \alpha \mid \sigma_1 \times \sigma_2 \mid \mathbf{1} \mid \sigma_1 \rightarrow \sigma_2 \mid \sigma_1 + \sigma_2 \mid \mu\alpha.\sigma \mid \forall\alpha.\sigma$
Expr. $e ::= x \mid \langle e_1, e_2 \rangle \mid \pi_1 e \mid \pi_2 e \mid \langle \rangle \mid e_1; e_2 \mid \lambda(x:\sigma).e \mid e_1 e_2 \mid$
 $\text{inj}_j e \mid \text{case } e' \text{ of } x_1.e_1 \mid x_2.e_2 \mid \text{fold}_{\mu\alpha.\sigma} e \mid \text{unfold } e \mid \Lambda\alpha.e \mid e[\sigma]$
Values $v ::= x \mid \langle v_1, v_2 \rangle \mid \langle \rangle \mid \lambda(x:\sigma).e \mid \text{inj}_1 v \mid \text{inj}_2 v \mid \text{fold}_{\mu\alpha.\sigma} v \mid \Lambda\alpha.v$
Contexts $\Gamma ::= \cdot \mid \Gamma, x:\sigma \mid \Gamma, \alpha$

Fig. 1. Unrestricted Language: Syntax

The unrestricted language λ^U is a run-of-the-mill idealized ML language with functions, pairs, sums, iso-recursive types and polymorphism. It is presented in its explicitly typed form—we will not discuss type inference in this work. The full syntax is described in Figure 1, and the typing rules in Figure 2. The dynamic semantics is completely standard. Having binary sums, binary products and iso-recursive types lets us express algebraic datatypes in the usual way.

$$\boxed{\Gamma \vdash_{\text{U}} e : \sigma}$$

$$\begin{array}{c}
\frac{x : \sigma \in \Gamma}{\Gamma \vdash_{\text{U}} x : \sigma} \qquad \frac{}{\Gamma \vdash_{\text{U}} \langle \rangle : \mathbf{1}} \qquad \frac{\Gamma \vdash_{\text{U}} e : \mathbf{1} \quad \Gamma \vdash_{\text{U}} e' : \sigma}{\Gamma \vdash_{\text{U}} e; e' : \sigma} \\
\frac{\Gamma \vdash_{\text{U}} e_1 : \sigma_1 \quad \Gamma \vdash_{\text{U}} e_2 : \sigma_2}{\Gamma \vdash_{\text{U}} \langle e_1, e_2 \rangle : \sigma_1 \times \sigma_2} \qquad \frac{\Gamma \vdash_{\text{U}} e : \sigma_1 \times \sigma_2}{\Gamma \vdash_{\text{U}} \pi_i e : \sigma_i} \\
\frac{\Gamma, x : \sigma \vdash_{\text{U}} e : \sigma'}{\Gamma \vdash_{\text{U}} \lambda(x : \sigma). e : \sigma \rightarrow \sigma'} \qquad \frac{\Gamma \vdash_{\text{U}} e : \sigma' \rightarrow \sigma \quad \Gamma \vdash_{\text{U}} e' : \sigma'}{\Gamma \vdash_{\text{U}} e e' : \sigma} \\
\frac{\Gamma \vdash_{\text{U}} e : \sigma_i}{\Gamma \vdash_{\text{U}} \text{inj}_i e : \sigma_1 + \sigma_2} \qquad \frac{\Gamma \vdash_{\text{U}} e : \sigma_1 + \sigma_2 \quad \Gamma, x_1 : \sigma_1 \vdash_{\text{U}} e_1 : \sigma \quad \Gamma, x_2 : \sigma_2 \vdash_{\text{U}} e_2 : \sigma}{\Gamma \vdash_{\text{U}} \text{case } e \text{ of } x_1. e_1 \mid x_2. e_2 : \sigma} \\
\frac{\Gamma \vdash_{\text{U}} e : \sigma[\mu\alpha. \sigma/\alpha]}{\Gamma \vdash_{\text{U}} \text{fold}_{\mu\alpha. \sigma} e : \mu\alpha. \sigma} \qquad \frac{\Gamma \vdash_{\text{U}} e : \mu\alpha. \sigma}{\Gamma \vdash_{\text{U}} \text{unfold } e : \sigma[\mu\alpha. \sigma/\alpha]} \\
\frac{\Gamma, \alpha \vdash_{\text{U}} v : \sigma}{\Gamma \vdash_{\text{U}} \Lambda\alpha. v : \forall\alpha. \sigma} \qquad \frac{\Gamma \vdash_{\text{U}} e : \forall\alpha. \sigma \quad \Gamma \vdash_{\text{U}} \sigma'}{\Gamma \vdash_{\text{U}} e[\sigma'] : \sigma[\sigma'/\alpha]}
\end{array}$$

Fig. 2. Unrestricted Language: Static Semantics

The novelty lies in the linear language λ^{L} , which we present in several steps. As is common in λ -calculi with references, the small-step operational semantics is given for a language that is not exactly the surface language in which programs are written, because memory allocation returns *locations* ℓ that are not in the grammar of surface terms. Reductions are defined on *configurations*, a local store paired with a term in a slightly larger *internal* language. We have two type systems, a type system on surface terms, that does not mention locations and stores—which is the one a programmer needs to know—and a type system on configurations, which contains enough static information to reason about the dynamics of our language and prove subject reduction. Again, this follows the standard structure of syntactic soundness proofs for languages with a mutable store.

We present the surface language and type system in [Section 2.1](#), except for the language fragment manipulating the linear store which is presented in [Section 2.2](#). Finally, the internal terms, their typing and reduction semantics are presented in [Section 2.3](#).

2.1 The Core of λ^{L}

[Figure 3](#) presents the surface syntax of our linear language λ^{L} . For the syntactic categories of types σ , and expressions e , the last line contains the constructions related to the linear store that we only discuss in [Section 2.2](#).

<i>Types</i>	$\sigma ::= \sigma_1 \otimes \sigma_2 \mid \mathbf{1} \mid \sigma_1 \multimap \sigma_2 \mid \sigma_1 \oplus \sigma_2 \mid \mu\alpha.\sigma \mid \alpha \mid !\sigma \mid \text{Box } 1 \sigma \mid \text{Box } 0$
<i>Expr.</i>	$e ::= x \mid \langle e_1, e_2 \rangle \mid \text{let } \langle v_1, v_2 \rangle = e_1 \text{ in } e_2 \mid \langle \rangle \mid e_1; e_2 \mid \lambda(x:\sigma).e \mid e_1 e_2 \mid$ $\text{inj}_1 e \mid \text{inj}_2 e \mid \text{case } e' \text{ of } x_1.e_1 \mid x_2.e_2 \mid \text{fold}_{\mu\alpha.\sigma} e \mid \text{unfold } e \mid$ $\text{share } e \mid \text{copy } e \mid \text{new } e \mid \text{free } e \mid \text{box } e \mid \text{unbox } e$
<i>Values</i>	$v ::= x \mid \langle v_1, v_2 \rangle \mid \langle \rangle \mid \lambda(x:\sigma).e \mid \text{inj}_1 v \mid \text{inj}_2 v \mid \text{fold}_{\mu\alpha.\sigma} v \mid \text{share } v$
<i>Contexts</i>	$\Gamma ::= \cdot \mid \Gamma, x:\sigma$

Fig. 3. Linear Language: Surface Syntax

In technical terms, our linear type system is exactly propositional intuitionistic linear logic, extended with iso-recursive types. For simplicity and because we did not need them, our current system also does not have polymorphism or additive/lazy pairs $\sigma_1 \& \sigma_2$. Additive pairs would be a trivial addition, but polymorphism would require more work when we define the multi-language semantics in [Section 3](#).

In less technical terms, our type system can enforce that values be used *linearly*, meaning that they cannot be duplicated or erased, they have to be deconstructed exactly once. Only some types have this linearity restriction; others allow duplication and sharing of values at will. We can think of linear values as *resources* to be spent wisely; for any linear value somewhere in a term, there can be only one way to access this value, so we can interpret the language as enforcing an *ownership* discipline where whoever points to a linear value owns it.

The types of linear values are the type of linear pairs $\sigma_1 \otimes \sigma_2$, of linear disjoint unions $\sigma_1 \oplus \sigma_2$, of linear functions $\sigma_1 \multimap \sigma_2$, and of the linear unit type $\mathbf{1}$. For example, a linear function must be called exactly once, and its result must in turn be consumed – such linear functions can safely capture linear resources. The expression-formers at these types use the same syntax as the unrestricted language λ^U , with the exception of linear pair deconstruction $\text{let } \langle v_1, v_2 \rangle = e_1 \text{ in } e_2$, which names both members of the deconstructed pair at once. A linear pair type with projection would only ever allow to observe one of the two members; this would correspond to the additive/lazy pairs $\sigma_1 \& \sigma_2$, where only one of the two members is ever computed.

The types of non-linear, duplicable values are the types of the form $!\sigma$ —the *exponential* modality of linear logic. If e has type σ , the term $\text{share } e$ has type $!\sigma$. Values of this type are not uniquely owned, they can be shared at will. If the term e has duplicable type $!\sigma$, then the term $\text{copy } e$ has type σ : this creates a local copy of the value that is uniquely-owned by its receiver and must be consumed linearly.

This resource-usage discipline is enforced by the surface typing rules of λ^L , presented in [Figure 4](#). They are exactly the standard (two-sided) logical rules of intuitionistic linear logic, annotated with program terms. The non-duplicability of linear values is enforced by the way contexts are merged by the inference

$$\boxed{\Gamma_1 \Downarrow \Gamma_2}$$

$$\begin{aligned} (\Gamma_1, x : !\sigma) \Downarrow (\Gamma_2, x : !\sigma) &\stackrel{\text{def}}{=} (\Gamma_1 \Downarrow \Gamma_2), x : !\sigma \\ (\Gamma_1, x : \sigma) \Downarrow \Gamma_2 &\stackrel{\text{def}}{=} (\Gamma_1 \Downarrow \Gamma_2), x : \sigma \quad (x \notin \Gamma_2) \\ \Gamma_1 \Downarrow (\Gamma_2, x : \sigma) &\stackrel{\text{def}}{=} (\Gamma_1 \Downarrow \Gamma_2), x : \sigma \quad (x \notin \Gamma_1) \end{aligned}$$

$$\boxed{\Gamma \vdash_L e : \sigma}$$

$$\begin{array}{c} \frac{}{! \Gamma, x : \sigma \vdash_L x : \sigma} \quad \frac{\Gamma_1 \vdash_L e_1 : \sigma_1 \quad \Gamma_2 \vdash_L e_2 : \sigma_2}{\Gamma_1 \Downarrow \Gamma_2 \vdash_L (e_1, e_2) : \sigma_1 \otimes \sigma_2} \\ \\ \frac{\Gamma \vdash_L e : \sigma_1 \otimes \sigma_2 \quad \Gamma', x_1 : \sigma_1, x_2 : \sigma_2 \vdash_L e' : \sigma}{\Gamma \Downarrow \Gamma' \vdash_L \text{let } (x_1, x_2) = e \text{ in } e' : \sigma} \\ \\ \frac{}{! \Gamma \vdash_L \langle \rangle : \mathbf{1}} \quad \frac{\Gamma \vdash_L e : \mathbf{1} \quad \Gamma' \vdash_L e' : \sigma}{\Gamma \Downarrow \Gamma' \vdash_L e; e' : \sigma} \quad \frac{\Gamma, x : \sigma \vdash_L e : \sigma'}{\Gamma \vdash_L \lambda(x : \sigma). e : \sigma \multimap \sigma'} \\ \\ \frac{\Gamma \vdash_L e : \sigma' \multimap \sigma \quad \Gamma' \vdash_L e' : \sigma'}{\Gamma \Downarrow \Gamma' \vdash_L e e' : \sigma} \quad \frac{\Gamma \vdash_L e : \sigma_i}{\Gamma \vdash_L \text{inj}_i e : \sigma_1 \oplus \sigma_2} \\ \\ \frac{\Gamma \vdash_L e : \sigma_1 \oplus \sigma_2 \quad \Gamma', x_1 : \sigma_1 \vdash_L e_1 : \sigma \quad \Gamma', x_2 : \sigma_2 \vdash_L e_2 : \sigma}{\Gamma \Downarrow \Gamma' \vdash_L \text{case } e \text{ of } x_1. e_1 \mid x_2. e_2 : \sigma} \quad \frac{}{! \Gamma \vdash_L \text{share } e : !\sigma} \quad \frac{\Gamma \vdash_L e : !\sigma}{\Gamma \vdash_L \text{copy } e : \sigma} \\ \\ \begin{array}{ccc} \text{unfold} & \text{new} & \text{unbox} \\ \mu\alpha. \sigma \quad \frac{}{\circ} \quad \sigma[\mu\alpha. \sigma / \alpha] & \mathbf{1} \quad \frac{}{\circ} \quad \text{Box } 0 & \text{Box } 1 \quad \sigma \quad \frac{}{\circ} \quad (\text{Box } 0) \otimes \sigma \\ \text{fold}_{\mu\alpha. \sigma} & \text{free} & \text{box} \end{array} \end{array}$$

Fig. 4. Linear Language: Surface Static Semantics

rules: if e_1 is type-checked in the context Γ_1 and e_2 in Γ_2 , then the linear pair $\langle e_1, e_2 \rangle$ is only valid in the combined context $\Gamma_1 \curlywedge \Gamma_2$. The (\curlywedge) operation is partial; this combined context is defined only if the variables shared by Γ_1 and Γ_2 are duplicable—their type is of the form $!\sigma$. In other words, a variable at a non-duplicable type in $\Gamma_1 \curlywedge \Gamma_2$ cannot possibly appear in both Γ_1 and Γ_2 : it must appear exactly once⁶.

A good way to think of the linear judgment $\Gamma \vdash_L e : \sigma$ is that the evaluation of e *consumes* the linear variables of Γ ; it is thus natural that the strict pair $\langle e_1, e_2 \rangle$ would need separate sets of resources Γ_1 and Γ_2 , as it evaluates both members to return a value. On the other hand, case elimination `case e of $x_1. e_1 \mid x_2. e_2$` reuses the same context Γ' in both branches e_1 and e_2 : only one will be evaluated, so they do not compete for resources.

The variable rule does not expect a context of the form $\Gamma, x : \sigma$ but of the form $!\Gamma, x : \sigma$. Here $!\Gamma$ is a notation for the pointwise application of the $(!)$ connective to all the types in Γ —i.e., all types in $!\Gamma$ are of the form $!\sigma$. This means that the variable rule can only be used when all variables in the context are duplicable, except maybe the variable that is being used. A context of the form $\Gamma, x : \sigma$ would allow us to forget some variable present in the context; in our judgment $\Gamma \vdash_L e : \sigma$, all non-duplicable variables in Γ must appear (once) in e .

The form $!\Gamma$ is also used in the typing rule for `share e`: a term can only be made duplicable if it does not depend on linear resources from the context. Otherwise, duplicating the shared value could break the unique-ownership discipline on these linear resources.

Finally, the linear isomorphism notation for `fold` and `unfold` in Figure 4 defines them as primitive functions, at the given linear function type, in the empty context – using them does not consume resources. This notation also means that, operationally, these two operations shall be inverses of each other. The rules for the linear store type `Box 1 σ` and `Box 0` are described in Section 2.2.

Lemma 1 (Context joining properties) *Context joining (\curlywedge) is partial but associative and commutative. In particular, if $(\Gamma_1 \curlywedge \Gamma_2) \curlywedge \Gamma'$ is defined, then both $\Gamma_i \curlywedge \Gamma'$ are defined.*

2.2 Linear Memory in λ^L

The surface typing rules for the linear store are given at the end of Figure 4. The linear type `Box 1 σ` represents a memory location that holds a value of type σ . The type `Box 0` represents a location that has been allocated, but does not currently hold a value. The primitive operations to act on this type are given as linear isomorphisms: `new` allocates, turning a unit value into an empty location; conversely, `free` reclaims an empty location. Putting a value into the location and taking it out are expressed by `box` and `unbox`, which convert between a pair of an empty location and a value, of type $(\text{Box } 0) \otimes \sigma$, and a full location, of type `Box 1 σ` .

⁶ Standard presentations of linear logic force contexts to be completely distinct, but have a separate rule to duplicate linear variables, which is less natural for programming.

For example, the following program takes a full reference and a value, and swaps the value with the content of the reference:

```
 $\lambda(p: (\text{Box } 1 \sigma) \otimes \sigma). \text{let } \langle r, x \rangle = p \text{ in let } \langle l, x_l \rangle = \text{unbox } r \text{ in } \langle \text{box } \langle l, x \rangle, x_l \rangle$ 
```

The programming style following from this presentation of linear memory is functional, or applicative, rather than imperative. Rather than insisting on the mutability of references—which is allowed by the linear discipline—we may think of the type $\text{Box } 1 \sigma$ as representing the indirection through the heap that is implicit in functional programs. In a sense, we are not writing imperative programs with a mutable store, but rather making explicit the allocations and dereferences happening in higher-level purely functional language. In this view, empty cells allow memory reuse.

This view that $\text{Box } 1 \sigma$ represents indirection through the memory suggests we can encode lists of values of type σ by the type $\text{LinList } \sigma \stackrel{\text{def}}{=} \mu\alpha. 1 \oplus \text{Box } 1 (\sigma \otimes \alpha)$. The placement of the box inside the sum mirrors the fact that empty list is represented as an immediate value in functional languages. From this type definition, one can write an in-place reverse function on lists of σ as follows:

```
 $\text{fix } \lambda(\text{rev\_into} : \text{LinList } \sigma \multimap \text{LinList } \sigma \multimap \text{LinList } \sigma).$   
 $\lambda(xs : \text{LinList } \sigma). \lambda(\text{acc} : \text{LinList } \sigma).$   
 $\text{case unfold } xs \text{ of}$   
 $| y. (y; \text{acc})$   
 $| y. \text{let } \langle l, p \rangle = \text{unbox } y \text{ in}$   
 $\text{let } \langle xs, x \rangle = p \text{ in}$   
 $\text{rev\_into } xs (\text{fold } (\text{inj}_2 (\text{box } \langle l, \langle x, \text{acc} \rangle)))$ 
```

This definition uses a fixpoint operator fix that can be defined, in the standard way, using the iso-recursive type $\mu\alpha. \alpha \multimap \sigma \multimap \sigma'$ of the strict fixpoint combinator on functions $\sigma \multimap \sigma'$.

Our linear language λ^L is a formal language that is not terribly convenient to program directly. We will not present a full surface language in this work, but one could easily define syntactic sugar to write the exact same function as follows:

```
 $\text{rev\_into Nil} \quad \text{acc} = \text{acc}$   
 $\text{rev\_into } (\text{Cons } \langle x, xs \rangle @ l) \text{ acc} = \text{rev\_into } xs (\text{Cons } \langle x, \text{acc} \rangle @ l)$ 
```

One can read this function as the usual functional rev_append function on lists, annotated with memory reuse information: if we assume we are the unique owner of the input list and won't need it anymore, we can reuse the memory of its cons cells (given in this example the name l) to store the reversed list. On the other hand, if you read the box and unbox as imperative operations, this code expresses the usual imperative pointer-reversal algorithm.

This double view of linear state occurs in other programming systems with linear state. It was recently emphasized in [O'Connor, Chen, Rizkallah, Amani, Lim, Murray, Nagashima, Sewell, and Klein \[2016\]](#), where the functional point of view is seen as easing formal verification, while the imperative view is used as a compilation technique to produce efficient C code from linear programs.

2.3 Internal λ^L Syntax and Typing

To give a dynamic semantics for λ^L and prove it sound, we need to extend the language with explicit stores and store locations. Indeed, the allocating term `new $\langle \rangle$` should reduce to a “fresh location” ℓ allocated in some store s , and neither are part of the surface-language syntax. The corresponding internal typing judgment is more complex, but note that users do not need to know about it to reason about correctness of surface programs. The internal typing is essential for the soundness proof, but also useful for defining the multi-language semantics in [Section 3](#).

The syntax of internal terms and the internal type system are presented in [Figure 5](#). Reduction will be defined on *configurations* $(s \mid e)$, which are pairs of a store s and a term e . Stores s map *locations* ℓ to either nothing (the location is empty), written $[\ell \mapsto \cdot]$, or a value paired with its own local store, noted $[\ell \mapsto (s \mid v)]$. Having local stores in this way, instead of a single global store as is typical in formalizations of ML, directly expresses the idea of “memory ownership” in the syntax: a term e “owns” the locations that appear in it, and a configuration $(s \mid e)$ is only well-typed if the domain of s is exactly those locations. Each store slot, in turn, may contain a value and the local store owned by the value; in particular, passing a full location of type **Box 1** σ transfers ownership of the location, but also of the store fragment captured by the value.

Our internal typing judgment $\Psi; \Gamma \vdash_L s \mid e : \sigma$ checks configurations, not just terms, and relies not only on a typing context for variables Γ but also on a *store typing* Ψ , which maps the locations of the configuration to typing assumptions of two forms: $(\cdot; \vdash \ell : \mathbf{Box\ 0})$ indicates that ℓ must be empty in the configuration, and $(\Gamma; \Psi \vdash \ell : \mathbf{Box\ 1\ } \sigma)$ indicates that ℓ is full, and that the value it contains owns a local store of type s and the resources in Γ .

Just as linear variables must occur exactly once in a term, locations have linear types and thus occur exactly once in a term. Our typing judgment uses disjoint store typings $\Psi_1 \uplus \Psi_2$ to enforce this linearity. Similarly, leaf rules such as the variable, unit, and location rules enforce that both the store typing and the store be empty, which enforces that all locations are used in the term.

Locations ℓ are always linear, never duplicable. To allow sharing terms that contain locations, the internal language uses the internal construction `share($s : \Psi$). e` , that *captures* a local store $s : \Psi$. This notation is a binding construct: the locations in s are bound by this shared term, and not visible outside this term. In particular, the typing rule for `share($s : \Psi$). e` checks the term e in the store s , but it is itself only valid paired with an empty store, under the empty store typing. When new copies of a shared term are made, the local store is copied as well: this is necessary to guarantee that locations remain linear—and for correctness of linear state update.

The typing rule for functions `$\lambda(x : \sigma).e$` lets function bodies use an arbitrary store typing Ψ . This would be unsound if our functions were duplicable, but it is a natural and expressive choice for linear, one-shot functions. To make a function duplicable, one can share it at type $!(\sigma \multimap \sigma')$, whose values are of the canonical

<i>Types</i>	σ	(unchanged from Figure 4)
<i>Expressions</i>	$e + ::= \dots \mid \ell \mid \text{share}(s:\Psi).e$	
		with $\text{share } e \stackrel{\text{def}}{=} \text{share}(\emptyset:\cdot).e$
<i>Values</i>	$v + ::= \dots \mid \ell \mid \text{share}(s:\Psi).v$	
<i>Store</i>	$s ::= \emptyset \mid s[\ell \mapsto (s \mid v)] \mid s[\ell \mapsto \cdot]$	
<i>Configurations</i>	$::= (s \mid e)$	
<i>Store typing</i>	$\Psi ::= \cdot \mid \Psi, (\cdot; \vdash \ell : \text{Box } 0)$	
		$\mid \Psi, (\Psi'; \Gamma \vdash \ell : \text{Box } 1 \sigma)$

$s_1 \# s_2$	Union of stores on disjoint locations	$\Psi_1 \uplus \Psi_2$
Union of store typings on disjoint locations		
$\Psi; \Gamma \vdash_L s \mid e : \sigma$	$\Gamma \vdash_L e : \sigma \stackrel{\text{def}}{=} \cdot; \Gamma \vdash_L \emptyset \mid e : \sigma$	
$\frac{\Psi_1; \Gamma_1 \vdash_L s_1 \mid e_1 : \sigma_1 \quad \Psi_2; \Gamma_2 \vdash_L s_2 \mid e_2 : \sigma_2}{\Psi_1 \uplus \Psi_2; \Gamma_1 \curlywedge \Gamma_2 \vdash_L s_1 \# s_2 \mid (e_1, e_2) : \sigma_1 \otimes \sigma_2}$		
$\frac{\Psi; \Gamma \vdash_L s \mid e : \sigma_1 \otimes \sigma_2 \quad \Psi'; \Gamma', x_1 : \sigma_1, x_2 : \sigma_2 \vdash_L s' \mid e' : \sigma}{\Psi \uplus \Psi'; \Gamma \curlywedge \Gamma' \vdash_L s \# s' \mid \text{let } (x_1, x_2) = e \text{ in } e' : \sigma}$		
$\frac{\cdot; !\Gamma, x : \sigma \vdash_L \emptyset \mid x : \sigma \quad \cdot; !\Gamma \vdash_L \emptyset \mid \langle \rangle : 1 \quad \Psi; \Gamma \vdash_L s \mid e : 1 \quad \Psi'; \Gamma' \vdash_L s' \mid e' : \sigma}{\Psi \uplus \Psi'; \Gamma \curlywedge \Gamma' \vdash_L s \# s' \mid e; e' : \sigma}$		
$\frac{\Psi; \Gamma, x : \sigma \vdash_L s \mid e : \sigma'}{\Psi; \Gamma \vdash_L s \mid \lambda(x:\sigma).e : \sigma \multimap \sigma'} \quad \frac{\Psi; \Gamma \vdash_L s \mid e : \sigma' \multimap \sigma \quad \Psi'; \Gamma' \vdash_L s' \mid e' : \sigma'}{\Psi \uplus \Psi'; \Gamma \curlywedge \Gamma' \vdash_L s \# s' \mid e; e' : \sigma}$		
$\frac{\Psi; \Gamma \vdash_L s \mid e : \sigma_i}{\Psi; \Gamma \vdash_L s \mid \text{inj}_i e : \sigma_1 \oplus \sigma_2}$		
$\frac{\Psi; \Gamma \vdash_L s \mid e : \sigma_1 \oplus \sigma_2 \quad \Psi'; \Gamma', x_1 : \sigma_1 \vdash_L s' \mid e_1 : \sigma \quad \Psi'; \Gamma', x_2 : \sigma_2 \vdash_L s' \mid e_2 : \sigma}{\Psi \uplus \Psi'; \Gamma \curlywedge \Gamma' \vdash_L s \# s' \mid \text{case } e \text{ of } x_1.e_1 \mid x_2.e_2 : \sigma}$		
$\frac{\Psi; !\Gamma \vdash_L s \mid e : \sigma}{\cdot; !\Gamma \vdash_L \emptyset \mid \text{share}(s:\Psi).e : !\sigma} \quad \frac{\Psi; \Gamma \vdash_L s \mid e : !\sigma}{\Psi; \Gamma \vdash_L s \mid \text{copy } e : \sigma}$		
$\frac{(\cdot; \cdot \vdash \ell : \text{Box } 0); !\Gamma \vdash_L [\ell \mapsto \cdot] \mid \ell : \text{Box } 0}{\Psi; \Gamma \vdash_L s \mid v : \sigma}$		
$\frac{\Psi; \Gamma \vdash_L s \mid v : \sigma}{(\Psi; \Gamma \vdash \ell : \text{Box } 1 \sigma); \Gamma \curlywedge \Gamma' \vdash_L [\ell \mapsto (s \mid v)] \mid \ell : \text{Box } 1 \sigma}$		
$\frac{\text{unfold}}{\text{fold}_{\mu\alpha.\sigma}} \mu\alpha.\sigma \overset{\circ}{\longleftarrow} \sigma[\mu\alpha.\sigma/\alpha]$	$\frac{\text{new}}{\text{free}} 1 \overset{\circ}{\longleftarrow} \text{Box } 0$	$\frac{\text{unbox}}{\text{box}} \text{Box } 1 \sigma \overset{\circ}{\longleftarrow} (\text{Box } 0) \otimes \sigma$

Fig. 5. Linear Language: Internal Static Semantics

form $\text{share}(s; \Psi). \lambda(x: \sigma). e$. It is the sharing construct, not the function itself, that closes over the local store.

With the macro-expansion $\text{share } e \stackrel{\text{def}}{=} \text{share}(\emptyset; \cdot). e$, any term e of the surface language (Figure 3) can be seen as a term of the internal language (Figure 5). In particular, we can prove that the surface and internal typing judgments coincide on surface terms.

Lemma 2 *If e is a surface term of λ^L , then the surface judgment $\Gamma \vdash_L e : \sigma$ holds if and only if the internal judgment $\cdot; \Gamma \vdash_L \emptyset \mid e : \sigma$ holds.*

The following technical results are used in the soundness proof for the language – the subject-reduction result.

Lemma 3 (Inversion principle for λ^L values) *In any complete derivation of $\Psi; \Gamma \vdash_L s \mid v : \sigma$, either v is a variable x , or the derivation starts with the introduction rule for σ .*

For example, if we have $\Psi; \Gamma \vdash_L s \mid v : !\sigma$, then we know that v is either a variable or of the form $\text{share}(s'; \Psi'). v'$ for some v' , but also that $s = \emptyset$, $\Psi = \cdot$ and that Γ is of the form $! \Gamma'$ for some Γ' . The latter is immediate if v is $\text{share}(s'; \Psi'). v'$, and also holds if v is a variable.

Lemma 4 (Weakening by duplicable contexts) *$\Psi; \Gamma' \vdash_L s \mid e : \sigma$ implies $\Psi; ! \Gamma, \Gamma' \vdash_L s \mid e : \sigma$.*

2.4 Reduction of Internal Terms

Figure 6 gives a small-step operational semantics for the internal terms of λ^L . We separate the head reductions (\rightsquigarrow) from reductions in depth (\xrightarrow{L}). The head reduction of the linear types of the core language do not involve the store and are standard. For the store primitives of Figure 4 acting on **Box 0**, **Box 1** σ , we reuse the isomorphism notation to emphasize that the related primitives are inverses of each other.

There are several reduction rules for **copy** ($\text{share } e$), one for each type connective. These reductions perform a deep copy of the value, stopping only on ground data (\diamond), function values, and shared sub-terms: when copying a $!\sigma$ into a $!\sigma$, there is no need for a deep copy. When it encounters a location, **copy** ($\text{share } \ell$) reduces to a new allocation. If the location contains a value, the new location is filled with a copy of this value.

The copying rule for functions performs a copy of the local store s of the shared function. The locations in s are bound on the left-hand-side of the reduction, and free on the right-hand-side: this reduction step allocates fresh locations, and the store typing of the term changes from \cdot on the left to Ψ on the right. The fact that reduction changes the store typing is not unique to this rule, it is also the case when directly copying locations. In ML languages with references, the store only grows during reduction. That is not the case for our linear store: our reduction may either allocate new locations or free existing ones.

head reduction $\boxed{e \xrightarrow{L} e'}$ $\boxed{(s \mid e) \xrightarrow{L} (s' \mid e')}$

$$\begin{array}{c}
\text{let } \langle x_1, x_2 \rangle = \langle v_1, v_2 \rangle \text{ in } e \xrightarrow{L} e[v_1/x_1][v_2/x_2] \\
\langle \rangle; e \xrightarrow{L} e \\
(\lambda(x : \sigma). e) v \xrightarrow{L} e[v/x] \\
\text{case } (\text{inj}_i v) \text{ of } x_1. e_1 \mid x_2. e_2 \xrightarrow{L} e_i[v/x_i] \\
\text{unfold } (\text{fold}_{\mu\alpha.\sigma} v) \xrightarrow{L} v \\
\frac{e \xrightarrow{L} e'}{(s \mid e) \xrightarrow{L} (s \mid e')} \\
\begin{array}{ccc}
\text{new} & & \text{box} \\
\frac{(\emptyset \mid \langle \rangle) \xrightarrow{L} ([\ell \mapsto \cdot] \mid \ell)}{\text{free}} & (s[\ell \mapsto \cdot] \mid \langle \ell, v \rangle) \xrightarrow{L} ([\ell \mapsto (s \mid v)] \mid \ell) & \frac{}{\text{unbox}}
\end{array} \\
\text{copy } (\text{share}(s_1 \# s_2 : \Psi_1 \uplus \Psi_2). \langle v_1, v_2 \rangle) \\
\xrightarrow{L} \text{if } \text{locs}(s_i) = \text{locs}(\Psi_i) = \text{locs}(v_i) \\
\langle \text{copy } \text{share}(s_1 : \Psi_1). v_1, \text{copy } \text{share}(s_2 : \Psi_2). v_2 \rangle \\
\text{copy } (\text{share}(\emptyset : \cdot). \langle \rangle) \xrightarrow{L} \langle \rangle \\
\text{copy } (\text{share}(s : \Psi). \text{inj}_i v) \xrightarrow{L} \text{inj}_i \text{copy } (\text{share}(s : \Psi). v) \\
\text{copy } (\text{share}(s : \Psi). \text{fold } v) \xrightarrow{L} \text{fold } (\text{copy } (\text{share}(s : \Psi). v)) \\
(\emptyset \mid \text{copy } (\text{share}(s : \Psi). \lambda(x : \sigma). e)) \xrightarrow{L} (s \mid \lambda(x : \sigma). e) \\
\text{copy } (\text{share}(\emptyset : \cdot). (\text{share}(s : \Psi). v)) \xrightarrow{L} \text{share}(s : \Psi). v \\
\text{copy } (\text{share}([\ell \mapsto \cdot] : (\cdot; \cdot \vdash \ell : \text{Box } 0)). \ell) \xrightarrow{L} \text{new } \langle \rangle \\
\text{copy } (\text{share}([\ell \mapsto (s \mid v)] : (\Psi; !\Gamma \vdash \ell : \text{Box } 1 \sigma)). \ell) \\
\xrightarrow{L} \text{box } \langle \text{new } \langle \rangle, \text{copy } (\text{share}(s : \Psi). v) \rangle
\end{array}$$

linear reduction contexts $\boxed{\Psi; \Gamma \vdash_L s \mid K[\square : \sigma] : \sigma'}$

$$\begin{array}{l}
K ::= \square : \sigma \mid \langle K, e_2 \rangle \mid \langle v, K \rangle \mid \text{let } \langle v_1, v_2 \rangle = K \text{ in } e_2 \mid \\
K; e \mid K e \mid v K \mid \text{copy } K \mid \\
\text{inj}_1 K \mid \text{inj}_2 K \mid \text{case } K \text{ of } x_1. e_1 \mid x_2. e_2 \mid \\
\text{fold}_{\mu\alpha.\sigma} K \mid \text{unfold } K \mid \\
\text{new } K \mid \text{free } K \mid \text{box } K \mid \text{unbox } K
\end{array}$$

typing rules of terms, plus: $\frac{}{\cdot; \vdash_L \emptyset \mid (\square : \sigma) : \sigma}$

reduction $\boxed{(s \mid e) \xrightarrow{L} (s' \mid e')}$

$$\begin{array}{c}
\frac{(s \mid e) \xrightarrow{L} (s' \mid e') \quad \Psi; \Gamma \vdash_L s'' \mid K[\square : \sigma] : \sigma' \quad (s \mid e) \xrightarrow{L} (s' \mid e')}{(s \mid e) \xrightarrow{L} (s' \mid e') \quad (s'' \# s \mid K[e]) \xrightarrow{L} (s'' \# s' \mid K[e'])} \\
\frac{\Psi; \Gamma \vdash_L s \mid e : \sigma \quad (s \mid e) \xrightarrow{L} (s' \mid e') \quad \Psi'; \Gamma \vdash_L s' \mid e' : \sigma}{(\emptyset \mid \text{share}(s : \Psi). e) \xrightarrow{L} (\emptyset \mid \text{share}(s' : \Psi'). e')}
\end{array}$$

Fig. 6. Linear Language: Operational Semantics

We define a grammar of (deterministic) reduction contexts, which contain exactly one hole \square in evaluation position. However, we only define *linear* contexts K that do not share their hole: we need a specific treatment of the $\text{share}(s : \Psi).e$ reduction. Its subterm e is reduced in the local store s , but may create or free locations in the store; so we need to update the local store and its store typing during the reduction.

Theorem 1 (Progress) *If $\Psi; \Gamma \vdash_L s \mid e : \sigma$, then either e is a value v or there exists $(s' \mid e')$ such that $(s \mid e) \xrightarrow{L} (s' \mid e')$.*

Proof. By induction on the typing derivation of e , using induction hypothesis in the evaluation order corresponding to the structure of contexts K . If one induction hypothesis returns a reduction, we build a bigger reduction (\xrightarrow{L}) for the whole term. If all induction hypotheses return a value, the proof depends on whether the head term-former is an introduction/construction form or an elimination/destruction form. An introduction form whose subterms are values is a value. For elimination forms, we use [Lemma 3 \(Inversion principle for \$\lambda^L\$ values\)](#) on the eliminated subterm (a value), to learn that it starts with an introduction form, and thus forms a head redex with the head elimination form, so we build a head reduction (\rightsquigarrow^L).

Lemma 5 (Non-store-escaping substitution principle) *If*

$$\Psi'; \Gamma, x : \sigma \vdash_L s' \mid e : \sigma' \quad \Psi; \Gamma' \vdash_L s \mid v : \sigma \quad \Gamma \Downarrow \Gamma' \quad x \notin \Psi$$

then

$$\Psi \uplus \Psi'; \Gamma \Downarrow \Gamma' \vdash_L s \# s' \mid e[v/x] : \sigma'$$

Proof. The proof, summarized below, proceeds by induction on the typing derivation of e .

Most cases need an additional case analysis on whether the substituted type σ is a duplicable type of the form $!\sigma''$, as it influence whether it may appear in zero or several subterms of e . (This is a price to pay for contraction and weakening happening in all rules for convenience, instead of being isolated in separate structural rules.)

For example, in the variable case, e may be the variable x itself, in which case we know that Γ is empty and conclude immediately. But e may also be another variable y if x is duplicable and has been dropped. In that case, we perform an inversion ([Lemma 3](#)) on the v premise to learn that Ψ is empty and Γ' is duplicable, and can thus use [Lemma 4 \(Weakening by duplicable contexts\)](#).

In the $\langle e_1, e_2 \rangle$ case, if x is a linear variable it only occurs in one subterm on which we apply our induction hypothesis. If x is duplicable, inversion on the v premises again tells us that Γ' is duplicable. We know by assumption that $(\Gamma_1 \Downarrow \Gamma_2) \Downarrow \Gamma'$; because Γ' is duplicable, we can deduce from [Lemma 1 \(Context joining properties\)](#) that the $\Gamma_i \Downarrow \Gamma'$ are also defined, which let us

apply an induction hypothesis on both subterms e_i . To conclude, we need the computation

$$\begin{aligned} & (\Gamma_1 \Downarrow \Gamma') \Downarrow (\Gamma_2 \Downarrow \Gamma') \\ &= \Gamma_1 \Downarrow \Gamma_2 \Downarrow (\Gamma' \Downarrow \Gamma') \\ &= \Gamma_1 \Downarrow \Gamma_2 \Downarrow \Gamma' \end{aligned}$$

which again comes from duplicability of Γ' .

The assumption $x \notin \Psi$ enforces that the resource x is consumed in the term e itself, not in one of the values $[\ell \mapsto (s \mid v)]$ in the store: otherwise x would appear in the store typing $(\Gamma; \ell \vdash \Psi : \mathbf{Box} \ 1)$ of this location in Ψ . It is used in the case where $e : \sigma$ is a full location $\ell : \mathbf{Box} \ 1 \ \sigma'$. If x could appear in the value of ℓ in the store, we would have substitute it in the store as well – in our substitution statement, only the term is modified. Here we know that this value is unused, so it has a duplicable type and we can perform an inversion in the other cases.

Lemma 6 (Context decomposition) *If $\Psi'; \Gamma' \vdash_L s' \mid K[\square : \sigma] : \sigma'$ holds, then $\Psi''; \Gamma'' \vdash_L s'' \mid K[e] : \sigma'$ holds if and only if there exists Ψ, Γ, s such that $\Psi'' = \Psi \uplus \Psi', \Gamma'' = \Gamma \Downarrow \Gamma', s'' = s \# s'$ and $\Psi; \Gamma \vdash_L s \mid e : \sigma$.*

Theorem 2 (Subject reduction for λ^L) *If $\Psi; \Gamma \vdash_L s \mid e : \sigma$ and $(s \mid e) \xrightarrow{L} (s' \mid e')$, then there exists a (unique) Ψ' such that $\Psi'; \Gamma \vdash_L s' \mid e' : \sigma$.*

Proof. The proof is done by induction on the reduction derivation.

The head-reduction rules involving substitutions rely on [Lemma 5 \(Non-store-escaping substitution principle\)](#); note that in each of them, for example $(\lambda(x : \sigma'). e') e''$, the substituted variable x is bound in the term e , and thus does not appear in the store s : the non-store-escaping hypothesis holds.

For the copy rule and the store operators, we build a valid derivation for the reduced configuration by inverting the typing derivation of the reducible configuration.

In the non-head-reduction cases, the [share](#) case is by direction, and the context case $K[e]$ uses [Lemma 6 \(Context decomposition\)](#) to obtain a typing derivation for e , and the same lemma again rebuild a derivation of the reduced term $K[e']$.

3 Multi-language semantics

To formally define our multi-language semantics we create a combined language $\lambda^{\mathbf{UL}}$ which lets us compose term fragments from both $\lambda^{\mathbf{U}}$ and $\lambda^{\mathbf{L}}$ together, and we give an operational semantics to this combined language. Interoperability is enabled by specifying how to transport values across the language boundaries.

Multi-language systems in the wild are not defined in this way: both languages are given a semantics, by interpretation or compilation, in terms of a shared lower-level language (C, assembly, the JVM or CLR bytecode, or Racket's core forms), and the two languages are combined at that level. Our formal multi-language description can be seen as a model of such combinations, that gives a specification of the expected observable behavior of this language combination.

Another difference from multi-languages in the wild is our use of very fine-grained language boundaries: a term written in one language can have its subterms written in the other, provided the type-checking rules allow it. Most multi-language systems, typically using Foreign Function Interfaces, offer coarser-grained composition at the level of compilation units. Fine-grained composition of existing languages, as done in the Eco project [Barrett, Bolz, Diekmann, and Tratt, 2016], is difficult because of semantic mismatches. In Section 5 (Hybrid program examples) we demonstrate that fine-grained composition is a rewarding language design, enabling new programming patterns.

3.1 Lump Type and Language Boundaries

The core components the multi-language semantics are shown Figure 7—the communication of values from one language to the other will be described in the next section. The multi-language λ^{UL} has two distinct syntactic categories of types, values, and expressions: those that come from λ^{U} and those that come from λ^{L} . Contexts, on the other hand, are mixed, and can have variables of both sorts. For a mixed context Γ , the notation $!\Gamma$ only applies (!) to its linear variables.

The typing rules of λ^{U} and λ^{L} are imported into our multi-language system, working on those two separate categories of program. They need to be extended to handle mixed contexts Γ instead of their original contexts Γ and Γ . In the linear case, the rules look exactly the same. In the ML case, the typing rules implicitly duplicate all the variables in the context. It would be unsound to extend them to arbitrary linear variables, so they use a duplicable context $!\Gamma$.

To build interesting multi-language programs, we need a way to insert a fragment coming from a language into a term written in another. This is done using *language boundaries*, two new term formers $\mathcal{LU}(\mathbf{e})$ and $\mathcal{UL}(s:\Psi \mid \mathbf{e})$ that inject an ML term into the syntactic category of linear terms, and a linear configuration into the syntactic category of ML terms.

Of course, we need new typing rules for these term-level constructions, clarifying when it is valid to send a value from λ^{U} into λ^{L} and vice versa. It would be incorrect to allow sending any type from one language into the other—for instance, by adding the counterpart of our language boundaries in the syntax of types—since values of linear types must be uniquely owned so they cannot possibly be sent to the ML side as the ML type system cannot enforce unique ownership.

On the other hand, any ML value could safely be sent to the linear world. For closed types, we could provide a corresponding linear type ($\mathbf{1}$ maps to $!\mathbf{1}$, etc.), but an ML value may also be typed by an abstract type variable α , in which case we can't know what the linear counterpart should be. Instead of trying to provide translations, we will send any ML type σ to the *lump type* $[\sigma]$, which embeds ML types into linear types. A lump is a blackbox, not a type translation: the linear language does not assume anything about the behavior of its values—the values of $[\sigma]$ are of the form $[\mathbf{v}]$, where $\mathbf{v} : \sigma$ is an ML value that the linear world cannot use. More precisely, we only propagate the information that ML values are all duplicable by sending σ to $![\sigma]$.

Types $\sigma \mid \sigma$
 σ (unchanged from Figure 1)
 $\sigma + ::= \dots \mid [\sigma]$

Values $v \mid v$
 v (unchanged from Figure 1)
 $v + ::= \dots \mid [v]$

Expressions $e \mid e$
 $e + ::= \dots \mid \mathcal{UL}(s:\Psi \mid e)$
with $\mathcal{UL}(e) \stackrel{\text{def}}{=} \mathcal{UL}(\emptyset:\cdot \mid e)$
 $e + ::= \dots \mid \mathcal{LU}(e)$

Contexts $\Gamma ::= \cdot \mid \Gamma, x:\sigma \mid \Gamma, \alpha \mid \Gamma, x:\sigma$

Typing rules $\boxed{\Gamma \vdash_{\text{LU}} e : \sigma} \quad \boxed{\Psi \mid \Gamma \vdash_{\text{UL}} s \mid e : \sigma}$

with $\Gamma \vdash_{\text{UL}} e : \sigma \stackrel{\text{def}}{=} \cdot \mid \Gamma \vdash_{\text{UL}} \emptyset \mid e : \sigma$

(Typing rules of $\Gamma \vdash_{\text{U}} e : \sigma$ reused, with mixed context $!\Gamma$)

(Typing rules of $\Psi; \Gamma \vdash_{\text{L}} s \mid e : \sigma$ reused, with mixed context Γ)

$$\frac{!\Gamma \vdash_{\text{LU}} e : \sigma}{\cdot \mid !\Gamma \vdash_{\text{UL}} \emptyset \mid \mathcal{LU}(e) : ![\sigma]} \quad \frac{\Psi \mid !\Gamma \vdash_{\text{UL}} s \mid e : ![\sigma]}{!\Gamma \vdash_{\text{LU}} \mathcal{UL}(s:\Psi \mid e) : \sigma}$$

Reduction rules

(Reduction rules of λ^{U} and λ^{L} reused unchanged)

$$\frac{e \xrightarrow{\text{U}} e'}{\mathcal{LU}(e) \xrightarrow{\text{L}} \mathcal{LU}(e')}$$

$$\frac{\overline{\mathcal{LU}(v) \xrightarrow{\text{L}} [v]} \quad \overline{\mathcal{UL}(\emptyset:\cdot \mid \text{share}[v]) \xrightarrow{\text{U}} v}}{\Psi \mid \Gamma \vdash_{\text{UL}} s \mid e : \sigma \quad (s \mid e) \xrightarrow{\text{L}} (s' \mid e') \quad \Psi' \mid \Gamma \vdash_{\text{UL}} s' \mid e' : \sigma}{\mathcal{UL}(s:\Psi \mid e) \xrightarrow{\text{U}} \mathcal{UL}(s':\Psi' \mid e')}$$

Fig. 7. Multi-language: Lump and Boundaries

The typing rules for language boundaries insert lumps when going from λ^U to λ^L , and remove them when going back from λ^L to λ^U . In particular, arbitrary linear types cannot occur at the boundary, they must be of the form $![\sigma]$.

Finally, boundaries have reduction rules: a term or configuration inside a boundary in reduction position is reduced until it becomes a value, and then a lump is added or removed depending on the boundary direction. Note that because the \mathbf{v} in $\mathcal{UL}(s:\Psi \mid \mathbf{v})$ is at a duplicable type $![\sigma]$, we know by inversion that the store is empty.

3.2 Interoperability: Static Semantics

If the linear language could not interact with lumped values at all, our multi-language programs would be rather boring, as the only way for the linear extension to provide a value back to ML would be to have received it from λ^U and pass it back unchanged (as in the lump embedding of [Matthews and Findler \[2009\]](#)). To provide a real interaction, we provide a way to extract values out of a lump $![\sigma]$, use it at some linear type σ , and put it back in before sending the result to λ^U .

Interoperability context $\Sigma ::= \cdot \mid \Sigma, \alpha \simeq \beta$

Compatibility relation $\boxed{\Sigma \vdash_{\text{UL}} \sigma \simeq \sigma}$ with $\sigma \simeq \sigma \stackrel{\text{def}}{=} \cdot \vdash_{\text{UL}} \sigma \simeq \sigma$

$$\begin{array}{c}
\frac{}{\Sigma \vdash_{\text{UL}} 1 \simeq !1} \qquad \frac{\Sigma \vdash_{\text{UL}} \sigma_1 \simeq !\sigma_1 \quad \Sigma \vdash_{\text{UL}} \sigma_2 \simeq !\sigma_2}{\Sigma \vdash_{\text{UL}} \sigma_1 \times \sigma_2 \simeq !(\sigma_1 \otimes \sigma_2)} \\
\frac{\Sigma \vdash_{\text{UL}} \sigma_1 \simeq !\sigma_1 \quad \Sigma \vdash_{\text{UL}} \sigma_2 \simeq !\sigma_2}{\Sigma \vdash_{\text{UL}} \sigma_1 + \sigma_2 \simeq !(\sigma_1 \oplus \sigma_2)} \qquad \frac{\Sigma \vdash_{\text{UL}} \sigma \simeq !\sigma \quad \Sigma \vdash_{\text{UL}} \sigma' \simeq !\sigma'}{\Sigma \vdash_{\text{UL}} \sigma \rightarrow \sigma' \simeq !(\sigma \multimap \sigma')} \\
\frac{}{\Sigma \vdash_{\text{UL}} \sigma \simeq ![\sigma]} \qquad \frac{\Sigma \vdash_{\text{UL}} \sigma \simeq !\sigma}{\Sigma \vdash_{\text{UL}} \sigma \simeq !!\sigma} \qquad \frac{\Sigma \vdash_{\text{UL}} \sigma \simeq !\sigma}{\Sigma \vdash_{\text{UL}} \sigma \simeq !(\text{Box } 1 \sigma)} \\
\frac{\Sigma, \alpha \simeq !\beta \vdash_{\text{UL}} \sigma \simeq !\sigma}{\Sigma \vdash_{\text{UL}} \mu\alpha. \sigma \simeq !(\mu\beta. \sigma)} \qquad \frac{(\alpha \simeq !\beta) \in \Sigma}{\Sigma \vdash_{\text{UL}} \alpha \simeq !\beta}
\end{array}$$

Interoperability primitives and derived constructs:

$$![\sigma] \begin{array}{c} \xrightarrow{\sigma \text{ unlump}} \\ \xrightarrow{\text{lump}^\sigma} \end{array} \sigma \quad \text{whenever } \cdot \vdash_{\text{UL}} \sigma \simeq \sigma \qquad \sigma \mathcal{LU}(e) \stackrel{\text{def}}{=} \sigma \text{ unlump } \mathcal{LU}(e)$$

$$\mathcal{UL}^\sigma(e) \stackrel{\text{def}}{=} \mathcal{UL}(\text{lump}^\sigma e)$$

Fig. 8. Multi-language: Static Interoperability Semantics

The correspondence between intuitionistic types σ and linear types σ is specified by a heterogeneous *compatibility relation* $\sigma \simeq \sigma$ defined in Figure 8 (Multi-language: Static Interoperability Semantics). The specification of this relation is that if $\sigma \simeq \sigma$ holds, then the space of values of $![\sigma]$ and σ are isomorphic: we can convert back and forth between them. When this relation holds, the term-formers lump^σ and ${}^\sigma\text{unlump}$ perform the conversion. (The position of the index σ emphasizes that the *input* e of $\text{lump}^\sigma e$ has type σ , while the *output* of ${}^\sigma\text{unlump } e$ has type σ .)

For example, we have $![(\sigma \rightarrow \sigma')] \simeq !(![\sigma] \multimap ![\sigma'])$. Given a lumped ML function, we can unlump it to see it as a linear function. We can call it from the linear side, but have to pass it a duplicable argument since an ML function may duplicate its argument. Conversely, we can convert a linear function into a lumped function type to pass it to the ML side, but it has to have a duplicable return type since the ML side may freely share the return value.

Our lump^σ and ${}^\sigma\text{unlump}$ primitives are only indexed by the linear type σ , because a compatible ML type σ can be uniquely recovered, as per the following result.

Lemma 7 (Determinism of the compatibility relation) *If $\sigma \simeq \sigma$ and $\sigma' \simeq \sigma$ then $\sigma = \sigma'$.*

Proof. By induction on the syntax of σ : the judgment $\Sigma \vdash_{\text{UL}} \sigma \simeq \sigma$ is syntax-directed in its σ component.

Note that the converse property does not hold: for a given σ , there are many σ such that $\sigma \simeq \sigma$. For example, we have $\mathbf{1} \simeq !\mathbf{1}$ but also $\mathbf{1} \simeq !!\mathbf{1}$. This corresponds to the fact that the linear types are more fine-grained, and make distinctions (inner duplicability, dereference of full locations) that are erased in the ML world. The $\sigma \simeq ![\sigma]$ case also allows you to (un)lump as deeply or as shallowly as you need: $\sigma_1 \times (\sigma_2 \times \sigma_3)$ is compatible with both $!(![\sigma_1] \otimes ![\sigma_2 \times \sigma_3])$ and $!(![\sigma_1] \otimes (![\sigma_2] \otimes ![\sigma_3]))$. We could not systematically translate the complete type σ , as type variables cannot be translated and need to remain lumped. Allowing lumps to “stop” the translation at arbitrary depth is a natural generalization.

Lemma 8 (Substitution of recursive hypotheses) *If $\Sigma, \alpha \simeq !\beta \vdash_{\text{UL}} \sigma \simeq \sigma$, $\alpha \notin \sigma$, and $\Sigma \vdash_{\text{UL}} \sigma' \simeq !\sigma'$ then $\Sigma \vdash_{\text{UL}} \sigma[\sigma'/\alpha] \simeq \sigma[\sigma'/\beta]$.*

Proof. By induction on the $\sigma \simeq \sigma$ derivation. There are two leaf cases: the case recursive hypotheses, which is immediate, and the case of lump $\sigma \simeq ![\sigma]$. In this latter case, notice that σ is a type of λ^{U} , so in particular it does not contain the variable β ; and we assumed $\alpha \notin \sigma$ so we also have $\alpha \notin \sigma$, so $\sigma[\sigma'/\alpha] = \sigma \simeq ![\sigma] = \sigma[\sigma'/\beta]$.

The term $\mathcal{LU}(e)$ turns a $e : \sigma$ into a lumped type $![\sigma]$, and we need to unlump it with some ${}^\sigma\text{unlump}$ for a compatible $\sigma \simeq \sigma$ to interact with it on the linear side. It is common to combine both operations and we provide syntactic sugar for it: ${}^\sigma\mathcal{LU}(e)$. Similarly $\mathcal{UL}^\sigma(e)$ first lumps a linear term then sends the result to the ML world.

3.3 Interoperability: Dynamic Semantics

$$\begin{array}{c}
\langle \rangle \leftrightarrow^{!1} \text{share } \langle \rangle \quad \frac{\begin{array}{l} v_1 \leftrightarrow^{! \sigma_1} \text{share}(s_1 : \Psi_1). v_1 \quad \text{locs}(s_1) = \text{locs}(\Psi_1) = \text{locs}(v_1) \\ v_2 \leftrightarrow^{! \sigma_2} \text{share}(s_2 : \Psi_2). v_2 \quad \text{locs}(s_2) = \text{locs}(\Psi_2) = \text{locs}(v_2) \end{array}}{\langle v_1, v_2 \rangle \leftrightarrow^{!(\sigma_1 \otimes \sigma_2)} \text{share}(s_1 \# s_2 : \Psi_1 \uplus \Psi_2). \langle v_1, v_2 \rangle} \\
\\
\frac{v \leftrightarrow^{! \sigma_i} \text{share}(s : \Psi). v}{\text{inj}_i v \leftrightarrow^{!(\sigma_1 \oplus \sigma_2)} \text{share}(s : \Psi). \text{inj}_i v} \quad \frac{\sigma \simeq \sigma' \quad \sigma' \simeq \sigma'}{e \rightarrow^{!(\sigma \multimap \sigma')} \text{share } \lambda(x : !\sigma). \sigma' \mathcal{L}\mathcal{U}(e \mathcal{U}\mathcal{L}^\sigma(x))} \\
\\
\frac{\sigma \simeq \sigma' \quad \sigma' \simeq \sigma'}{\lambda(x : \sigma). \mathcal{U}\mathcal{L}^{\sigma'}(\text{copy } e^\sigma \mathcal{L}\mathcal{U}(x)) \leftarrow^{!(\sigma \multimap \sigma')} e} \quad \frac{}{v \leftrightarrow^{![\sigma]} \text{share } [v]} \\
\\
\frac{v \leftrightarrow^{! \sigma} \text{share } v}{v \leftrightarrow^{! \sigma} \text{share } (\text{share } v)} \quad \frac{v \leftrightarrow^{! \sigma} \text{share}(s : \Psi). v}{v \leftrightarrow^{! \text{Box } 1 \sigma} \text{share}([\ell \mapsto (s \mid v)] : (\cdot; \ell \vdash \Psi : \text{Box } 1 \sigma)). \ell} \\
\\
\frac{v \leftrightarrow^{! \sigma[\mu\alpha.\sigma/\alpha]} \text{share}(s : \Psi). v}{\text{fold}_{\mu\alpha.\sigma} v \leftrightarrow^{! \mu\alpha.\sigma} \text{share}(s : \Psi). (\text{fold}_{\mu\alpha.\sigma} v)} \\
\\
\begin{array}{ccc}
& \sigma \text{ unlump} & \text{whenever } v \rightarrow^\sigma v \\
& \downarrow \rightsquigarrow & \\
(\emptyset \mid \text{share } [v]) & & (\emptyset \mid v) \\
& \uparrow \rightsquigarrow & \\
& \text{lump}^\sigma & \text{whenever } v \leftarrow^\sigma v
\end{array}
\end{array}$$

Fig. 9. Multi-language: Dynamic Interoperability Semantics

We were careful to define the compatibility relation such that $\sigma \simeq \sigma'$ only holds when $![\sigma]$ and σ' are isomorphic, in the sense that any value of one can be converted into a value of another. Figure 9 defines the operational semantics of the lumping and unlumping operations precisely as realizing these isomorphisms. For concision, we specify the isomorphisms as relations, following the inductive structure of the compatibility judgment itself. We write (\leftrightarrow) when a rule can be read bidirectionally to convert in either directions (assuming the same direction holds of the premises), and (\leftarrow) or (\rightarrow) for rules that only describe how to convert values in one direction.

Lemma 9 (Substitution of polymorphic variables) *If $\Sigma \vdash_{\text{UL}} \sigma \simeq \sigma'$ and $\alpha \notin \Sigma$, then $\Sigma \vdash_{\text{UL}} \sigma[\sigma'/\alpha] \simeq \sigma'[\sigma'/\alpha]$ and their lumping operations coincide on all values.*

Proof. By induction on $\sigma \simeq \sigma'$. In the variable leaf case, we know $\alpha \notin \Sigma$. In the lump leaf case $\sigma \simeq ![\sigma]$, the goal $\sigma[\sigma'/\alpha] \simeq ![\sigma[\sigma'/\alpha]]$ is immediate.

Theorem 3 (Value translations are functional) *If $\sigma \simeq \sigma'$, then for any closed value $v : \sigma$ there is a unique $v' : \sigma'$ such that $v \rightarrow^\sigma v'$, and conversely for any closed value $v' : \sigma'$ there is a unique $v : \sigma$ such that $v \leftarrow^\sigma v'$.*

Proof. Remark that in the statement of the term, when we quantify over all closed values $\mathbf{v} : \sigma$, we implicitly assume that in the general case values of \mathbf{v} live in the empty global store – otherwise we would have a value of the form $\Psi; \cdot \vdash_L s \mid \mathbf{v} : \sigma$. This is valid because all types σ in the image of the type-compatibility relation are duplicable types of the form $!\sigma'$, so by Lemma 3 (Inversion principle for λ^L values) we know that \mathbf{v} is in fact of the form $\text{share}(s : \Psi).e'$, living in the empty store.)

The two sides of the result are proved simultaneously by induction on $\sigma \simeq \sigma$, using inversion to reason on the shapes of \mathbf{v} and \mathbf{v} . Note that the inductive cases remain on closed values: the only variable-binder constructions, λ -abstractions, do not use the recursion hypothesis.

In the recursive case $\mu\alpha. \sigma \simeq !(\mu\alpha. \sigma)$, to use the induction hypothesis on the folded values we need to know that the unfolded types $\sigma[\mu\alpha. \sigma/\alpha]$ and $\sigma[\mu\alpha. \sigma/\alpha]$ are compatible. This is exactly Lemma 8 (Substitution of recursive hypotheses), using the hypothesis $\mu\alpha. \sigma \simeq !(\mu\alpha. \sigma)$ itself.

Lemma 10 (Lumping cancellation) *The lump conversions lump^σ and σunlump cancel each other modulo $\beta\eta$. In particular,*

$${}^\sigma \mathcal{LU}(\mathcal{UL}^\sigma(\mathbf{v})) =_{\beta\eta} \mathbf{v} \qquad \mathcal{UL}^\sigma({}^\sigma \mathcal{LU}(\mathbf{v})) =_{\beta\eta} \mathbf{v}$$

Proof. By induction on σ , and then by parallel induction on the derivations of type compatibility and value compatibility. The parallel cases are symmetric by definition, only the function case $!(\sigma \rightarrow !\sigma')$ needs to be checked. A simple computation, using the induction hypothesis on the smaller types $!\sigma$ and $!\sigma'$, shows that composing the two function translations gives an η -expansion – plus the $\beta\eta$ -steps from the induction hypotheses.

Note that the β - η rules used (for functions and $!_*$) have been proved admissible for the logical relation, so in particular sound with respect with contextual equivalence.

Implementation consideration In a realistic implementation of this multi-language system, we would expect the representation choices made for λ^U and λ^L to be such that, for some but not all compatible pairs $\sigma \simeq \sigma$, the types σ and σ actually have the exact same representation, making the conversion an efficient no-op. An implementation could even restrict the compatibility relation to accept only the pairs that can be implemented in this way. That is, it would reject some λ^{UL} programs, but the “graceful interoperability” result that is our essential contribution would still hold.

3.4 Full Abstraction from λ^U into λ^{UL}

We can now state and prove the major meta-theoretical result of this work, which is the proposed multi-language design extends the simple language λ^U in a way that provably has, in a certain sense, “no abstraction leaks”.

Definition 1 (Contextual equivalence in λ^U) We say that e, e' such that $\Gamma \vdash_U e, e' : \sigma$ are contextually equivalent, written $e \approx_U^{ctx} e'$, if, for any expression context $C[\square]$ such that $\cdot \vdash_U C[e] : \mathbf{1}$, the closed terms $C[e]$ and $C[e']$ are equi-terminating.

Definition 2 (Contextual equivalence in λ^{LU}) We say that e, e' such that $\Gamma \vdash_{LU} e, e' : \sigma$ are contextually equivalent, written $e \approx_{LU}^{ctx} e'$, if, for any expression context $C[\square]$ such that $\cdot \vdash_{LU} C[e] : \mathbf{1}$, the closed terms $C[e]$ and $C[e']$ are equi-terminating.

4 Multi-language parametricity

We discussed the design choice of manipulating lumps $[\sigma]$ of any ML type, not just the type variable that motivates them. In the presence of polymorphism, this generalization is also an important design choice to preserve parametricity.

Let us define $\text{id}^\sigma(e) \stackrel{\text{def}}{=} \text{lump}^\sigma(\sigma \text{unlump } e)$, and consider a polymorphic term of the form $\lambda\alpha. \mathcal{UL}(\dots \text{id}^{![\alpha]} \dots)$. The (un)lumping operations on a lumped type such as $![\alpha]$ are just the identity: the lumped value is passed around unchanged, so $\text{id}^{![\alpha]}(v)$ will reduce to v . Now, if we instantiate this polymorphic term with a ML type σ , it will reduce to a term $\mathcal{UL}(\dots \text{id}^{![\sigma]} \dots)$ whose unlumping operation is still on a lumped type, so is still exactly the identity.

On the contrary, if we allowed lumps only on type variables, we would have to push the lump inside σ , and the (un)lumping operations would become more complex: if σ starts with an ML product type $_ \times _$, it would be turned into a shared linear pair $!(_ \otimes _)$ by unlumping, and back into an ML pair by lumping. In general, id^σ may perform deep η -expansions of lumped values. The fact that, after instantiation of the polymorphic term, we get a monomorphic term that has different (but η -equivalent) computational behavior would cause meta-theoretic difficulties; this is the approach that was adopted in previous work on multi-languages with polymorphism by Perconti and Ahmed [2014], and it made some of their proofs using a logical relations argument substantially more complicated. In the logical relation, polymorphism is obtained by allowing each polymorphic variable to be replaced by two types related by an “admissible” relation R , and the notion of admissibility of this previous work had to force relations to be compatible with η -expansion, which complicates the proofs.

In contrast, our handling of lump types as turning arbitrary types into black-boxes makes type instantiation obviously parametric. To formally demonstrate this aspect of our design, we develop a step-indexed logical relation (Figure 10) that proves that our multi-language satisfies a strong parametricity property that is not disrupted by the linear sublanguage or the cross-language boundaries.

The logical relation is a family of relations indexed by closed “relational types” which extend the grammar of σ, σ , which include a case for admissible relations R that is used to enable parametric arguments. The step index j in the definitions decreases strictly whenever related values of a type σ or σ are defined in terms of a non-strictly-smaller type; this happens in the definition of the relation at

recursive types $\mathcal{V} \llbracket \mu\alpha. \sigma \rrbracket^j$. Because the language is non-terminating, our relation does not define an equivalence but an approximation: two expressions (e_1, e_1) are related in $\mathcal{E} \llbracket \sigma \rrbracket^j$ if e_1 approximates e_2 : if e_1 reduces to a value in less than j steps, then e_2 must reduce to a related value.

The definition of admissible relations $\text{Rel}[\sigma_1, \sigma_2]$, used to define when λ^U values of polymorphic types are related, $\mathcal{V} \llbracket \forall\alpha. \sigma \rrbracket$, is completely standard, which demonstrates that our notion of boundaries preserves simple parametricity reasoning.

Although we have a stateful linear language, the logical relations for the linear types have more in common with a language with explicit closures—this is another consequence of the remark in Section 2.2 (Linear Memory in λ^L) that the language can also be interpreted using a functional semantics. The relations for closed λ^L values and expressions, $\mathcal{V} \llbracket \sigma \rrbracket$ and $\mathcal{E} \llbracket \sigma \rrbracket$, are indexed by a type but do not depend on a store typing: the related values may have different, non-empty store typings. This allows to relate two programs that are equivalent but allocate different references in different ways. Furthermore, since all state is *linear*, we don't need additional machinery to relate stores, since all the values in a store owned by a value will be reflected in the value. For example, the relation for empty locations $\mathcal{V} \llbracket \text{Box } 0 \rrbracket$ relates any two arbitrary (empty) locations, and the relation for non-empty locations $\mathcal{V} \llbracket \text{Box } 1 \ \sigma \rrbracket$ relates possibly-distinct locations that contain related values.

Logical relations effectively translate global invariants of the system into properties of type connectives. For example, consider the reduction rule for lumped values: $\mathcal{UL}(\emptyset : \cdot \mid \text{share } [v]) \xrightarrow{U} v$. In this rule we implicitly assumed that a linear value of the shape $[v]$ at type $!\sigma$ would occur in an empty local store. The term $\text{share } [v]$ desugars into $\text{share}(\emptyset : \cdot). [v]$, but it is not immediately obvious that this should always be the case since it is possible to compute a value of type $!\sigma$ by allocating references and using them. The intuitive reason why the store becomes empty when a value $[v]$ is reached is that linear sub-terms e within v may only occur within a language boundary $\mathcal{UL}(s : \Psi \mid e)$: linear sub-terms have their own local store, so there are no globally visible linear locations for $[v]$ to refer to. This global reasoning is elegantly expressed in a type-directed way in our logical relation by the definition of related values at lump type, which encodes the invariant that they always have an empty store:

$$\mathcal{V} \llbracket \sigma \rrbracket^j \stackrel{\text{def}}{=} \{((\emptyset \mid [v_1]), (\emptyset \mid [v_2])) \mid (v_1, v_2) \in \mathcal{V} \llbracket \sigma \rrbracket^j\}$$

From the logical relation defined on closed λ^{UL} terms and values, we define logical approximation relations on open terms $\Gamma \vdash e_1 \lesssim^{log} e_2 : \sigma$ and $\Gamma \vdash_L (s_1 \mid e_1) \lesssim^{log} (s_2 \mid e_2) : \sigma$ by asking for the open terms and values to be related under all related environments in the standard way—see Appendix A for full details. This lets us demonstrate the Fundamental Property to validate the construction of our logical relation, showing that all typing rules are admissible.

Theorem 4 (Fundamental Property)

1. If $!\Gamma \vdash_v e : \sigma$ then $!\Gamma \vdash e \lesssim^{log} e : \sigma$

$$\begin{aligned}
\text{Atom}[\sigma] &\stackrel{\text{def}}{=} \{v \mid \cdot \vdash_v v : \sigma\} \\
\text{Rel}[\sigma_1, \sigma_2] &\stackrel{\text{def}}{=} \{R : \mathbb{N} \rightarrow \mathcal{P}(\text{Atom}[\sigma_1] \times \text{Atom}[\sigma_2]) \mid \forall j \leq j'. R^{j'} \subset R^j\} \\
\mathcal{V}[[R]]^j &\stackrel{\text{def}}{=} R^j \\
\mathcal{V}[[\sigma_1 \rightarrow \sigma_2]]^j &\stackrel{\text{def}}{=} \{(\lambda(x_1 : (\sigma_1)_1). e_1, \lambda(x_2 : (\sigma_1)_2). e_2) \mid \forall j' \leq j, (v_1, v_2) \in \mathcal{V}[[\sigma_1]]^{j'} \cdot (e_1[v_1/x_1], e_2[v_2/x_2]) \in \mathcal{E}[[\sigma_2]]^{j'}\} \\
\mathcal{V}[[\forall \alpha. \sigma]]^j &\stackrel{\text{def}}{=} \{(\Lambda \alpha. v_1, \Lambda \alpha. v_2) \mid \forall \sigma_1, \sigma_2, R \in \text{Rel}[\sigma_1, \sigma_2]. (v_1, v_2) \in \mathcal{V}[[\sigma[R/\alpha]]]^j\} \\
\mathcal{V}[[\mathbf{1}]]^j &\stackrel{\text{def}}{=} \{((\emptyset \mid \langle \rangle), (\emptyset \mid \langle \rangle))\} \\
\mathcal{V}[[\sigma \otimes \sigma']]^j &\stackrel{\text{def}}{=} \{((s_1 \# s'_1 \mid \langle v_1, v'_1 \rangle), (s_2 \# s'_2 \mid \langle v_2, v'_2 \rangle)) \mid ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[[\sigma]]^j \wedge ((s'_1 \mid v'_1), (s'_2 \mid v'_2)) \in \mathcal{V}[[\sigma']]^j\} \\
\mathcal{V}[[\sigma' \multimap \sigma]]^j &\stackrel{\text{def}}{=} \{((s_1 \mid \lambda(x : \sigma'). e_1), (s_2 \mid \lambda(x : \sigma'). e_2)) \mid \\
&\quad \forall j' \leq j, s'_1, s'_2, ((s'_1 \mid v_1), (s'_2 \mid v_2)) \in \mathcal{V}[[\sigma']]^{j'} \cdot \\
&\quad s'_1 = s_1 \# s''_1 \wedge s'_2 = s_2 \# s''_2 \Rightarrow ((s'_1 \mid e_1[v_1/x]), (s'_2 \mid e_2[v_2/x])) \in \mathcal{E}[[\sigma]]^{j'}\} \\
\mathcal{V}[[\mu \alpha. \sigma]]^j &\stackrel{\text{def}}{=} \{((s_1 \mid \text{fold}_{\mu \alpha. \sigma} v_1), (s_2 \mid \text{fold}_{\mu \alpha. \sigma} v_2)) \mid \forall j' < j. ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[[\sigma[\mu \alpha. \sigma/\alpha]]]^{j'}\} \\
\mathcal{V}[[! \sigma]]^j &\stackrel{\text{def}}{=} \{((\emptyset \mid \text{share}(s_1 : \Psi_1). v_1), (\emptyset \mid \text{share}(s_2 : \Psi_2). v_2)) \mid ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[[\sigma]]^j\} \\
\mathcal{V}[[\text{Box } 0]]^j &\stackrel{\text{def}}{=} \{(([\ell_1 \mapsto \cdot] \mid \ell_1), ([\ell_2 \mapsto \cdot] \mid \ell_2))\} \\
\mathcal{V}[[\text{Box } 1 \sigma]]^j &\stackrel{\text{def}}{=} \{(([\ell_1 \mapsto (s_1 \mid v_1)] \mid \ell_1), ([\ell_2 \mapsto (s_2 \mid v_2)] \mid \ell_2)) \mid ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[[\sigma]]^j\} \\
\mathcal{V}[[[\sigma]]]^j &\stackrel{\text{def}}{=} \{((\emptyset \mid [v_1]), (\emptyset \mid [v_2])) \mid (v_1, v_2) \in \mathcal{V}[[\sigma]]^j\} \\
\mathcal{E}[[\sigma]]^j &\stackrel{\text{def}}{=} \{(e_1, e_2) \mid \forall j' \leq j. e_1 \xrightarrow{u^{j'}} v_1 \Rightarrow \exists v_2. e_2 \xrightarrow{u^*} v_2 \wedge (v_1, v_2) \in \mathcal{V}[[\sigma]]^{j-j'}\} \\
\mathcal{E}[[\sigma]]^j &\stackrel{\text{def}}{=} \{((s_1 \mid e_1), (s_2 \mid e_2)) \mid \forall j' \leq j, (s'_1 \mid v_1). (s_1 \mid e_1) \xrightarrow{L^{j'}} (s'_1 \mid v_1) \Rightarrow \\
&\quad \exists (s'_2 \mid v_2). (s_2 \mid e_2) \xrightarrow{L^*} (s'_2 \mid v_2) \wedge ((s'_1 \mid v_1), (s'_2 \mid v_2)) \in \mathcal{V}[[\sigma]]^{j-j'}\}
\end{aligned}$$

Fig. 10. Multi-language Logical Relation (excerpt)

2. If $\Psi; \Gamma \vdash_L s \mid e : \sigma$ then $\Gamma \vdash_L (s \mid e) \lesssim^{log} (s \mid e) : \sigma$

We also prove that the logical relation is sound with respect to contextual equivalence. For this, we define contextual approximation relations $\Gamma \vdash e_1 \lesssim^{ctx} e_2 : \sigma$ and $\Gamma \vdash_L (s_1 \mid e_1) \lesssim^{ctx} (s_2 \mid e_2) : \sigma$, by asking that $e_1, (s_1 \mid e_1)$ terminate more often than $e_2, (s_2 \mid e_2)$ when run under arbitrary contexts—see Appendix A for details.

Theorem 5 (Soundness of Logical Relation with respect to Contextual Equivalence)
 $(\lesssim^{log}) \subset (\lesssim^{ctx})$.

5 Hybrid program examples

5.1 In-Place Transformations

In Section 2.2 (Linear Memory in λ^L) we proposed a program for in-place reversal of linear lists defined by the type $\text{LinList } \sigma \stackrel{\text{def}}{=} \mu\alpha. \mathbf{1} \oplus \text{Box } \mathbf{1} (\sigma \otimes \alpha)$. We can also define a type of ML lists $\text{List } \sigma \stackrel{\text{def}}{=} \mu\alpha. \mathbf{1} + \sigma \times \alpha$. Note that ML lists are compatible with shared linear lists, in the sense that $\text{List } \sigma \simeq !(\text{LinList } ![\sigma])$. This enables writing in-place list-manipulation functions in λ^L , and exposing them to beginners at a λ^U type:

$$\text{rev xs} \stackrel{\text{def}}{=} \mathcal{UL}^{\text{LinList } ![\sigma]}(\text{share}(\text{rev_into copy } (\text{LinList } ![\sigma] \mathcal{LU}(\text{xs})) \text{ Nil}))$$

This example is arguably silly, as the allocations that are avoided by doing an in-place traversal are paid when copying the shared list to obtain a uniquely-owned version. A better example of list operations that can profitably be sent on the linear side is quicksort, whose code we give in Figure 11 (Quicksort). An ML implementation allocates intermediary lists for each recursive call, while the surprisingly readable λ^U implementation only allocates for the first copy.

5.2 Typestate Protocols

Linear types can enforce proper allocation and deallocation of resources, and in general any automata/typestate-like protocols on their usage by encoding the state transitions as linear transformations. In the simple example of file-descriptor handling in the introduction, additional safety compared to ML programming can be obtained by exposing file-handling functions on the λ^U side, with linear types. We assumed the following API for linear file handling, which enforces a correct usage protocol:

```
open : !([Path]  $\multimap$  Handle)
line  : !(Handle  $\multimap$  (Handle  $\oplus$  (![String]  $\otimes$  Handle)))
close : !(Handle  $\multimap$  1)
```

Another interesting example of protocol usage for which linear types help is the use of *transient* versions of persistent data structures, as popularized by Clojure. An unrestricted type $\text{Set } \alpha$ may represent persistent sets as balanced trees

```

partition : !( $\alpha \multimap \text{Bool}$ )  $\multimap$  LList ! $\alpha \multimap$  LList ! $\alpha \oplus$  LList ! $\alpha$ 
partition p li = partition_aux p (Nil, Nil) li
partition_aux p (yes, no) Nil = (yes, no)
partition_aux p (yes, no) (Cons l x xs) =
  let (yes, no) =
    if copy p x
    then (Cons l x yes, no)
    else (yes, Cons l x no) in
  partition_aux p (yes, no) xs

lin_quicksort : LList ! $\alpha \multimap$  LList ! $\alpha$ 
lin_quicksort li = quicksort_aux li Nil
quicksort_aux Nil acc = acc
quicksort_aux (Cons l head li) acc =
  let p = share (fun x -> x < head) in
  let (below, above) = partition p li in
  quicksort_aux below (Cons l head (quicksort_aux above acc))

quicksort li UL(li) = UL(share (lin_quicksort (copy li)))

```

Fig. 11. Quicksort

with logarithmic operations performing path-copying. A **transient** call returns a mutable version of the structure that supports efficient batch in-place updates, before a **persistent** call freezes this transient structure back into a persistent tree. To preserve a purely functional semantics, we must enforce that the intermediate transient value is uniquely owned. We can do this by using the linear types for the transient API:

```

type Set  $\alpha$ 
val add : Set  $\alpha \rightarrow \alpha \rightarrow$  Set  $\alpha$  (* path copy *)
...
type MutSet  $\alpha$ 
val add : !(MutSet  $\alpha \multimap \alpha \multimap$  MutSet  $\alpha$ ) (* in-place update *)
...
val transient : !( $[\text{Set } \alpha] \multimap$  MutSet ! $[\alpha]$ )
val persistent : !(MutSet ! $[\alpha] \multimap$  ! $[\text{Set } \alpha]$ )

```

6 Conclusion

In our proposed multi-language design, a simple linear type system mirroring the standard rules of intuitionistic linear logic can be equipped with linear state and usefully complement a general-purpose functional ML language, without breaking equational reasoning or parametricity—and without requiring a significantly more complex meta-theory.

Fine-grained language boundaries allow interesting programming patterns to emerge, and full abstraction provides a novel rigorous specification of what it means for multi-language design to avoid *abstraction leaks* from advanced features into the general-purpose or beginner-friendly languages.

6.1 Related Work

Having a stack of usable, interoperable languages, extensions or dialects is at the forefront of the Racket approach to programming environments, in particular for teaching [Felleisen, Findler, Flatt, and Krishnamurthi, 2004].

Our multi-language semantics builds on the seminal work by Matthews and Findler [2009], who gave a formal semantics of interoperability between a dynamically and a statically typed language. Others have followed the Matthews-Findler approach of designing multi-language systems with fine-grained boundaries—for instance, formalizing interoperability between a simply and dependently typed language [Osera, Sjöberg, and Zdancewic, 2012]; between a functional and typed assembly language [Patterson, Perconti, Dimoulas, and Ahmed, 2017]; between an ML-like and an affinely typed language, where linearity is enforced at runtime on the ML side using stateful contracts [Tov and Pucella, 2010]; and between the source and target languages of compilation to specify compiler correctness [Perconti and Ahmed, 2014]. However, all these papers address only the question of soundness of the multi-language; we propose a formal treatment of *usability* and absence of abstraction leaks.

The only work to establish that a language embeds into a multi-language in a fully abstract way is the work on fully abstract compilation by Ahmed and Blume [2011] and New, Bowman, and Ahmed [2016] who show that their compiler’s source language embeds into their source-target multi-language in a fully abstract way. But the focus of this work was on fully abstract compilation, not on usability of user-facing languages.

The Eco project [Barrett, Bolz, Diekmann, and Tratt, 2016] is studying multi-language systems where user-exposed languages are combined in a very fine-grained way; it is closely related in that it studies the user experience in a multi-language system. The choice of an existing dynamic language creates delicate interoperability issues (conflicting variable scoping rules, etc.) as well as performance challenges. We propose a different approach, to design new multi-languages from scratch with interoperability in mind to avoid legacy obstacles.

We are not aware of existing systems exploiting the simple idea of using promotion to capture uniquely-owned state and dereliction to copy it—common formulations would rather perform copies on the contraction rule.

The general idea that linear types can permit reuse of unused allocated cells is not new. In Wadler [1990], a system is proposed with both linear and non-linear types to attack precisely this problem. It is however more distant from standard linear logic and somewhat ad-hoc; for example, there is no way to permanently turn a uniquely-owned value into a shared value, it provides instead a local *borrowing* construction that comes with ad-hoc restrictions necessary for safety. (The inability to *give up* unique ownership, which is essential in our

list-programming examples, seems to also be missing from Rust, where one would need to perform a costly operation of traversing the graph of the value to turn all pointers into `Arc` nodes.)

The RAML project [Hoffmann, Aehlig, and Hofmann, 2012] also combines linear logic and memory reuse: its *destructive match* operator will implicitly reuse consumed cells in new allocations occurring within the match body. Multi-languages give us the option to explore more explicit, flexible representations of those low-level concern, without imposing the complexity to all programmers.

A recent related work is the Cogent language [O’Connor, Chen, Rizkallah, Amani, Lim, Murray, Nagashima, Sewell, and Klein, 2016], in which linear state is also viewed as both functional and imperative – the latter view enabling memory reuse. The language design is interestingly reversed: in Cogent, the linear layer is the simple language that everyone uses, and the non-linear language is a complex but powerful language that is used when one really has to, named C.

Our linear language λ^L is sensibly simpler, and in several ways less expressive, than advanced programming languages based on linear logic [Tov and Pucella, 2011], separation logic [Balabonski, Pottier, and Protzenko, 2016], fine-grained permissions [Garcia, Tanter, Wolff, and Aldrich, 2014]: it is not designed to stand on its own, but to serve as a useful side-kick to a functional language, allowing safer resource handling.

One major simplification of our design compared to more advanced linear or separation-logic-based languages is that we do not separate physical locations from the logical capability/permission to access them (e.g., as in Ahmed, Fluet, and Morrisett [2007]). This restricts expressiveness in well-understood ways [Fahndrich and DeLine, 2002]: shared values cannot point to linear values.

Alms [Tov and Pucella, 2011], Quill [Morris, 2016] and Linear Haskell [Bernardy, Boespflug, Newton, Jones, and Spiwack, 2018] add linear types to a functional language, trying hard not to lose desirable usability property, such as type inference or the genericity of polymorphic higher-order functions. This is very challenging; for example, Linear Haskell gives up on principality of inference⁷. Our multi-language design side-steps this issue as the general-purpose language remains unchanged. Language boundaries are more rigid than an ideal no-compromise language, as they force users to preserve the distinction between the general-purpose and the advanced features; it is precisely this compromise that gives a design of reduced complexity.

Finally, on the side of the semantics, our system is related to LNL [Benton, 1994], a calculus for linear logic that, in a sense, is itself built as a multi-language system where (non-duplicable) linear types and (duplicable) intuitionistic types interact through a boundary. It is not surprising that our design contains an instance of this adjunction: for any σ there is a unique σ such that $\sigma \simeq !\sigma$, and converting a σ value to this σ and back gives a $!\sigma$ and is provably equivalent, by boundary cancellation, to just using `share`.

⁷ Thanks to Stephen Dolan for pointing out that $\lambda f.\lambda x.f x$ has several incompatible Linear Haskell types.

Acknowledgments

We thank our anonymous reviewers for their feedback, as well as Neelakantan Krishnaswami, François Pottier, Jennifer Paykin, Sylvie Boldo and Simon Peyton-Jones for our discussions on this work.

This work was supported in part by the National Science Foundation under grants CCF-1422133 and CCF-1453796, and the European Research Council under ERC Starting Grant SECOMP (715753). Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of our funding agencies.

Bibliography

- Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000. 2
- Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *International Conference on Functional Programming (ICFP), Victoria, British Columbia, Canada*, pages 157–168, September 2008. 2
- Amal Ahmed and Matthias Blume. An equivalence-preserving CPS translation via multi-language semantics. In *International Conference on Functional Programming (ICFP), Tokyo, Japan*, pages 431–444, September 2011. 2, 27
- Amal Ahmed, Matthew Fluet, and Greg Morrisett. L3 : A linear language with locations. *Fundamenta Informaticae*, 77(4):397–449, June 2007. 28
- Thibaut Balabonski, François Pottier, and Jonathan Protzenko. The design and formalization of Mezzo, a permission-based programming language. *ACM Transactions on Programming Languages and Systems*, 38(4):14:1–14:94, August 2016. 28
- Edd Barrett, Carl Friedrich Bolz, Lukas Diekmann, and Laurence Tratt. Fine-grained language composition: A case study. In *ECOOP*, July 2016. 16, 27
- P Nick Benton. A mixed linear and non-linear logic: Proofs, terms and models. In *International Workshop on Computer Science Logic*, 1994. 28
- Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. Linear haskell: practical linearity in a higher-order polymorphic language. *PACMPL*, 2(POPL):5:1–5:29, 2018. doi: 10.1145/3158093. URL <http://doi.acm.org/10.1145/3158093>. 28
- Robert Cartwright and Matthias Felleisen. Observable sequentiality and full abstraction. In *ACM Symposium on Principles of Programming Languages (POPL), Albuquerque, New Mexico*, pages 328–342, 1992. 2
- Dominique Devriese, Marco Patrignani, and Frank Piessens. Fully-abstract compilation by approximate back-translation. In *ACM Symposium on Principles of Programming Languages (POPL), St. Petersburg, Florida*, 2016. 2
- Manuel Fahndrich and Robert DeLine. Adoption and focus: Practical linear types for imperative programming. In *PLDI'02, PLDI '02*, 2002. 28
- Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. The teachscheme! project: Computing and programming for every student. *Computer Science Education*, 14(1):55–77, 2004. 27
- Ronald Garcia, Éric Tanter, Roger Wolff, and Jonathan Aldrich. Foundations of tpestate-oriented programming. *TOPLAS*, 2014. 28
- Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Resource Aware ML. In *24rd International Conference on Computer Aided Verification (CAV'12)*, volume 7358 of *Lecture Notes in Computer Science*, pages 781–786. Springer, 2012. 28

- Alan Jeffrey and Julian Rathke. Java Jr.: Fully abstract trace semantics for a core Java language. In *European Symposium on Programming (ESOP)*, 2005. 2
- Jacob Matthews and Robert Bruce Findler. Operational semantics for multi-language programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 31(3):12, 2009. 18, 27
- A. R. Meyer and K. Sieber. Towards fully abstract semantics for local variables. In *ACM Symposium on Principles of Programming Languages (POPL)*, San Diego, California, pages 191–203, 1988. 2
- Robin Milner. Fully abstract models of typed lambda calculi. *Theoretical Computer Science*, 4(1), 1977. 2
- J. Garrett Morris. The best of both worlds: Linear functional programming without compromise. In *ICFP*, 2016. 28
- Max S. New, William J. Bowman, and Amal Ahmed. Fully abstract compilation via universal embedding. In *International Conference on Functional Programming (ICFP)*, Nara, Japan, September 2016. 2, 27
- Liam O’Connor, Zilin Chen, Christine Rizkallah, Sidney Amani, Japheth Lim, Toby Murray, Yutaka Nagashima, Thomas Sewell, and Gerwin Klein. Refinement through restraint: Bringing down the cost of verification. In *ICFP*, 2016. 9, 28
- Peter-Michael Osera, Vilhelm Sjöberg, and Steve Zdancewic. Dependent interoperability. In *Programming Languages meets Program Verification (PLPV)*, January 2012. 27
- Marco Patrignani, Pieter Agten, Raoul Strackx, Bart Jacobs, Dave Clarke, and Frank Piessens. Secure compilation to protected module architectures. *ACM Transactions on Programming Languages and Systems*, 37(2):6:1–6:50, April 2015. 2
- Daniel Patterson, Jamie Perconti, Christos Dimoulas, and Amal Ahmed. FunTAL: Reasonably mixing a functional language with assembly. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Barcelona, Spain, June 2017. To appear. Available at <http://www.ccs.neu.edu/home/amal/papers/funtal.pdf>. 27
- James T Perconti and Amal Ahmed. Verifying an open compiler using multi-language semantics. In *ESOP*, 2014. 22, 27
- Gordon Plotkin and Martín Abadi. A logic for parametric polymorphism. In *Typed Lambda Calculi and Applications*, 1993. 32
- Jesse Tov and Riccardo Pucella. Stateful contracts for affine types. In *European Symposium on Programming (ESOP)*, March 2010. 27
- Jesse A Tov and Riccardo Pucella. Practical affine types. In *POPL*, 2011. 28
- Philip Wadler. Linear Types Can Change the World! In *Programming Concepts and Methods*, 1990. 27

A Logical relation

To define the logical relation in a precise way, we introduce a new grammar of “relational types” ρ, ρ , that simply extend the grammar of σ, σ to include a case for a relation on λ^U types (R, σ_1, σ_2) .

Introducing this lightweight syntax for relations here is a middle way between definitions that use an explicit relational substitution (cluttering the relation with yet another index) and a full-blown logic for parametricity as in [Plotkin and Abadi \[1993\]](#)

$$\begin{aligned} \rho &::= \dots \text{ all cases for } \sigma, \text{ but using } \rho, \rho \text{ recursively} \mid (R, \sigma_1, \sigma_2) \\ \rho &::= \text{ all cases for } \sigma, \text{ but using } \rho, \rho \text{ recursively} \end{aligned}$$

Fig. 12. Relation Type Syntax

Every ρ has two associated types, the types of terms that it relates, which we denote $(\rho)_1, (\rho)_2$. It is defined as follows:

$$\begin{array}{ll} ((R, \sigma_1, \sigma_2))_1 \stackrel{\text{def}}{=} \sigma_1 & (\rho_1 \otimes \rho_2)_i \stackrel{\text{def}}{=} \rho_1 \otimes \rho_2 \\ ((R, \sigma_1, \sigma_2))_2 \stackrel{\text{def}}{=} \sigma_2 & (1)_i \stackrel{\text{def}}{=} 1 \\ (\alpha)_i \stackrel{\text{def}}{=} \alpha & (\rho_1 \multimap \rho_2)_i \stackrel{\text{def}}{=} (\rho_1)_i \multimap (\rho_2)_i \\ (\rho_1 \times \rho_2)_i \stackrel{\text{def}}{=} (\rho_1)_i \times (\rho_2)_i & (\rho_1 \oplus \rho_2)_i \stackrel{\text{def}}{=} (\rho_1)_i \oplus (\rho_2)_i \\ (1)_i \stackrel{\text{def}}{=} 1 & (\mu\alpha. \rho)_i \stackrel{\text{def}}{=} \mu\alpha. (\rho)_i \\ (\rho_1 \rightarrow \rho_2)_i \stackrel{\text{def}}{=} (\rho_1)_i \rightarrow (\rho_2)_i & (\alpha)_i \stackrel{\text{def}}{=} \alpha \\ (\rho_1 + \rho_2)_i \stackrel{\text{def}}{=} (\rho_1)_i + (\rho_2)_i & (!\rho)_i \stackrel{\text{def}}{=} !(\rho)_i \\ (\mu\alpha. \rho)_i \stackrel{\text{def}}{=} \mu\alpha. (\rho)_i & (\text{Box } 1 \ \rho)_i \stackrel{\text{def}}{=} \text{Box } 1 \ (\rho)_i \\ (\forall\alpha. \rho)_i \stackrel{\text{def}}{=} \forall\alpha. (\rho)_i & (\text{Box } 0)_i \stackrel{\text{def}}{=} \text{Box } 0 \end{array}$$

First, we define when closed values are related at each type, indexing by a natural number to break the circularity of recursive types. The relations are defined by nested induction on j, ρ, ρ , any time a bigger type is used in a definition, the step-index j is decremented.

The definition of $\mathcal{E} \llbracket \sigma \rrbracket^j$ shows that this is an asymmetric relation capturing a notion of *approximation*, not equivalence.

Next, we extend the relations to open terms by defining open terms to be related when they are related when closed by related substitutions.

The Fundamental Property is the key to proving parametricity results.

Lemma 11 (Fundamental Property)

1. If $! \Gamma \vdash_v v : \sigma$ then $! \Gamma \vdash_v v \lesssim^{\log} v : \sigma$
2. If $! \Gamma \vdash_v e : \sigma$ then $! \Gamma \vdash e \lesssim^{\log} e : \sigma$
3. If $\Psi; \Gamma \vdash_L s \mid e : \sigma$ then $! \Gamma \vdash_L (s \mid e) \lesssim^{\log} (s \mid e) : \sigma$

$$\begin{aligned}
\text{Atom}[\sigma] &\stackrel{\text{def}}{=} \{v \mid \vdash_v v : \sigma\} \\
\text{Rel}[\sigma_1, \sigma_2] &\stackrel{\text{def}}{=} \{R : \mathbb{N} \rightarrow \mathcal{P}(\text{Atom}[\sigma_1] \times \text{Atom}[\sigma_2]) \mid \forall j \leq j'. R^{j'} \subset R^j\} \\
\mathcal{V}[(R, \sigma_1, \sigma_2)]^j &\stackrel{\text{def}}{=} R^j \\
\mathcal{V}[\mathbf{1}]^j &\stackrel{\text{def}}{=} \{(\langle \rangle, \langle \rangle)\} \\
\mathcal{V}[\rho \times \rho']^j &\stackrel{\text{def}}{=} \{(\langle v_1, v'_1 \rangle, \langle v_2, v'_2 \rangle) \mid (v_1, v_2) \in \mathcal{V}[\rho]^j \wedge (v'_1, v'_2) \in \mathcal{V}[\rho']^j\} \\
\mathcal{V}[\rho_1 + \rho_2]^j &\stackrel{\text{def}}{=} \{(\text{inj}_i v_1, \text{inj}_i v_2) \mid (v_1, v_2) \in \mathcal{V}[\rho_i]^j\} \\
\mathcal{V}[\mu\alpha. \rho]^j &\stackrel{\text{def}}{=} \{(\text{fold}_{(\mu\alpha. \rho)_1} v_1, \text{fold}_{(\mu\alpha. \rho)_2} v_1) \mid \forall j' < j. (v_1, v_2) \in \mathcal{V}[\rho[\mu\alpha. \rho/\alpha]]^{j'}\} \\
\mathcal{V}[\rho_1 \rightarrow \rho_2]^j &\stackrel{\text{def}}{=} \{(\lambda(x_1 : (\rho_1)_1). e_1, \lambda(x_2 : (\rho_2)_2). e_2) \mid \\
&\quad \forall j' \leq j, (v_1, v_2) \in \mathcal{V}[\rho_1]^{j'}. (e_1[v_1/x_1], e_2[v_2/x_2]) \in \mathcal{E}[\rho_2]^{j'}\} \\
\mathcal{V}[\forall\alpha. \rho]^j &\stackrel{\text{def}}{=} \{(\Lambda\alpha. v_1, \Lambda\alpha. v_2) \mid \forall \sigma_1, \sigma_2, R \in \text{Rel}[\sigma_1, \sigma_2]. (v_1, v_2) \in \mathcal{V}[\rho[(R, \sigma_1, \sigma_2)/\alpha]]^j\} \\
\mathcal{V}[\mathbf{1}]^j &\stackrel{\text{def}}{=} \{((\emptyset \mid \langle \rangle), (\emptyset \mid \langle \rangle))\} \\
\mathcal{V}[\rho \otimes \rho']^j &\stackrel{\text{def}}{=} \{((s_1 \# s'_1 \mid \langle v_1, v'_1 \rangle), (s_2 \# s'_2 \mid \langle v_2, v'_2 \rangle)) \mid \\
&\quad ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[\rho]^j \wedge ((s'_1 \mid v'_1), (s'_2 \mid v'_2)) \in \mathcal{V}[\rho']^j\} \\
\mathcal{V}[\rho_1 \oplus \rho_2]^j &\stackrel{\text{def}}{=} \{((s_1 \mid \text{inj}_i v_1), (s_2 \mid \text{inj}_i v_2)) \mid \\
&\quad ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[\rho_i]^j\} \\
\mathcal{V}[\mu\alpha. \rho]^j &\stackrel{\text{def}}{=} \{((s_1 \mid \text{fold}_{\mu\alpha. \rho} v_1), (s_2 \mid \text{fold}_{\mu\alpha. \rho} v_2)) \mid \\
&\quad \forall j' < j. ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[\rho[\mu\alpha. \rho/\alpha]]^{j'}\} \\
\mathcal{V}[\rho' \multimap \rho]^j &\stackrel{\text{def}}{=} \{((s_1 \mid \lambda(x : \rho'). e_1), (s_2 \mid \lambda(x : \rho'). e_2)) \mid \\
&\quad \forall j' \leq j, s'_1, s'_2, ((s'_1 \mid v_1), (s'_2 \mid v_2)) \in \mathcal{V}[\rho']^{j'}. \\
&\quad s'_1 = s_1 \# s'_1 \wedge s'_2 = s_2 \# s'_2 \Rightarrow \\
&\quad ((s'_1 \mid e_1[v_1/x]), (s'_2 \mid e_2[v_2/x])) \in \mathcal{E}[\rho]^{j'}\} \\
\mathcal{V}[\!|\rho|^j &\stackrel{\text{def}}{=} \{((\emptyset \mid \text{share}(s_1 : \Psi_1). v_1), (\emptyset \mid \text{share}(s_2 : \Psi_2). v_2)) \mid \\
&\quad ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[\rho]^j\} \\
\mathcal{V}[\text{Box } 0]^j &\stackrel{\text{def}}{=} \{([\ell_1 \mapsto \cdot] \mid \ell_1), ([\ell_2 \mapsto \cdot] \mid \ell_2)\} \\
\mathcal{V}[\text{Box } 1 \rho]^j &\stackrel{\text{def}}{=} \{([\ell_1 \mapsto (s_1 \mid v_1)] \mid \ell_1), ([\ell_2 \mapsto (s_2 \mid v_2)] \mid \ell_2)) \mid \\
&\quad ((s_1 \mid v_1), (s_2 \mid v_2)) \in \mathcal{V}[\rho]^j\} \\
\mathcal{V}[[\rho]^j &\stackrel{\text{def}}{=} \{((\emptyset \mid [v_1]), (\emptyset \mid [v_2])) \mid (v_1, v_2) \in \mathcal{V}[\rho]^j\} \\
\mathcal{E}[\rho]^j &\stackrel{\text{def}}{=} \{(e_1, e_2) \mid \forall j' \leq j. e_1 \xrightarrow{U}^{j'} v_1 \Rightarrow \\
&\quad \exists v_2. e_2 \xrightarrow{U}^* v_2 \wedge (v_1, v_2) \in \mathcal{V}[\rho]^{j-j'}\} \\
\mathcal{E}[\rho]^j &\stackrel{\text{def}}{=} \{((s_1 \mid e_1), (s_2 \mid e_2)) \mid \\
&\quad \forall j' \leq j, (s'_1 \mid v_1). (s_1 \mid e_1) \xrightarrow{L}^{j'} (s'_1 \mid v_1) \Rightarrow \\
&\quad \exists (s'_2 \mid v_2). (s_2 \mid e_2) \xrightarrow{L}^* (s'_2 \mid v_2) \wedge \\
&\quad ((s'_1 \mid v_1), (s'_2 \mid v_2)) \in \mathcal{V}[\rho]^{j-j'}\}
\end{aligned}$$

Fig. 13. Logical Approximation for Closed Terms

$$\begin{aligned}
\mathcal{G}[\cdot]^j &\stackrel{\text{def}}{=} \{((\emptyset, \emptyset) \mid \emptyset)\} \\
\mathcal{G}[\Gamma, \mathbf{x} : \sigma]^j &\stackrel{\text{def}}{=} \{((s_1 \# s'_1, s_2 \# s'_2) \mid \gamma[\mathbf{x} \mapsto (\mathbf{v}_1, \mathbf{v}_2)]) \mid \\
&\quad ((s_1, s_2) \mid \gamma) \in \mathcal{G}[\Gamma]^j \wedge ((s'_1 \mid \mathbf{v}_1), (s'_2 \mid \mathbf{v}_2)) \in \mathcal{V}[(\gamma)_R(\sigma)]^j\} \\
\mathcal{G}[\Gamma, \mathbf{x} : \sigma]^j &\stackrel{\text{def}}{=} \{((s_1, s_2) \mid \gamma[\mathbf{x} \mapsto (\mathbf{v}_1, \mathbf{v}_2)]) \mid \\
&\quad ((s_1, s_2) \mid \gamma) \in \mathcal{G}[\Gamma]^j \wedge (\mathbf{v}_1, \mathbf{v}_2) \in \mathcal{V}[(\gamma)_R(\sigma)]^j\} \\
\mathcal{G}[\Gamma, \alpha]^j &\stackrel{\text{def}}{=} \{((s_1, s_2) \mid \gamma[\alpha \mapsto (\mathbf{R}, \sigma_1, \sigma_2)]) \mid \\
&\quad \mathbf{R} \in \text{Rel}[\sigma_1, \sigma_2] \wedge ((s_1, s_2) \mid \gamma) \in \mathcal{G}[\Gamma]^j\} \\
!\Gamma \vdash \mathbf{e}_1 \lesssim^{\text{log}} \mathbf{e}_2 : \sigma &\stackrel{\text{def}}{=} \forall j \geq 0, ((\emptyset, \emptyset) \mid \gamma) \in \mathcal{G}[\Gamma]^j \cdot ((\gamma)_1(\mathbf{e}_1), (\gamma)_2(\mathbf{e}_2)) \in \mathcal{E}[(\gamma)_R(\sigma)]^j \\
!\Gamma \vdash_v \mathbf{v}_1 \lesssim^{\text{log}} \mathbf{v}_2 : \sigma &\stackrel{\text{def}}{=} \forall j \geq 0, ((\emptyset, \emptyset) \mid \gamma) \in \mathcal{G}[\Gamma]^j \cdot ((\gamma)_1(\mathbf{v}_1), (\gamma)_2(\mathbf{v}_2)) \in \mathcal{V}[(\gamma)_R(\sigma)]^j \\
\Gamma \vdash_L (s_1 \mid \mathbf{e}_1) \lesssim^{\text{log}} (s_2 \mid \mathbf{e}_2) : \sigma &\stackrel{\text{def}}{=} \\
\forall j \geq 0, ((s'_1, s'_2) \mid \gamma) \in \mathcal{G}[\Gamma]^j \cdot \\
((s_1 \# s'_1 \mid (\gamma)_1(\mathbf{e}_1)), (s_2 \# s'_2 \mid (\gamma)_2(\mathbf{e}_2))) &\in \mathcal{E}[(\gamma)_R(\sigma)]^j
\end{aligned}$$

Fig. 14. Logical Approximation for Open Terms

$$\begin{aligned}
!\Gamma \vdash \mathbf{e}_1 \lesssim^{\text{ctx}} \mathbf{e}_2 : \sigma &\stackrel{\text{def}}{=} \forall C. \cdot \vdash_v C[\mathbf{e}_1] : \mathbf{1} \wedge \cdot \vdash_v C[\mathbf{e}_2] : \mathbf{1} \wedge C[\mathbf{e}_1] \overset{\text{U}}{\hookrightarrow^*} \langle \rangle \implies C[\mathbf{e}_2] \overset{\text{U}}{\hookrightarrow^*} \langle \rangle \\
\Gamma \vdash_L (s_1 \mid \mathbf{e}_1) \lesssim^{\text{ctx}} (s_2 \mid \mathbf{e}_2) : \sigma &\stackrel{\text{def}}{=} \forall C. \cdot \vdash_v C[(s_1 \mid \mathbf{e}_1)] : \mathbf{1} \wedge \cdot \vdash_v C[\mathbf{e}_1](s_2 \mid \mathbf{e}_2) : \\
\mathbf{1} \wedge C[(s_1 \mid \mathbf{e}_1)] \overset{\text{U}}{\hookrightarrow^*} \langle \rangle &\implies C[(s_2 \mid \mathbf{e}_2)] \overset{\text{U}}{\hookrightarrow^*} \langle \rangle
\end{aligned}$$

Fig. 15. Contextual Approximation

Finally, we prove our logical relation is sound with respect to contextual equivalence, that is, it can be used as a more tractable way to prove contextual equivalence results, such as lump/unlump cancellation.

Theorem 6 (Soundness of Logical Relation) $\lesssim^{log} \subset \lesssim^{ctx}$

The proof is by induction on contexts, showing that every term formation rule preserves logical relatedness. These “compatibility” lemmas are extensive, but their proofs are simple. Their proofs are in the extended technical report.