

# State-Dependent Representation Independence

Amal Ahmed

TTI-C  
amal@tti-c.org

Derek Dreyer

MPI-SWS  
dreyer@mpi-sws.mpg.de

Andreas Rossberg

MPI-SWS  
rossberg@mpi-sws.mpg.de

## Abstract

Mitchell’s notion of *representation independence* is a particularly useful application of Reynolds’ relational parametricity — two different implementations of an abstract data type can be shown contextually equivalent so long as there exists a relation between their type representations that is preserved by their operations. There have been a number of methods proposed for proving representation independence in pure extensions of System F, as well as in impure Java-like languages. However, none of the existing work on representation independence addresses the interaction of existential type abstraction and mutable state. Specifically, none shows how to prove equivalence of *generative* ADTs, for which each instance of the ADT must introduce a distinct abstract type because the relational interpretation of that type depends on some local instance-specific state.

In this work, we present a method for proving *state-dependent representation independence*. Our method extends Ahmed’s previous work on step-indexed logical relations for recursive and quantified types in order to handle ML-style unrestricted state. We use this method to prove a number of interesting contextual equivalences, some drawn from the literature on type generativity, that involve a close interaction between existentials and state. In these examples, the relational interpretations of the abstract types may be initially empty and then grow over time in a manner that is tightly coupled with changes to some local state. We encode such state-dependent type representations using a novel possible-worlds model that is inspired by several recent approaches but relies on step-indexing to enable a well-founded quasi-circular definition of possible worlds.

## 1. Introduction

Reynolds’ notion of *relational parametricity* [20] is the essence of type abstraction — clients of an abstract type behave uniformly across *all* relational interpretations of that type and thus cannot depend in any way on how the type is represented. Mitchell’s notion of *representation independence* [15] is a particularly useful application of relational parametricity — two different implementations of an abstract data type can be shown contextually equivalent so long as there *exists* a relation between their type representations that is preserved by their operations. This is useful even when the type representations of the two ADTs are the same, because the choice of an arbitrary relational interpretation for the abstract type allows one to establish the existence of local invariants.

Originally these ideas were developed in the context of (variants of) System F, but over the last two decades there has been a great deal of work on extending them to the setting of more realistic languages, such as those with recursive functions [17], general recursive types [14, 2, 10], selective strictness [26], etc. Others have considered representation independence in an impure Java-like object-oriented setting, in which abstraction is achieved through nominal class typing and `private` access modifiers [6, 7, 13].

None of the existing work, however, studies representation independence in the context of a language supporting *both* existential type abstraction (as in System F) and mutable state (as in ML or Java). Why should one care about such a language?

### 1.1 Reasoning About Generative Type Abstraction

Existential type abstraction provides *type generativity* — every unpacking of an existential package generates a *fresh* abstract type that is distinct from any other. This is similar to the behavior of Standard ML’s *generative functors*, which generate fresh abstract types at each application, and indeed the semantics of SML-style functors may be understood as a stylized use of existential type abstraction [22]. The clearest motivation for type generativity is in the definition of ADTs that encapsulate some local state. In such instances, generativity is sometimes *necessary* to achieve the proper degree of data abstraction.

As a simple motivating example, consider the SML module code in Figure 1, which is adapted from an example of Dreyer *et al.* [11]. (Later in the paper, we will develop a similar example using existential types.) Here, the signature `SYMBOL` describes a module implementing a mutable symbol table, which maps “symbols” to strings. The module provides an abstract type `t` describing the symbols currently in its table; a function `eq` for comparing symbols for equality; a function `insert`, which adds a given string to the table and returns a fresh symbol mapped to it; and a function `lookup`, which looks up a given symbol in the table and returns the corresponding string.

The functor `Symbol` implements the symbol type `t` as an integer index into a (mutable) stack of strings. When applied, `Symbol` creates a fresh `table` (represented as a pointer to an empty list) and a mutable counter `size` (representing the size of the table). The implementations of the various functions are straightforward, and the body of the functor is sealed with the signature `SYMBOL`, thus hiding access to the local state (`table` and `size`).

The call to `List.nth` in the `lookup` function might in general raise a `Subscript` exception if the input `n` were an arbitrary integer. However, we “know” that this cannot happen because `lookup` is exported with argument type `t`, and the only values of type `t` that a client could possibly have gotten hold of are the values returned by `insert`, *i.e.*, integers that are between 1 and the current size of `table`. Therefore, the implementation of the `lookup` function need not bother handling the `Subscript` exception.

This kind of reasoning is commonplace in modules that encapsulate local state. But what justifies it? Intuitively, the answer is

---

```

signature SYMBOL = sig
  type t
  val eq : t * t -> bool
  val insert : string -> t
  val lookup : t -> string
end
functor Symbol () :> SYMBOL = struct
  type t = int
  val size = ref 0
  val table = ref nil
  fun eq (x,y) = x = y
  fun insert str = (
    size := !size + 1;
    table := str :: !table;
    !size
  )
  fun lookup n =
    List.nth (!table, !size - n)
end

```

**Figure 1.** Generativity Example

type generativity. Each instantiation of the `Symbol` functor creates a fresh symbol type `t`, which represents the type of symbols that are valid in its own `table` (but not any other). Were `Symbol` not generative, each application of the `Symbol` functor would produce a module with distinct local state but the *same* symbol type. It would then be easy to induce a `Subscript` error by accidentally passing a value of one `table`'s symbol type to another's `lookup` function.<sup>1</sup>

While this intuition about the importance of generativity is very appealing, it is also completely informal, and we are not aware of any prior work that actually formalizes it. The goal of this paper is to develop a formal framework for reasoning about the interaction of generative type abstraction and mutable state.

In the case of an example like the `Symbol` functor, we will be able to show that the implementation of `Symbol` shown in Figure 1 is *contextually equivalent* to one whose `lookup` function is replaced by:

```

fun lookup n =
  if n > 0 andalso n <= !size
    andalso !size = length(!table)
  then List.nth (!table, !size - n)
  else "Hell freezes over"

```

In other words, there is no observable difference between the original `Symbol` functor and one that dynamically checks the various invariants we claim to “know.” Hence, the checks are unnecessary.

This kind of result can be understood as an instance of representation independence, albeit a somewhat degenerate one in that the ADTs in question share the same type representation. As with most such results, the proof hinges on the construction of an appropriate relational interpretation of the abstract type `t`, which serves to impose an invariant on the possible values of type `t`. In this case, we wish to assert that for a particular structure `S` defined by `Symbol()`, the only values of type `S.t` are integers between 1 and the current size of `S.table`. This will allow us to prove that any range check on the argument to `S.lookup` function is superfluous.

But ay, there's the rub: the relational interpretation we wish to assign to `S.t` depends on the *current* values stored in `S`'s local state. In effect, as `S.insert` function is called repeatedly over time, its `table` grows larger, and the relational interpretation of `S.t` must grow accordingly to include more and more integers.

Thus, what we need is an account of *state-dependent representation independence*, in which the relational interpretations of ab-

<sup>1</sup>This is the case, for example, in OCaml, which only supports *applicative* (i.e., non-generative) functors.

stract types are permitted to *grow* over time, in a manner that is tightly coupled with changes to some local state. To our knowledge, such an account does not currently exist.

## 1.2 Overview

In this paper, we present a novel method for proving state-dependent representation independence results. Our method extends previous work by Ahmed on syntactic *step-indexed* logical relations for recursive and quantified types [2]. We extend her technique with support for reasoning about local state, and demonstrate its effectiveness on a variety of small but representative examples.

In order to reason about local state, we build into our logical relation a notion of *possible worlds*. While several aspects of our possible worlds are derived from and inspired by prior work, other aspects are quite novel:

1. We enrich our possible worlds with *populations* and *laws*, which allow us to evolve the relational interpretation of an abstract type over time in a controlled, state-dependent fashion. For instance, we can use a population to grow a set of values (e.g., the integers between 1 and some  $n$ ), together with a law that explains what the current population implies about the current machine state (e.g., that the symbol table has size  $n$ ).
2. Second, our method provides the ability to reason locally about references to higher-order values. While ours is not the first method to handle higher-order state, our approach is novel and depends critically on step-indexing to support a quasi-circular definition of possible worlds.

The remainder of the paper is structured as follows. In Section 2, we present  $F^{\mu^1}$ , our language under consideration, which is essentially System F extended with general recursive types and general ML-style references. In Section 3, we explain at a high level how our method works and what is novel about it. In Section 4, we present the details of our logical relation and prove it sound (but not complete) with respect to contextual equivalence of  $F^{\mu^1}$  programs. In Section 5, we show how to use our method to prove a number of interesting contextual equivalences. Finally, in Section 6, we conclude with a thorough comparison to related work, as well as directions for future work.

## 2. The Language $F^{\mu^1}$

We consider  $F^{\mu^1}$ , a call-by-value  $\lambda$ -calculus with impredicative polymorphism, iso-recursive types, and general ML-style references. The syntax of  $F^{\mu^1}$  is shown in Figure 2, together with excerpts of the static and dynamic semantics. We assume an infinite set of locations *Loc* ranged over by  $l$ . Our term language includes equality on references ( $e_1 == e_2$ ), but is otherwise standard.

We define a small-step operational semantics for  $F^{\mu^1}$  as a relation between configurations  $(s, e)$ , where  $s$  is a global store mapping locations  $l$  to values  $v$ . We use evaluation contexts  $E$  to lift the primitive reductions to a standard left-to-right call-by-value semantics for the language. We elide the syntax of evaluation contexts as it is completely standard, and we show only some of the reduction rules in Figure 2.

$F^{\mu^1}$  typing judgments have the form  $\Delta; \Gamma; \Sigma \vdash e : \tau$  where the contexts  $\Delta$ ,  $\Gamma$ , and  $\Sigma$  are defined as in Figure 2. The type context  $\Delta$  is used to track the set of type variables in scope; the value context  $\Gamma$  is used to track the term variables in scope (along with their types  $\tau$ , which must be well formed in context  $\Delta$ , written  $\Delta \vdash \tau$ ); and the store typing  $\Sigma$  tracks the types of the contents of locations in the store. Note that  $\Sigma$  maps locations to closed types. We write  $FTV(\tau)$  to denote the set of type variables that appear free in type  $\tau$ . The typing rules are entirely standard, so we show only

Types	$\tau ::= \alpha \mid \text{unit} \mid \text{int} \mid \text{bool} \mid \tau_1 \times \tau_2 \mid \tau_1 \rightarrow \tau_2 \mid \forall \alpha. \tau \mid \exists \alpha. \tau \mid \mu \alpha. \tau \mid \text{ref } \tau$
Prim Ops	$o ::= + \mid - \mid = \mid < \mid \leq \mid \dots$
Terms	$e ::= x \mid () \mid l \mid n \mid o(e_1, \dots, e_n) \mid \text{true} \mid \text{false} \mid \text{if } e \text{ then } e_1 \text{ else } e_2 \mid \langle e_1, e_2 \rangle \mid \text{fst } e \mid \text{snd } e \mid \lambda x : \tau. e \mid e_1 e_2 \mid \Lambda \alpha. e \mid e[\tau] \mid \text{pack } \tau, e \text{ as } \exists \alpha. \tau' \mid \text{unpack } e_1 \text{ as } \alpha, x \text{ in } e_2 \mid \text{fold } e \mid \text{unfold } e \mid \text{ref } e \mid !e \mid e_1 := e_2 \mid e_1 == e_2$
Values	$v ::= () \mid l \mid n \mid \text{true} \mid \text{false} \mid \langle v_1, v_2 \rangle \mid \lambda x : \tau. e \mid \Lambda \alpha. e \mid \text{pack } \tau_1, v \text{ as } \exists \alpha. \tau \mid \text{fold } v$
$\frac{s, (\lambda x : \tau. e) v \mapsto s, [v/x]e \quad s, (\Lambda \alpha. e)[\tau] \mapsto s, [\tau/\alpha]e \quad s, \text{unpack } (\text{pack } \tau, v \text{ as } \exists \alpha. \tau') \text{ as } \alpha, x \text{ in } e \mapsto s, [\tau/\alpha][v/x]e \quad s, \text{unfold } (\text{fold } v) \mapsto s, v \quad s, \text{ref } v \mapsto s[l \mapsto v], l \quad \text{where } l \notin \text{dom}(s) \quad s, !l \mapsto s, v \quad \text{where } s(l) = v \quad s, l := v \mapsto s[l \mapsto v], () \quad \text{where } l \in \text{dom}(s) \quad s, l == l \mapsto s, \text{true} \quad s, l == l' \mapsto s, \text{false} \quad \text{where } l \neq l' \quad s, e \mapsto s', e' \quad s, E[e] \mapsto s', E[e']}{s, E[e] \mapsto s', E[e']}$	
Type Contexts	$\Delta ::= \cdot \mid \Delta, \alpha$
Value Contexts	$\Gamma ::= \cdot \mid \Gamma, x : \tau$
Store Typing	$\Sigma ::= \cdot \mid \Sigma, l : \tau \quad \text{where } FTV(\tau) = \emptyset$
$\frac{\Gamma(x) = \tau \quad \dots \quad \Sigma(l) = \tau \quad \Delta; \Gamma; \Sigma \vdash e : \tau}{\Delta; \Gamma; \Sigma \vdash x : \tau} \quad \frac{\Delta; \Gamma; \Sigma \vdash e : \text{ref } \tau \quad \Delta; \Gamma; \Sigma \vdash e_1 : \text{ref } \tau \quad \Delta; \Gamma; \Sigma \vdash e_2 : \tau}{\Delta; \Gamma; \Sigma \vdash !e : \tau} \quad \frac{\Delta; \Gamma; \Sigma \vdash e_1 : \text{ref } \tau \quad \Delta; \Gamma; \Sigma \vdash e_2 : \text{ref } \tau}{\Delta; \Gamma; \Sigma \vdash e_1 == e_2 : \text{bool}}$	
Well-typed Store:	$\frac{\forall l \in \text{dom}(\Sigma). \cdot; \Sigma \vdash s(l) : \Sigma(l)}{\vdash s : \Sigma}$

**Figure 2.**  $F^{\mu 1}$  Syntax + Dynamic and Static Semantics (excerpts)

a few rules in Figure 2. We refer the reader to the online technical appendix [4] for full details of  $F^{\mu 1}$ .

### 2.1 Contextual Equivalence

A context  $C$  is an expression with a single hole  $[\cdot]$  in it. Typing judgments for contexts have the form  $\vdash C : (\Delta; \Gamma; \Sigma \vdash \tau) \Rightarrow (\Delta'; \Gamma'; \Sigma' \vdash \tau')$ , where  $(\Delta; \Gamma; \Sigma \vdash \tau)$  indicates the type of the hole. Essentially, this judgment says that if  $e$  is an expression such that  $\Delta; \Gamma; \Sigma \vdash e : \tau$ , then  $\Delta'; \Gamma'; \Sigma' \vdash C[e] : \tau'$ . The typing rule for a hole  $[\cdot]$  is as follows:

$$\frac{\Delta \subseteq \Delta' \quad \Gamma \subseteq \Gamma' \quad \Sigma \subseteq \Sigma'}{\vdash [\cdot] : (\Delta; \Gamma; \Sigma \vdash \tau) \Rightarrow (\Delta'; \Gamma'; \Sigma' \vdash \tau')}$$

The other rules are straightforward (see our online appendix [4]).

#### Definition 2.1 (Contextual Approximation & Equivalence)

Let  $\Delta; \Gamma; \Sigma \vdash e_1 : \tau$  and  $\Delta; \Gamma; \Sigma \vdash e_2 : \tau$ .

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

$$\Delta; \Gamma; \Sigma \vdash e_1 \approx^{ctx} e_2 : \tau \stackrel{\text{def}}{=} \Delta; \Gamma; \Sigma \vdash e_1 \preceq^{ctx} e_2 : \tau \wedge \Delta; \Gamma; \Sigma \vdash e_2 \preceq^{ctx} e_1 : \tau$$

## 3. The Main Ideas

In this section we give an informal overview of the main novel ideas in our method, and how it compares to some previous approaches.

### 3.1 Logical Relations

Broadly characterized, our approach is a *logical relations* method. We define a relation  $\mathcal{V} \llbracket \tau \rrbracket \rho$ , which relates pairs of values at a type  $\tau$ , where the free type variables of  $\tau$  are given relational interpretations in  $\rho$ . The relation is “logical” in the sense that functions are defined to be related (at an arrow type) iff, when applied to related arguments, they return related results. We will show that this logical relation is sound with respect to contextual equivalence for  $F^{\mu 1}$ . This is useful because, for many examples, it is much easier to show that two programs are in the logical relation than to show that they are contextually equivalent directly.

Logical relations methods are among the oldest techniques for proving representation independence results. We will assume the reader is generally familiar with the flavor of these techniques, and instead focus on what is distinctive and original about ours.

### 3.2 Local Reasoning via Possible Worlds and Islands

As explained in Section 1, our core contribution is the idea of *state-dependent* relational interpretations of abstract types. That is, whether two values are related by some abstract type’s relational interpretation may depend on the current state of the heap. But when defining such a relational interpretation, how can we characterize the “current state of the heap?”

As a starting point, we review the general approach taken by a number of prior works on reasoning about local state [18, 19, 8, 9]. This approach, which utilizes a *possible worlds* model, has influenced us greatly, and constitutes the foundation of our method. However, the form it has taken in prior work is insufficient for our purposes, and it is instructive to see why.

The general approach of these prior works is to index the logical relation not only by a type  $\tau$  but by a *world*  $W$ . Instead of characterizing the current state of the heap,  $W$  characterizes the properties we expect the heap to have. In other words, it is a relation on machine stores, and we restrict attention to pairs of stores that satisfy it. If two values  $v_1$  and  $v_2$  are in the logical relation at type  $\tau$  and world  $W$ , then it means they are related when considered under any two stores  $s_1$  and  $s_2$ , respectively, that satisfy  $W$ .

Worlds in turn are constructed as a separating conjunction of smaller worlds  $\langle w_1, \dots, w_n \rangle$ , sometimes called *islands*, where each island is a relation that “concerns” a disjoint piece of the store. Intuitively, this means that each island distinguishes between pairs of stores only on the basis of a particular set of memory locations, and the set of locations that one island cares about is disjoint from the set that any other one cares about.

Exactly how the separation criterion on islands is formalized is immaterial; the important point is that it enables local reasoning. Suppose we want to prove that one expression is related to another in world  $W$ . Each may allocate some fresh piece of the store, and before showing that the resulting values of the expressions are related, we are permitted to extend  $W$  with a new island  $w$  describing how these fresh pieces of the store relate to each other. World extension is sound here precisely because the new island is (due to freshness of allocation) separate from the others. So long as the expressions in question do not make the locations in their local stores publicly accessible, no other part of the program is capable of mutating the store in such a manner as to violate  $w$ .

To make things concrete and to observe the limitations of possible worlds (as we have described them), let us consider the motivating example from Section 1. To prove that the two implementations of the `Symbol` functor are contextually equivalent, we will show that their bodies are logically related in an arbitrary ini-

tial world  $W_0$ . Both functors allocate local state in the same way, namely by allocating one pointer for `table` and one for `size`, so we will want to extend  $W_0$  with an island  $w_{\text{sym}}$  describing the local invariants on `table` and `size`. How should we define  $w_{\text{sym}}$ ?

One useful invariant that  $w_{\text{sym}}$  can enforce is that, for both implementations of `Symbol`, the integer pointed to by `size` is equal to the length of the list pointed to by `table`. By incorporating this property into  $w_{\text{sym}}$ , we will be guaranteed that, in any *future world* (i.e., any extension of  $W_0 \uplus w_{\text{sym}}$ ) in which the `lookup` function is called, the dynamic check `!size = length(!table)` in the second implementation of `Symbol` will always evaluate to `true`.

We can also use  $w_{\text{sym}}$  to enforce that `!size` is the same in the stores of both programs, and similarly for `!table`. Unfortunately, while this is a necessary condition, it is not sufficient to prove that the range check on the argument of `lookup` in the second `Symbol` implementation always evaluates to `true`. For that, we need a way of correlating the value of `!size` and the possible values of type `t`, but the islands we have developed thus far do not provide one.

### 3.3 Populating the Islands and Enforcing the Laws

The problem with islands is that they are static entities with no potential for development. To address this limitation, we enrich islands with *populations*. A population is a set of values that “inhabit” an island and affect the definition of the store relation for that island. An island’s population may *grow* over time (i.e., as we move to future worlds), and its store relation may change accordingly. In order to control population growth, we equip every island with an immutable *law* governing the connection between its population and its store relation. We denote populations by  $V$ , store relations by  $\psi$ , and laws by  $\mathcal{L}$ .

Consider the `Symbol` example. Let us define  $V_n = \{1, \dots, n\}$ , and let  $\psi_n$  be the store relation containing pairs of stores that obey the properties concerning `table` and `size` described in Section 3.2 and that, in addition, both map the location `size` to  $n$ . The idea is that  $V_n$  describes the set of values of type `t`, given that the current stores satisfy  $\psi_n$ . Thus, when we extend the initial world  $W_0$  with an island  $w_{\text{sym}}$  governing `Symbol`’s local state, we will choose that  $w_{\text{sym}}$  to comprise population  $V_0$ , store relation  $\psi_0$ , and a law  $\mathcal{L}$  defined as  $\{(\psi_n, V_n) \mid n \geq 0\}$ . Here,  $V_0$  and  $\psi_0$  characterize the initial world, in which there are no values of type `t` and the size of the table is 0. The law  $\mathcal{L}$  describes what future populations and store relations on this island may look like. In particular,  $\mathcal{L}$  enforces that future populations may contain 1 to  $n$  (for any  $n$ ), but only in conjunction with stores that map `size` to  $n$ . (Of course, the initial population and store relation must also obey the law, which they do.) An island’s law is established when the island is first added to the world and may not be amended in future worlds.

Having extended the world  $W_0$  with this new island  $w_{\text{sym}}$ , we are now able to define a relational interpretation for the type `t`, namely: values  $v_1$  and  $v_2$  are related at type `t` in world  $W$  if  $v_1 = v_2 = n$ , where  $n$  belongs to the population of  $w_{\text{sym}}$  in  $W$ . In proving equivalence of the two versions of the `lookup` function, we can assume that we start with stores  $s_1$  and  $s_2$  that are related by some world  $W$ , where  $W$  is a future world of  $W_0 \uplus w_{\text{sym}}$ , and that the arguments to `lookup` are related at type `t` in  $W$ . Consequently, given the law that we established for  $w_{\text{sym}}$  together with the interpretation of `t`, we know that the arguments to `lookup` must both equal some  $n$ , that the current population of  $w_{\text{sym}}$  must be some  $V_m$ , where  $1 \leq n \leq m$ , and that the current store relation must be  $\psi_m$ . Since  $s_1$  and  $s_2$  satisfy  $W$ , they must satisfy  $\psi_m$ , which means they map `size` to  $m \geq n$ . Hence, the dynamic range check in the second version of `Symbol` must evaluate to `true`.

For the above relational interpretation of `t` to make sense, we clearly need to be able to refer to a particular island in a world (e.g.,  $w_{\text{sym}}$ ) by some unique identifier that works in all future worlds. We

achieve this by insisting that a world be an ordered list of islands, and that new islands only be added to the end of the list. This allows us to access islands by their position in the list, which stays the same in future worlds.

In addition, an important property of the logical relation, which relational interpretations of abstract types must thus obey as well, is closure under world extension, i.e., that if two values are related in world  $W$ , then they are related in any future world of  $W$ . To ensure closure under world extension for relations that depend on their constituents’ inhabitation of a particular island (such as the relation used above to interpret `t`), we require that island populations can only grow larger in future worlds, not smaller.

### 3.4 References to Higher-Order Values

Most prior possible-worlds logical-relation approaches to reasoning about local state impose serious restrictions on what can be stored in memory. Pitts and Stark [18], for example, only allow references to integers. Reddy and Yang [19] and Benton and Leperchey [8] additionally allow references to *data*, which include integers and pointers but not functions. While the exploration of relational reasoning in less expressive languages is certainly a worthwhile endeavor, in this work we are concerned with contextual equivalence in the presence of unrestricted state.

To see what (we think) is tricky about handling references to higher-order values, suppose we have two programs that both maintain some local state, and we are trying to prove these programs equivalent. Say the invariant on this local state, which we will enforce using an island  $w$ , is very simple: the value that the first program stores in location  $l_1$  is logically related to the value that the second program stores in  $l_2$ . If these values were just integers, we could write the law for  $w$  (as we did in the `Symbol` example) so that in any future world,  $w$ ’s store relation  $\psi$  must demand that two stores  $s_1$  and  $s_2$  are related only if  $s_1(l_1) = s_2(l_2)$ . This works because at type `int`, the logical relation coincides with equality.

If the locations have some higher type  $\tau$ , however, the definition of  $w$ ’s store relation  $\psi$  will need to relate  $s_1(l_1)$  and  $s_2(l_2)$  using the *logical relation* at type  $\tau$ , not mere syntactic equality. But the problem is: logical relations are indexed by worlds. In order for  $\psi$  to say that  $s_1(l_1)$  and  $s_2(l_2)$  are related at type  $\tau$ , it needs to specify the world  $W$  in which their relation is being considered.

Bohr and Birkedal [9] address this issue by imposing a rigid structure on their store relations. Specifically, instead of having a single store relation per island, they employ a *local parameter*, which is roughly a set of pairs of the form  $(P, LL)$ , where  $P$  is a store relation and  $LL$  is a finite set of pairs of locations (together with a closed type). The way to interpret this local parameter is that the current stores must satisfy *one* of the  $P$ ’s, and all the pairs of locations in the corresponding  $LL$  must be related by the logical relation in the *current* world. In the case of our example with  $l_1$  and  $l_2$ , they would define a local parameter  $\{(P, LL)\}$ , where  $P$  is the universal store relation and  $LL = \{(l_1, l_2, \tau)\}$ . Bohr and Birkedal’s approach effectively uses the  $LL$ ’s to abstract away explicit references to the world-indexed logical relation within the store relation. This avoids the need to refer to a specific world inside a store relation, but it only works for store relations that are expressible in the highly stylized form of these local parameters.

Instead, our approach is to *parameterize* store relations over the world in which they will be considered. Then, in defining what it means for two stores  $s_1$  and  $s_2$  to satisfy some world  $W$ , we require that for every  $\psi$  in  $W$ ,  $(s_1, s_2) \in \psi[W]$ , i.e.,  $s_1$  and  $s_2$  obey  $\psi$  when it is instantiated to the current world  $W$ . The astute reader will have noticed, however, that this parameterization introduces a circularity: worlds are defined to be collections of store relations, which are now parameterized by worlds. To break this circularity, we employ *step-indexed* logical relations.

### 3.5 Step-Indexed Logical Relations and Possible Worlds

Appel and McAllester [5] introduced the step-indexed model as a way to express *semantic* type soundness proofs for languages with general recursive and polymorphic types. Although its original motivation was tied to foundational proof-carrying code, the technique has proven useful in a variety of applications. In particular, Ahmed [2] has used a binary version of Appel and McAllester’s model for relational reasoning about System F extended with general recursive types, and it is her work that we build on.

The basic idea is closely related to classic constructions from domain theory. We define the logical relation  $\mathcal{V} \llbracket \tau \rrbracket \rho$  as the limit of an infinite chain of approximation relations  $\mathcal{V}_n \llbracket \tau \rrbracket \rho$ , where  $n \geq 0$ . Informally, values  $v_1$  and  $v_2$  are related by the  $n$ -th approximation relation only if they are indistinguishable in any context for  $n$  steps of computation. (They might be distinguishable *after*  $n$  steps, but we don’t care because the “clock” has run out.) Thus, values are related in the limit only if they are indistinguishable in any context for *any* finite number of steps, *i.e.*, if they are really indistinguishable.

The step-indexed stratification makes it possible to define the semantics of recursive types quite easily. Two values `fold`  $v_1$  and `fold`  $v_2$  are defined to be related by  $\mathcal{V}_n \llbracket \mu\alpha. \tau \rrbracket \rho$  iff  $v_1$  and  $v_2$  are related by  $\mathcal{V}_k \llbracket [\mu\alpha. \tau / \alpha] \tau \rrbracket \rho$  for all  $k < n$ . Even though the unfolded type is larger (usually a deal breaker for logical relations, which are typically defined by induction on types), the step index gets smaller, so the definition of the logical relation is well-founded. Moreover, it makes sense for the step index to get smaller, since it takes a step of computation to extract  $v_i$  from `fold`  $v_i$ .

Just as we use steps to stratify logical relations, we can also use them to stratify our quasi-circular possible worlds. We define an “ $n$ -level world” inductively to be one whose constituent store relations (the  $\psi$ ’s) are parameterized by  $(n-1)$ -level worlds. The intuition behind this stratification of worlds is actually very simple: an  $n$ -level world describes properties of the current stores that may affect the relatedness of pairs of values for  $n$  steps of computation. Since it takes one step of computation just to *inspect* the stores (via a pointer dereference), the relatedness of pairs of values for  $n$  steps can only possibly depend on the relatedness of the current stores for  $n-1$  steps. Thus, it is fine for an  $n$ -level world to be defined as a collection of  $(n-1)$ -level store relations, *i.e.*,  $\psi$ ’s that only guarantee relatedness of memory contents for  $n-1$  steps. And these  $(n-1)$ -level  $\psi$ ’s, in turn, need only be parameterized by  $(n-1)$ -level worlds.

## 4. Step-Indexed Logical Relations for $F^{\mu 1}$

In this section, we present the details of our logical relation for  $F^{\mu 1}$  and prove it sound with respect to contextual equivalence.

The basic idea is to give a relational interpretation  $\mathcal{V} \llbracket \tau \rrbracket$  of a (closed) type  $\tau$  as a set of tuples of the form  $(k, W, v_1, v_2)$ , where  $k$  is a natural number (called the *step index*),  $W$  is a world as motivated in Section 3, and  $v_1$  and  $v_2$  are values. Informally,  $(k, W, v_1, v_2) \in \mathcal{V} \llbracket \tau \rrbracket$  says that in any computation running for no more than  $k$  steps,  $v_1$  approximates  $v_2$  at the type  $\tau$  in world  $W$ . An important point is that to determine if  $v_1$  approximates  $v_2$  for  $k$  steps (at type  $\tau$ ), it suffices for the world  $W$  to be a  $k$ -level world. That is, the store relations  $\psi$  in  $W$  need only guarantee relatedness of memory contents for  $k-1$  steps as discussed in Section 3.5. We make the notion of a “ $k$ -level world” precise in Section 4.1.

**Preliminaries** In the rest of the paper, the metavariables  $i, j, k, m$ , and  $n$  all range over natural numbers. We use the metavariable  $\chi$  to denote sets of tuples of the form  $(k, W, e_1, e_2)$  where  $k$  is a step index,  $W$  is a world, and  $e_1$  and  $e_2$  are closed terms (*i.e.*, terms that may contain locations, but no free type or term variables). Given a set  $\chi$  of this form, we write  $\chi^{\text{val}}$  to denote the subset of  $\chi$  such that  $e_1$  and  $e_2$  are values.

As mentioned in Section 3.3, a *world*  $W$  is an ordered list (written  $\langle w_1, \dots, w_n \rangle$ ) of *islands*. An island  $w$  is a pair of some current *knowledge*  $\eta$  and a *law*  $\mathcal{L}$ . The knowledge  $\eta$  for each island represents the current “state” of the island. It comprises four parts: a *store relation*  $\psi$  (which is a set of tuples of the form  $(k, W, s_1, s_2)$ , where  $k$  is a step index,  $W$  is a world, and  $s_1$ , and  $s_2$  are stores); a *population*  $V$  (which is a set of closed values); and two store typings  $\Sigma_1$  and  $\Sigma_2$ . The domains of  $\Sigma_1$  and  $\Sigma_2$  give us the sets of locations that the island “cares about” (a notion we mentioned in Section 3.2). Meanwhile, a law  $\mathcal{L}$  is a set of pairs  $(k, \eta)$ . If  $(k, \eta) \in \mathcal{L}$ , it means that, at “time”  $k$  (representing the number of steps left on the clock), the knowledge  $\eta$  represents an acceptable state for the island to be in. Below we summarize our notation for ease of reference.

<i>Type Interpretation</i>	$\chi$	::=	$\{(k, W, e_1, e_2), \dots\}$
<i>Store Relation</i>	$\psi$	::=	$\{(k, W, s_1, s_2), \dots\}$
<i>Population</i>	$V$	::=	$\{v_1, \dots\}$
<i>Knowledge</i>	$\eta$	::=	$(\psi, V, \Sigma_1, \Sigma_2)$
<i>Law</i>	$\mathcal{L}$	::=	$\{(k, \eta), \dots\}$
<i>Island</i>	$w$	::=	$(\eta, \mathcal{L})$
<i>World</i>	$W$	::=	$\langle w_1, \dots, w_n \rangle$

If  $W = \langle w_1, \dots, w_n \rangle$  and  $1 \leq j \leq n$ , we write  $W[j]$  as shorthand for  $w_j$ . If  $w = (\eta_i, \mathcal{L}_i)$  where  $\eta_i = (\psi_i, V_i, \Sigma_{i1}, \Sigma_{i2})$ , we use the following shorthand to extract various elements out of the island  $w$ :

$w.\eta$	$\equiv$	$\eta_i$	$w.V$	$\equiv$	$V_i$
$w.\mathcal{L}$	$\equiv$	$\mathcal{L}_i$	$w.\Sigma_1$	$\equiv$	$\Sigma_{i1}$
$w.\psi$	$\equiv$	$\psi_i$	$w.\Sigma_2$	$\equiv$	$\Sigma_{i2}$

If  $W$  is a world with  $n$  islands, we also use the following shorthand:

$$\begin{aligned} \Sigma_1(W) &\stackrel{\text{def}}{=} \bigcup_{1 \leq j \leq n} W[j].\Sigma_1 \\ \Sigma_2(W) &\stackrel{\text{def}}{=} \bigcup_{1 \leq j \leq n} W[j].\Sigma_2 \end{aligned}$$

We write *Val* for the set of all values, *Store* for the set of all stores (finite maps from locations to values), and *StoreTy* for the set of store typings (finite maps from locations to closed types). We write *Population* for the set of all subsets of *Val*. Finally, we write  $S_1 \# S_2$  to denote that the sets  $S_1$  and  $S_2$  are disjoint.

### 4.1 Well-Founded, Well-Formed Worlds and Relations

Notice that we cannot naïvely construct a set-theoretic model based on the above intentions since the worlds we wish to construct are (effectively) lists of store relations and store relations are themselves parameterized by worlds (as discussed in Section 3.4). If we ignore islands, laws, populations, and store typings for the moment, and simply model worlds as lists of store relations, we are led to the following specification which captures the essence of the problem:

$$\begin{aligned} \text{StoreRel} &= \mathcal{P}(\mathbb{N} \times \text{World} \times \text{Store} \times \text{Store}) \\ \text{World} &= \text{StoreRel}^n \end{aligned}$$

A simple diagonalization argument shows that the set *StoreRel* has an inconsistent cardinality (*i.e.*, it is an ill-founded recursive definition).

We eliminate the inconsistency by stratifying our definition via the step index. To do so, we first construct *candidate* sets, which are well-founded sets of our intended form. We then construct proper notions of worlds, islands, laws, store relations, and so on, by filtering the candidate sets through some additional well-foundedness constraints.

Figure 3 (top left) defines our candidate sets by induction on  $k$ . First, note that elements of *CandAtom* $_k$  and *CandStoreAtom* $_k$  are tuples with step index  $j$  strictly less than  $k$ . Hence, our candidate sets are well-defined at all steps. Next, note that elements of *CandLawAtom* $_k$  are tuples with step index  $j \leq k$ . Informally, this is because a  $k$ -level law should be able to govern the current knowledge (*i.e.*, the knowledge at the present time when we have  $k$

---

$\begin{aligned} \text{CandAtom}_k &\stackrel{\text{def}}{=} \{(j, W, e_1, e_2) \mid j < k \wedge W \in \text{CandWorld}_j\} \\ \text{CandType}_k &\stackrel{\text{def}}{=} \mathcal{P}(\text{CandAtom}_k^{\text{val}}) \\ \text{CandStoreAtom}_k &\stackrel{\text{def}}{=} \{(j, W, s_1, s_2) \mid j < k \wedge W \in \text{CandWorld}_j\} \\ \text{CandStoreRel}_k &\stackrel{\text{def}}{=} \mathcal{P}(\text{CandStoreAtom}_k) \\ \text{CandKnowledge}_k &\stackrel{\text{def}}{=} \text{CandStoreRel}_k \times \text{Population} \\ &\quad \times \text{StoreTy} \times \text{StoreTy} \\ \text{CandLawAtom}_k &\stackrel{\text{def}}{=} \{(j, \eta) \mid j \leq k \wedge \eta \in \text{CandKnowledge}_j\} \\ \text{CandLaw}_k &\stackrel{\text{def}}{=} \mathcal{P}(\text{CandLawAtom}_k) \\ \text{CandIsland}_k &\stackrel{\text{def}}{=} \text{CandKnowledge}_k \times \text{CandLaw}_k \\ \text{CandWorld}_k &\stackrel{\text{def}}{=} \{W \in (\text{CandIsland}_k)^n \mid n \geq 0\} \end{aligned}$	$\begin{aligned} \text{CandAtom}_\omega &\stackrel{\text{def}}{=} \bigcup_{k \geq 0} \text{CandAtom}_k \\ \text{CandType}_\omega &\stackrel{\text{def}}{=} \mathcal{P}(\text{CandAtom}_\omega^{\text{val}}) \supseteq \bigcup_{k \geq 0} \text{CandType}_k \\ \\ [\chi]_k &\stackrel{\text{def}}{=} \{(j, W, e_1, e_2) \mid j < k \wedge (j, W, e_1, e_2) \in \chi\} \\ [\psi]_k &\stackrel{\text{def}}{=} \{(j, W, s_1, s_2) \mid j < k \wedge (j, W, s_1, s_2) \in \psi\} \\ [\eta]_k &\stackrel{\text{def}}{=} ([\psi]_k, V, \Sigma_1, \Sigma_2) \quad \text{where } \eta = (\psi, V, \Sigma_1, \Sigma_2) \\ [\mathcal{L}]_k &\stackrel{\text{def}}{=} \{(j, \eta) \mid j \leq k \wedge (j, \eta) \in \mathcal{L}\} \\ [w]_k &\stackrel{\text{def}}{=} ([\eta]_k, [\mathcal{L}]_k) \quad \text{where } w = (\eta, \mathcal{L}) \\ [W]_k &\stackrel{\text{def}}{=} \langle [w_1]_k, \dots, [w_n]_k \rangle \quad \text{where } W = \langle w_1, \dots, w_n \rangle \end{aligned}$
---	---

---

$\begin{aligned} (\psi', V', \Sigma'_1, \Sigma'_2) \sqsupseteq (\psi, V, \Sigma_1, \Sigma_2) &\stackrel{\text{def}}{=} V' \supseteq V \wedge \Sigma'_1 \supseteq \Sigma_1 \wedge \Sigma'_2 \supseteq \Sigma_2 \\ (\eta', \mathcal{L}') \sqsupseteq (\eta, \mathcal{L}) &\stackrel{\text{def}}{=} \eta' \supseteq \eta \wedge \mathcal{L}' = \mathcal{L} \\ \langle w'_1, \dots, w'_{n+m} \rangle \sqsupseteq \langle w_1, \dots, w_n \rangle &\stackrel{\text{def}}{=} m \geq 0 \wedge \forall i \in \{1, \dots, n\}. w'_i \sqsupseteq w_i \\ (j, W') \sqsupseteq (k, W) &\stackrel{\text{def}}{=} j \leq k \wedge W' \sqsupseteq [W]_j \wedge W' \in \text{World}_j \wedge W \in \text{World}_k \\ (j, W') \sqsupseteq (k, W) &\stackrel{\text{def}}{=} j < k \wedge (j, W') \sqsupseteq (k, W) \end{aligned}$	$\begin{aligned} \text{Atom}[\tau_1, \tau_2]_k &\stackrel{\text{def}}{=} \{(j, W, e_1, e_2) \in \text{CandAtom}_k \mid W \in \text{World}_j \wedge \Sigma_1(W) \vdash e_1 : \tau_1 \wedge \Sigma_2(W) \vdash e_2 : \tau_2\} \\ \text{Type}[\tau_1, \tau_2]_k &\stackrel{\text{def}}{=} \{\chi \in \mathcal{P}(\text{Atom}[\tau_1, \tau_2]_k^{\text{val}}) \mid \forall (j, W, v_1, v_2) \in \chi. \forall (j', W') \sqsupseteq (j, W). (j', W', v_1, v_2) \in \chi\} \\ \text{StoreAtom}_k &\stackrel{\text{def}}{=} \{(j, W, s_1, s_2) \in \text{CandStoreAtom}_k \mid W \in \text{World}_j\} \\ \text{StoreRel}_k &\stackrel{\text{def}}{=} \{\psi \in \mathcal{P}(\text{StoreAtom}_k) \mid \forall (j, W, s_1, s_2) \in \psi. \forall (i, W') \sqsupseteq (j, W). (i, W', s_1, s_2) \in \psi\} \\ \text{Knowledge}_k &\stackrel{\text{def}}{=} \{(\psi, V, \Sigma_1, \Sigma_2) \in \text{CandKnowledge}_k \mid \psi \in \text{StoreRel}_k \wedge \\ &\quad \forall s_1, s_2, s'_1, s'_2. \\ &\quad (\forall l \in \text{dom}(\Sigma_1). s_1(l) = s'_1(l) \wedge \forall l \in \text{dom}(\Sigma_2). s_2(l) = s'_2(l)) \implies \\ &\quad \forall j, W. (j, W, s_1, s_2) \in \psi \iff (j, W, s'_1, s'_2) \in \psi\} \\ \text{LawAtom}_k &\stackrel{\text{def}}{=} \{(j, \eta) \in \text{CandLawAtom}_k \mid \eta \in \text{Knowledge}_j\} \\ \text{Law}_k &\stackrel{\text{def}}{=} \{\mathcal{L} \in \mathcal{P}(\text{LawAtom}_k) \mid \forall (j, \eta) \in \mathcal{L}. \forall i < j. (i, [\eta]_i) \in \mathcal{L}\} \\ \text{Island}_k &\stackrel{\text{def}}{=} \{(\eta, \mathcal{L}) \in \text{Knowledge}_k \times \text{Law}_k \mid (k, \eta) \in \mathcal{L}\} \\ \text{World}_k &\stackrel{\text{def}}{=} \{W \in (\text{Island}_k)^n \mid n \geq 0 \wedge \\ &\quad \forall a, b \in \{1, \dots, n\}. a \neq b \implies \\ &\quad \text{dom}(W[a].\Sigma_1) \# \text{dom}(W[b].\Sigma_1) \wedge \text{dom}(W[a].\Sigma_2) \# \text{dom}(W[b].\Sigma_2)\} \\ \text{Type}[\tau_1, \tau_2] &\stackrel{\text{def}}{=} \{\chi \in \text{CandType}_\omega \mid \forall k \geq 0. [\chi]_k \in \text{Type}[\tau_1, \tau_2]_k\} \supseteq \bigcup_{k \geq 0} \text{Type}[\tau_1, \tau_2]_k \end{aligned}$
---	---

**Figure 3.** Auxiliary Definitions: Candidate Sets,  $k$ -Approximation, World Extension, and Well-Formedness Conditions

steps left to execute), not just the knowledge in the future when we have strictly fewer steps left.

While our candidate sets establish the existence of sets of our intended form, our worlds and type relations will need to be well-behaved in other ways. There are key constraints associated with atoms, types, store relations, knowledge, laws, islands, and worlds that will be enforced in our final definitions. To specify these constraints we need some additional functions and predicates.

For any set  $\chi$  and any set  $\psi$ , we define the  $k$ -approximation of the set (written  $[\chi]_k$  and  $[\psi]_k$ , respectively) as the subset of its elements whose indices are *strictly less* than  $k$  (see Figure 3, top right). Meanwhile, for any set  $\mathcal{L}$ , we define the  $k$ -approximation of the set (written  $[\mathcal{L}]_k$ ) as the subset of its elements whose indices are *less than or equal to*  $k$ . We extend these  $k$ -approximation notions to knowledge  $\eta$ , islands  $w$ , and worlds  $W$  (written  $[\eta]_k$ ,  $[w]_k$ , and  $[W]_k$ , respectively) by applying  $k$ -approximation to their constituent parts. Note that each of the  $k$ -approximation functions yield elements of  $\text{Cand}X_k$  where  $X$  denotes the appropriate semantic object.

Next, we define the notion of *world extension* (see Figure 3, middle). We write  $(j, W') \sqsupseteq (k, W)$  (where  $\sqsupseteq$  is pronounced “extends”) if  $W$  is a world that is good for  $k$  steps (*i.e.*,  $W \in \text{World}_k$ , see below),  $W'$  is a good world for  $j \leq k$  steps ( $W' \in \text{World}_j$ ), and  $W'$  extends  $[W]_j$  (written  $W' \sqsupseteq [W]_j$ ). Recall from Section 3.3 that future worlds accessible from  $W$  may have new islands added to the end of the list. Furthermore, for each island  $w \in [W]_j$ , the island  $w'$  in the same position in  $W'$  must extend  $w$ . Here we require that  $w'.\mathcal{L} = w.\mathcal{L}$  since an island’s law cannot be amended in future worlds (see Section 3.3). We also require that  $w'.\eta \sqsupseteq w.\eta$ , which says that the island’s population may grow ( $w'.V \supseteq w.V$ ), as may the sets of locations that the island cares about ( $w'.\Sigma_1 \supseteq w.\Sigma_1$  and  $w'.\Sigma_2 \supseteq w.\Sigma_2$ ). Though it may seem from the definition of knowledge extension in Figure 3 that we do not impose any constraints on  $w'.\psi$ , this is not the case. As explained in Section 3.3, an island’s law should govern what the island’s future store relations, populations, and locations of concern may look like. The requirement  $W' \in \text{World}_j$  (which we discuss below) ensures that the future knowledge  $w'.\eta$  obeys the law  $w'.\mathcal{L}$ .

Figure 3 (bottom) defines our various semantic objects, again by induction on  $k$ . These definitions serve to filter their corresponding candidate sets. We proceed now to discuss each of these filtering constraints.

Following Pitts [17], our model is built from syntactically well-typed terms. Thus, we define  $Atom[\tau_1, \tau_2]_k$  as the set of tuples  $(j, W, e_1, e_2)$  where  $\Sigma_1(W) \vdash e_1 : \tau_1$  and  $\Sigma_2(W) \vdash e_2 : \tau_2$ , and  $j < k$ . (Recall that  $\Sigma_1(W)$  denotes the “global” store typing—*i.e.*, the union of the  $\Sigma_1$  components of all the islands in  $W$ .) We also require the world  $W$  to be a member of  $World_j$ .

We define  $Type[\tau_1, \tau_2]_k$  as those sets  $\chi \subseteq Atom[\tau_1, \tau_2]_k^{\text{val}}$  that are closed under world extension. Informally, if  $v_1$  and  $v_2$  are related for  $k$  steps in world  $W$ , then  $v_1$  and  $v_2$  should also be related for  $j$  steps in any future world  $W'$  such that  $(j, W')$  is accessible from (*i.e.*, extends)  $(k, W)$ . We define  $StoreRel_k$  as the set of all  $\psi \subseteq StoreAtom_k \subseteq CandStoreAtom_k$  that are closed under world extension. Again, the intuition for this property is that if two stores  $s_1$  and  $s_2$  are related for  $k$  steps in world  $W$ , then they should continue to be related in all future worlds. This makes sense since stores are essentially tuples of values.

$Knowledge_k$  is the set of all tuples of the form  $(\psi, V, \Sigma_1, \Sigma_2) \in CandKnowledge_k$  such that  $\psi \in StoreRel_k$ . As mentioned above, the domains of  $\Sigma_1$  and  $\Sigma_2$  contain the locations that an island cares about. What this means is that when determining whether two stores  $s_1$  and  $s_2$  belong to the store relation  $\psi$ , we cannot depend upon the contents of any location  $l$  in store  $s_1$  that is not in  $\text{dom}(\Sigma_1)$  or on the contents of any location in  $s_2$  that is not in  $\text{dom}(\Sigma_2)$ . Thus,  $\Sigma_1$  and  $\Sigma_2$  essentially serve as *accessibility maps* [8]. While Benton and Leperchey’s accessibility maps are functions from stores to subsets of  $Loc$ , our accessibility maps are sets of locations that are allowed to grow over time.

We define  $Law_k$  as the set of  $\mathcal{L}$  such that, for all  $(j, \eta) \in \mathcal{L}$ , we have that  $\eta \in Knowledge_j$ . Furthermore, we require that the sets  $\mathcal{L}$  be closed under decreasing step index—that is, if some knowledge  $\eta$  obeys law  $\mathcal{L}$  for  $j$  steps, then it must be the case that at any future time, when we have  $i < j$  steps left, the  $i$ -approximation of knowledge  $\eta$  still obeys the law  $\mathcal{L}$ .

$Island_k$  is the set of all pairs  $(\eta, \mathcal{L}) \in (Knowledge_k \times Law_k)$  such that the knowledge  $\eta$  obeys the law  $\mathcal{L}$  at the current time denoted by step index  $k$ —*i.e.*,  $(k, \eta) \in \mathcal{L}$ .

Finally, we define  $World_k$  as the set of all  $W \in (Island_k)^n$ . We also require that the sets of locations that each island  $W[a]$  cares about are disjoint from the sets of locations that any other island  $W[b]$  cares about, thus ensuring separation of islands.

## 4.2 Relational Interpretations of Types

Figure 4 (top) gives the definition of our logical relations for  $F^{\mu!}$ . The relations  $\mathcal{V}_n[\tau] \rho$  are defined by induction on  $n$  and nested induction on the type  $\tau$ . We use the metavariable  $\rho$  to denote type substitutions. A type substitution  $\rho$  is a finite map from type variables  $\alpha$  to triples  $(\chi, \tau_1, \tau_2)$  where  $\tau_1$  and  $\tau_2$  are closed types, and  $\chi$  is a relational interpretation in  $Type[\tau_1, \tau_2]$ . If  $\rho(\alpha) = (\chi, \tau_1, \tau_2)$ , then  $\rho_1(\alpha)$  denotes  $\tau_1$  and  $\rho_2(\alpha)$  denotes  $\tau_2$ .

Note that, by the definition of  $\mathcal{V}_n[\tau] \rho$ , if  $(k, W, v_1, v_2) \in \mathcal{V}_n[\tau] \rho$ , then  $k < n$ ,  $W \in World_k$ ,  $\Sigma_1(W) \vdash v_1 : \rho_1(\tau)$ , and  $\Sigma_2(W) \vdash v_2 : \rho_2(\tau)$ . Most of the relations  $\bar{\mathcal{V}}_n[\tau] \rho$  are straightforward. For instance, the logical relation at type  $\text{int}$  says that two integers are logically related for any number of steps  $k$  and in any world  $W$  as long as they are equal. The relations for the other base types  $\text{unit}$  and  $\text{bool}$  are similar.

Two functions  $\lambda x : \rho_1(\tau). e_1$  and  $\lambda x : \rho_2(\tau). e_2$  are related for  $k$  steps in world  $W$  at the type  $\tau \rightarrow \tau'$  if in any future world  $W'$  where there are  $j < k$  steps left to execute (since beta-reduction consumes a step), and we have arguments  $v_1, v_2$  that are related for  $j$  steps in world  $W'$  at the argument type  $\tau$ , then the terms  $[v_1/x]e_1$

$$\begin{aligned}
s_1, s_2 :_k W &\stackrel{\text{def}}{=} \vdash s_1 : \Sigma_1(W) \wedge \vdash s_2 : \Sigma_2(W) \wedge \\
&\quad \forall w \in W. \forall j < k. (j, [W]_j, s_1, s_2) \in w.\psi \\
\mathcal{V}_n[\tau] \rho &= \bar{\mathcal{V}}_n[\tau] \rho \cap Atom[\rho_1(\tau), \rho_2(\tau)]_n^{\text{val}} \\
\bar{\mathcal{V}}_n[\alpha] \rho &= \chi \quad \text{where } \rho(\alpha) = (\chi, \tau_1, \tau_2) \\
\bar{\mathcal{V}}_n[\text{unit}] \rho &= \{(k, W, (), ())\} \\
\bar{\mathcal{V}}_n[\text{int}] \rho &= \{(k, W, v, v) \mid v \in \mathbb{N}\} \\
\bar{\mathcal{V}}_n[\text{bool}] \rho &= \{(k, W, v, v) \mid v = \text{true} \vee v = \text{false}\} \\
\bar{\mathcal{V}}_n[\tau \times \tau'] \rho &= \{(k, W, \langle v_1, v'_1 \rangle, \langle v_2, v'_2 \rangle) \mid \\
&\quad (k, W, v_1, v_2) \in \mathcal{V}_n[\tau] \rho \wedge \\
&\quad (k, W, v'_1, v'_2) \in \mathcal{V}_n[\tau'] \rho\} \\
\bar{\mathcal{V}}_n[\tau \rightarrow \tau'] \rho &= \{(k, W, \lambda x : \rho_1(\tau). e_1, \lambda x : \rho_2(\tau). e_2) \mid \\
&\quad \forall (j, W') \sqsupset (k, W). \forall v_1, v_2. \\
&\quad (j, W', v_1, v_2) \in \mathcal{V}_n[\tau] \rho \implies \\
&\quad (j, W', [v_1/x]e_1, [v_2/x]e_2) \in \mathcal{E}_n[\tau'] \rho\} \\
\bar{\mathcal{V}}_n[\forall \alpha. \tau] \rho &= \{(k, W, \Lambda \alpha. e_1, \Lambda \alpha. e_2) \mid \\
&\quad \forall (j, W') \sqsupset (k, W). \forall \tau_1, \tau_2, \chi \in Type[\tau_1, \tau_2]. \\
&\quad (j, W', [\tau_1/\alpha]e_1, \\
&\quad [\tau_2/\alpha]e_2) \in \mathcal{E}_n[\tau] \rho[\alpha \mapsto (\chi, \tau_1, \tau_2)]\} \\
\bar{\mathcal{V}}_n[\exists \alpha. \tau] \rho &= \{(k, W, \text{pack } \tau_1, v_1 \text{ as } \exists \alpha. \rho_1(\tau), \\
&\quad \text{pack } \tau_2, v_2 \text{ as } \exists \alpha. \rho_2(\tau)) \mid \\
&\quad \exists \chi \in Type[\tau_1, \tau_2]. \\
&\quad (k, W, v_1, v_2) \in \mathcal{V}_n[\tau] \rho[\alpha \mapsto (\chi, \tau_1, \tau_2)]\} \\
\bar{\mathcal{V}}_n[\mu \alpha. \tau] \rho &= \{(k, W, \text{fold } v_1, \text{fold } v_2) \mid k < n \wedge \\
&\quad \forall j < k. (j, [W]_j, v_1, v_2) \in \mathcal{V}_k[\mu \alpha. \tau/\alpha] \tau] \rho\} \\
\bar{\mathcal{V}}_n[\text{ref } \tau] \rho &= \{(k, W, l_1, l_2) \mid k < n \wedge w_{\text{ref}}(k, \rho, \tau, l_1, l_2) \in W\} \\
w_{\text{ref}}(k, \rho, \tau, l_1, l_2) &= (\eta, \mathcal{L}) \\
&\quad \text{where } \eta = (\psi, \{\}, \{l_1 : \rho_1(\tau)\}, \{l_2 : \rho_2(\tau)\}) \\
&\quad \psi = \{(j, W', s_1, s_2) \mid (j, W', s_1(l_1), s_2(l_2)) \in \mathcal{V}_k[\tau] \rho\} \\
&\quad \mathcal{L} = \{(j, [\eta]_j) \mid j \leq k\} \\
\mathcal{E}_n[\tau] \rho &= \{(k, W, e_1, e_2) \in Atom[\rho_1(\tau), \rho_2(\tau)]_n \mid \\
&\quad \forall j < k. \forall s_1, s_2, s'_1, v_1. \\
&\quad s_1, e_1 \xrightarrow{j} s'_1, v_1 \wedge s_1, s_2 :_k W \implies \\
&\quad \exists s_2, v_2, W'. (k - j, W') \sqsupset (k, W) \wedge \\
&\quad s_2, e_2 \xrightarrow{*} s'_2, v_2 \wedge s'_1, s'_2 :_{k-j} W' \wedge \\
&\quad (k - j, W', v_1, v_2) \in \mathcal{V}_n[\tau] \rho\}
\end{aligned}$$

$$\begin{aligned}
\mathcal{V}[\tau] \rho &= \bigcup_{n \geq 0} \mathcal{V}_n[\tau] \rho & \mathcal{E}[\tau] \rho &= \bigcup_{n \geq 0} \mathcal{E}_n[\tau] \rho \\
\mathcal{D}[\cdot] &= \{\emptyset\} \\
\mathcal{D}[\Delta, \alpha] &= \{\rho[\alpha \mapsto (\chi, \tau_1, \tau_2)] \mid \rho \in \mathcal{D}[\Delta] \wedge \chi \in Type[\tau_1, \tau_2]\} \\
\mathcal{G}[\cdot] \rho &= \{(k, W, \emptyset) \mid W \in World_k\} \\
\mathcal{G}[\Gamma, x : \tau] \rho &= \{(k, W, \gamma[x \mapsto (v_1, v_2)]) \mid \\
&\quad (k, W, \gamma) \in \mathcal{G}[\Gamma] \rho \wedge (k, W, v_1, v_2) \in \mathcal{V}[\tau] \rho\} \\
\mathcal{S}[\Sigma] &= \{(k, W) \mid \forall (l : \tau) \in \Sigma. (k, W, l, l) \in \mathcal{V}[\text{ref } \tau] \emptyset\} \\
\Delta; \Gamma; \Sigma \vdash e_1 \leq^{\text{log}} e_2 : \tau &\stackrel{\text{def}}{=} \Delta; \Gamma; \Sigma \vdash e_1 : \tau \wedge \Delta; \Gamma; \Sigma \vdash e_2 : \tau \wedge \\
&\quad \forall k \geq 0. \forall \rho, \gamma, W. \rho \in \mathcal{D}[\Delta] \wedge (k, W, \gamma) \in \mathcal{G}[\Gamma] \rho \wedge \\
&\quad (k, W) \in \mathcal{S}[\Sigma] \implies \\
&\quad (k, W, \rho_1(\gamma_1(e_1)), \rho_2(\gamma_2(e_2))) \in \mathcal{E}[\tau] \rho \\
\Delta; \Gamma; \Sigma \vdash e_1 \approx^{\text{log}} e_2 : \tau &\stackrel{\text{def}}{=} \Delta; \Gamma; \Sigma \vdash e_1 \leq^{\text{log}} e_2 : \tau \wedge \\
&\quad \Delta; \Gamma; \Sigma \vdash e_2 \leq^{\text{log}} e_1 : \tau
\end{aligned}$$

Figure 4. Step-Indexed Logical Relations for  $F^{\mu!}$

and  $[v_2/x]e_2$  are logically related for  $j$  steps in world  $W'$  at the type  $\tau'$  (i.e., they are in the relation  $\mathcal{E}_n \llbracket \tau' \rrbracket \rho$ ). Parameterizing over an arbitrary future world  $W'$  is necessary here in order to ensure closure of the logical relation under world extension.

The cases of the logical relation for  $\forall \alpha. \tau$  and  $\exists \alpha. \tau$  are essentially standard. The former involves *parameterizing* over an arbitrary relational interpretation  $\chi$  of  $\alpha$ , and the latter involves *choosing* an arbitrary relational interpretation  $\chi$  of  $\alpha$ . The way the worlds are manipulated follows in the style of the other rules. The logical relation for  $\mu \alpha. \tau$  is very similar to previous step-indexed accounts of recursive types, as described in Section 3.5. (Note that, although the type gets larger on the r.h.s. of the definition, the step index gets smaller, so the definition is well-founded.)

We say that two stores  $s_1$  and  $s_2$  are related for  $k$  steps at the world  $W$  (see Figure 4, top) if the stores are well-typed with respect to the store typings  $\Sigma_1(W)$  and  $\Sigma_2(W)$ , respectively, and if the stores are considered acceptable by—i.e., they are in the store relations of—all the islands in  $W$  at all future times when  $j < k$ .

Any two locations related at a type  $\text{ref } \tau$  are publicly accessible references. For reasoning about such *visible* locations, existing logical relations methods usually employ some mechanism that is distinct from the machinery used to reason about local or *hidden* state. Since there always exists a bijection between the visible locations of the two computations, the mechanism usually involves having a special portion of the world that tracks the bijection between visible locations as well as the type  $\tau$  of their contents. Unlike previous methods, our worlds have no specialized machinery for reasoning about visible locations. Our technique for modeling (publicly accessible) references is simply a mode of use of our mechanism for reasoning about local state.

Intuitively, two locations  $l_1$  and  $l_2$  should be related at the type  $\text{ref } \tau$  in world  $W$  for  $k$  steps if, given any two stores  $s_1$  and  $s_2$  that are related for  $k$  steps at world  $W$ , the contents of these locations, i.e.,  $s_1(l_1)$  and  $s_2(l_2)$ , are related for  $k - 1$  steps at the type  $\tau$ . To enforce this requirement, we simply install a special island  $w_{\text{ref}}$  that only cares about the one location  $l_1$  in  $s_1$  and the one location  $l_2$  in  $s_2$ . Furthermore,  $w_{\text{ref}}$  has an empty population and a law that says the population should remain empty in future worlds. Finally, the island's *fixed* store relation  $\psi$  relates all stores  $s_1$  and  $s_2$  whose contents at locations  $l_1$  and  $l_2$ , respectively, are related at type  $\tau$  for  $j < k$  steps. Here  $j < k$  suffices since pointer dereferencing consumes a step (see Section 3.5).

The relation  $\mathcal{E}_n \llbracket \tau \rrbracket \rho$  defines when two *terms* are logically related. Two well-typed terms  $e_1$  and  $e_2$  are related for  $k$  steps at the type  $\tau$  in world  $W$  if, given two starting stores  $s_1$  and  $s_2$  that are related for  $k$  steps at world  $W$ , if  $s_1, e_1$  steps to a configuration  $s'_1, v_1$  in  $j < k$  steps then the following conditions hold. First,  $s_2, e_2$  must step to some  $s'_2, v_2$  in any number of steps. Second, there must exist a world  $W' \in \text{World}_{k-j}$  that extends the world  $W$ . Third, the final stores  $s'_1$  and  $s'_2$  must be related for the remaining  $k - j$  steps at world  $W'$ . Fourth, the values  $v_1$  and  $v_2$  must be related for  $k - j$  steps in the world  $W'$  at the type  $\tau$ .

The definitions of logical approximation and equivalence for open terms are given at the bottom of Figure 4. Here we state some of the main properties of our logical relation. Further details of the meta-theory will appear in the online technical appendix [4].

**Lemma 4.1 (Logical Relations Are Valid Type Interpretations)**

If  $\Delta \vdash \tau$  and  $\rho \in \mathcal{D} \llbracket \Delta \rrbracket$ , then  $\mathcal{V} \llbracket \tau \rrbracket \rho \in \text{Type}[\rho_1(\tau), \rho_2(\tau)]$ .

**Theorem 4.2 (Fundamental Property)**

If  $\Delta; \Gamma; \Sigma \vdash e : \tau$  then  $\Delta; \Gamma; \Sigma \vdash e \preceq^{\text{log}} e : \tau$ .

**Theorem 4.3 (Soundness w.r.t. Contextual Equivalence)**

If  $\Delta; \Gamma; \Sigma \vdash e_1 \preceq^{\text{log}} e_2 : \tau$  then  $\Delta; \Gamma; \Sigma \vdash e_1 \preceq^{\text{ctx}} e_2 : \tau$ .

## 5. Examples

In this section we present a number of examples demonstrating applications of our method. Our examples do not make use of recursive types (or even recursion), but Ahmed's prior work, which we build on, gives several examples that do [2]. We will walk through the proof for the first example in detail. For the remaining ones, we only sketch the central ideas, mainly by giving suitable island definitions and type interpretations. Full proofs for these examples and others appear in the online technical appendix [4].

### 5.1 Name Generator

Our first example is perhaps the simplest possible state-dependent ADT, a generator for fresh names. Nevertheless, it captures the essence of the `Symbol` example from the introduction:

$$e = \text{let } x = \text{ref } 0 \text{ in} \\ \text{pack int}, \langle \lambda z : \text{unit}. (++)x, \lambda z : \text{int}. (z \leq !x) \rangle \text{ as } \sigma$$

where  $\sigma = \exists \alpha. (\text{unit} \rightarrow \alpha) \times (\alpha \rightarrow \text{bool})$  and  $(++)x$  abbreviates the expression  $(x := !x + 1; !x)$ , and `let` is encoded in the standard way (using function application). The package defines an abstract type  $\alpha$  of names and provides two operations: the first one returns a fresh name on each invocation, and the second one checks that any value of type  $\alpha$  it is given is a “valid” name, i.e., one that was previously generated by the first operation.

Names are represented as integers, and the local counter  $x$  stores the highest value that has been used so far. The intended invariant of this implementation is that no value of type  $\alpha$  ever has a representation that is greater than the current content of  $x$ . Under this invariant, we should be able to prove that the second operation, which dynamically checks this property, always returns `true`.

To prove this, we show that  $e$  is equivalent to a second expression  $e'$ , identical to  $e$ , except that the dynamic check  $(z \leq !x)$  is eliminated and replaced by `true`. We only show the one direction,  $\vdash e \preceq^{\text{log}} e' : \sigma$ . The other direction is proven analogously.

Because the terms are closed, this only requires showing that  $(k_0, W_0, e, e') \in \mathcal{E} \llbracket \sigma \rrbracket \emptyset$  for all  $k_0 \geq 0$  and worlds  $W_0$ . Assume stores  $s_0, s'_0 :_{k_0} W_0$  and the existence of a reduction sequence  $s_0, e \mapsto^{k_1} s_1, v_1$  with  $k_1 < k_0$ . According to the definition of  $\mathcal{E} \llbracket \sigma \rrbracket \emptyset$ , we need to come up with a reduction  $s'_0, e' \mapsto^* s'_1, v'_1$  and a world  $W_1$  such that  $(k_0 - k_1, W_1) \sqsupseteq (k_0, W_0)$  and:

$$s_1, s'_1 :_{k_0 - k_1} W_1 \wedge (k_0 - k_1, W_1, v_1, v'_1) \in \mathcal{V} \llbracket \sigma \rrbracket \emptyset$$

By inspecting the definition of reduction, we see that

$$s_1 = s_0[l \mapsto 0], \quad v_1 = \text{pack int}, \langle \lambda z. (++)l, \lambda z. (z \leq !l) \rangle \text{ as } \sigma$$

for some  $l \notin \text{dom}(s_0)$ . In the same manner,  $s'_0, e'$  obviously can choose some  $l' \notin \text{dom}(s'_0)$  and reduce to:

$$s'_1 = s'_0[l' \mapsto 0], \quad v'_1 = \text{pack int}, \langle \lambda z. (++)l', \lambda z. \text{true} \rangle \text{ as } \sigma$$

We now need to define a suitable island that enables us to show that  $v_1$  and  $v'_1$  are related. We know  $W_0$  has the form  $\langle w_1, \dots, w_p \rangle$  for some  $p$ . Let  $W_1$  be  $\langle w_1, \dots, w_p, w_{p+1} \rangle$  where:

$$\begin{aligned} w_{p+1} &= (\eta_{k_0 - k_1}^0, \mathcal{L}_{k_0 - k_1}) \\ \eta_k^n &= (\psi_k^n, V_n, \{l : \text{int}\}, \{l' : \text{int}\}) \\ \psi_k^n &= \{(j, W, s, s') \in \text{StoreAtom}_k \mid s(l) = s'(l') = n\} \\ V_n &= \{i \mid 1 \leq i \leq n\} \\ \mathcal{L}_k &= \{(j, \eta_j^n) \in \text{LawAtom}_k \mid n \in \mathbb{N}\} \end{aligned}$$

The population  $V_n$  consists of all integers that are “valid” names in a world, i.e., not greater than the current value of  $x$ . We have to show  $(k_0 - k_1, W_1) \sqsupseteq (k_0, W_0)$  and  $s_1, s'_1 :_{k_0 - k_1} W_1$ . Both are straightforward.

By definition of  $\mathcal{V}[\exists\alpha.\tau]$ , we need to continue by giving a relation  $\chi_\alpha \in \text{Type}[\text{int}, \text{int}]$ , such that:

$$(k_0 - k_1, W_1, \langle \lambda z. (++l), \lambda z. (z \leq !l) \rangle, \langle \lambda z. (++l'), \lambda z. \text{true} \rangle) \in \mathcal{V}[\langle (\text{unit} \rightarrow \alpha) \times (\alpha \rightarrow \text{bool}) \rangle] \rho$$

with  $\rho = [\alpha \mapsto (\chi_\alpha, \text{int}, \text{int})]$ . We choose the following one:

$$\chi_\alpha = \{(j, W, i, i) \in \text{Atom}[\text{int}, \text{int}]_{k_0 - k_1} \mid i \in W[p + 1].V\}$$

This interpretation of  $\alpha$  depends on the (valid) assumption that it will only be considered at  $W$ 's that are future worlds of  $W_1$  (in particular, it assumes that the  $(p+1)$ -th island in  $W$ , written  $W[p + 1]$ , is a future version of the  $w_{p+1}$  we defined above). We could build this assumption explicitly into the definition of  $\chi_\alpha$ , but as we will see it is simply not necessary to do so. By virtue of this assumption, a value  $i$  is only a valid inhabitant of type  $\alpha$  in worlds whose  $(p+1)$ -th island population contains  $i$ , that is, where  $!l \geq i$ . Note that the relation is closed under world extension because  $V$  may only grow over time, as explained in Section 3.3.

By definition of  $\mathcal{V}[\tau \times \tau']$ , it remains to be shown that:

1.  $(k_0 - k_1, W_1, \lambda z. (++l), \lambda z. (++l')) \in \mathcal{V}[\langle \text{unit} \rightarrow \alpha \rangle] \rho$
2.  $(k_0 - k_1, W_1, \lambda z. (z \leq !l), \lambda z. \text{true}) \in \mathcal{V}[\langle \alpha \rightarrow \text{bool} \rangle] \rho$

For each of these, we assume we begin in some strictly future world  $W_2$  in which  $(k_2, W_2) \sqsupset (k_0 - k_1, W_1)$  and  $s_2, s'_2 :_{k_2} W_2$ .

First consider (1). We are given  $s_2, (++l) \mapsto^{k_3} s_3, v_3$  for some  $k_3 < k_2$ , and it remains to show that  $s'_2, (++l') \mapsto^* s'_3, v'_3$ , such that  $s_3, s'_3$  and  $v_3, v'_3$  are related in some future world  $W_3$  such that  $(k_2 - k_3, W_3) \sqsupseteq (k_2, W_2)$ .

From  $(k_2, W_2) \sqsupseteq (k_0 - k_1, W_1)$  we know that  $W_2[p + 1].\mathcal{L} = [W_1[p + 1].\mathcal{L}]_{k_2} = \mathcal{L}_{k_2}$ . From that  $(k_2, W_2[p + 1].\eta) \in \mathcal{L}_{k_2}$  follows, and hence there exists  $n$  such that  $W_2[p + 1].\eta = \eta_{k_2}^n$ . That is,  $W_2[p + 1].\psi = \psi_{k_2}^n$  and  $W_2[p + 1].V = V_n$ . From  $s_2, s'_2 :_{k_2} W_2$  and  $k_3 < k_2$  we can conclude  $(k_3, [W_2]_{k_3}, s_2, s'_2) \in \psi_{k_2}^n$  and thus  $s_2(l) = s'_2(l') = n$ . Consequently,  $v_3 = v'_3 = n + 1$ ,  $s_3 = s_2[l \mapsto n + 1]$ , and  $s'_3 = s'_2[l' \mapsto n + 1]$ .

Now we choose  $W_3$  to be  $W_2$  with its  $(p + 1)$ -th island updated to  $(\eta_{k_2 - k_3}^{n+1}, \mathcal{L}_{k_2 - k_3})$ . Again, we have to check the relevant properties,  $(k_2 - k_3, W_3) \sqsupseteq (k_2, W_2)$  and  $s_3, s'_3 :_{k_2 - k_3} W_3$ , which are straightforward. Last, we have to show that the results  $v_3, v'_3$  are related in  $\mathcal{V}[\langle \alpha \rangle] \rho$  under this world, i.e.,  $(k_2 - k_3, W_3, n + 1, n + 1) \in \chi_\alpha$ . Since  $(n + 1, n + 1) \in V_{n+1} = W_3[p + 1].V$ , this is immediate from the definition of  $\chi_\alpha$ .

Now consider (2). The proof is similar to that for part (1), but simpler. We are given that  $(k_2, W_2, v_2, v'_2) \in \mathcal{V}[\langle \alpha \rangle] \rho = \chi_\alpha$ , and  $s_2, (v_2 \leq !l) \mapsto^{k_3} s_3, v_3$  for some  $k_3 < k_2$ . The main thing to show is that  $v_3 = \text{true}$  (we can pick the end world  $W_3$  to just be  $W_2$ ). As in part (1), we can reason that  $W_2[p + 1].\eta = \eta_{k_2}^n$  for some  $n$ , and therefore that  $s_2(l) = n$  and, by definition of  $\chi_\alpha$ , also that  $v_2 \leq n$ . Hence,  $v_2 \leq s_2(l)$ , and the desired result follows easily.

## 5.2 Using ref As a Name Generator

An alternative way to implement a name generator is to represent names by locations and rely on generativity of the `ref` operator.

$$e = \text{pack ref unit}, \langle \lambda z : \text{unit}. (\text{ref } ()) \rangle, \lambda p : (\text{ref unit} \times \text{ref unit}). (\text{fst } p == \text{snd } p) \text{ as } \sigma$$

where  $\sigma = \exists\alpha. (\text{unit} \rightarrow \alpha) \times (\alpha \times \alpha \rightarrow \text{bool})$ . Here, the second function implements a proper equality operator on names. We want to prove this implementation contextually equivalent to one using integers, as in the previous example:

$$e' = \text{let } x = \text{ref } 0 \text{ in} \\ \text{pack int}, \langle \lambda z : \text{unit}. (++x) \rangle, \lambda p : (\text{int} \times \text{int}). (\text{fst } p = \text{snd } p) \text{ as } \sigma$$

Here are a suitable island definition and type interpretation for  $\alpha$ :

$$\begin{aligned} w_{p+1} &= (\eta_{k_0}^{\langle \rangle}, \mathcal{L}_{k_0}) \\ \eta_k^{\langle l_1, \dots, l_n \rangle} &= (\psi_k^n, V_{\langle l_1, \dots, l_n \rangle}, \{\{l_i : \text{unit} \mid 1 \leq i \leq n\}, \{l' : \text{int}\}\}) \\ \psi_k^n &= \{(j, W, s, s') \in \text{StoreAtom}_k \mid s'(l') = n\} \\ V_{\langle l_1, \dots, l_n \rangle} &= \{\langle l_i, i \rangle \mid 1 \leq i \leq n\} \\ \mathcal{L}_k &= \{(j, \eta_j^{\langle l_1, \dots, l_n \rangle}) \in \text{LawAtom}_k \mid n \in \mathbb{N}\} \\ \chi_\alpha &= \{(j, W, l, i) \in \text{Atom}[\text{ref unit}, \text{int}]_{k_0} \mid \\ &\quad \langle l, i \rangle \in W[p + 1].V\} \end{aligned}$$

Here, and in the examples that follow,  $k_0$  represents the current step level, and  $p$  the number of islands in the current world  $W_0$ , at the point in the proof where we extend  $W_0$  with the island  $w_{p+1}$  governing the example's local state. In this example, we assume that all labels in a list  $\langle l_1, \dots, l_n \rangle$  are pairwise disjoint, and  $l'$  is a distinguished label, namely the one that has been allocated for  $x$  (as in the previous example).

In the definitions above, the population not only records the valid names for  $e'$  (as in Section 5.1), but also relates them to the locations allocated by  $e$ . The latter are not guessable ahead of time, due to nondeterminism of memory allocation, but the law  $\mathcal{L}_k$  is flexible enough to permit any partial bijection between  $\{1, \dots, n\}$  and  $\text{Loc}$  to evolve over time. We (ab)use term-level pairs  $\langle l, i \rangle$  to encode this partial bijection in  $V$ . This is sufficient to deduce  $i = j$  iff  $l_i = l_j$  when proving equivalence of the equality operators.

## 5.3 Twin Abstraction

Another interesting variation on the generator theme involves the definition of *two* abstract types (we write `pack`  $\tau_1, \tau_2, e$  as  $\exists\alpha, \beta. \tau$  to abbreviate two nested existentials in the obvious way):

$$\begin{aligned} e &= \text{let } x = \text{ref } 0 \text{ in} \\ &\quad \text{pack int, int}, \langle \lambda z : \text{unit}. (++x), \\ &\quad \quad \lambda z : \text{unit}. (++x) \rangle, \\ &\quad \lambda p : (\text{int} \times \text{int}). (\text{fst } p = \text{snd } p) \text{ as } \sigma \end{aligned}$$

where  $\sigma = \exists\alpha, \beta. (\text{unit} \rightarrow \alpha) \times (\text{unit} \rightarrow \beta) \times (\alpha \times \beta \rightarrow \text{bool})$ . Here we use a single counter to generate names of two types,  $\alpha$  and  $\beta$ , and a comparison operator that takes as input names of *different* type. Because both types share the same counter, it appears impossible for a name to belong to both types (either it was generated as a name of type  $\alpha$  or of type  $\beta$  but not of both). The example is interesting, however, in that we have no way of knowing the interpretations of  $\alpha$  and  $\beta$  ahead of time, since calls to the name generation functions can happen in arbitrary combinations. We can verify our intuition by proving that  $e$  is equivalent to an  $e'$  where the comparison operator is replaced by  $\lambda p : (\text{int} \times \text{int}). \text{false}$ .

The following  $w$  and  $\chi$  definitions enable such a proof:

$$\begin{aligned} w_{p+1} &= (\eta_{k_0}^{0,0}, \mathcal{L}_{k_0}) \\ \eta_k^{n,S} &= (\psi_k^n, V_{n,S}, \{\{l : \text{int}\}, \{l' : \text{int}\}\}) \\ \psi_k^n &= \{(j, W, s, s') \in \text{StoreAtom}_k \mid s(l) = s'(l') = n\} \\ V_{n,S} &= \{\langle 1, i \rangle \mid i \in S\} \cup \{\langle 2, i \rangle \mid i \in \{1, \dots, n\} \setminus S\} \\ \mathcal{L}_k &= \{(j, \eta_j^{n,S}) \in \text{LawAtom}_k \mid S \subseteq \{1, \dots, n\}\} \\ \chi_\alpha &= \{(j, W, i, i) \in \text{Atom}[\text{int}, \text{int}]_{k_0} \mid \langle 1, i \rangle \in W[p + 1].V\} \\ \chi_\beta &= \{(j, W, i, i) \in \text{Atom}[\text{int}, \text{int}]_{k_0} \mid \langle 2, i \rangle \in W[p + 1].V\} \end{aligned}$$

The population here is partitioned into the valid names for  $\alpha$  and the valid names for  $\beta$ , basically recording the history of calls to the two generator functions. To encode such a disjoint union in  $V$ , each value is wrapped in a pair with the first component marking the type it belongs to (1 for  $\alpha$ , 2 for  $\beta$ ). When proving equivalence of the two comparison operators, the definitions of  $\chi_\alpha, \chi_\beta$  and  $W[p + 1].V$  directly imply that the arguments must be from disjoint sets.

## 5.4 Cell Class

The next example is a more richly-typed variation of the higher-order cell object example of Koutavas and Wand [12]:

$$e = \Lambda\alpha. \text{pack } \text{ref } \alpha, \langle \lambda x : \alpha. \text{ref } x, \\ \lambda r : \text{ref } \alpha. !r, \\ \lambda \langle r, x \rangle : \text{ref } \alpha \times \alpha. (r := x) \rangle \text{ as } \sigma$$

where  $\sigma = \exists\beta. (\alpha \rightarrow \beta) \times (\beta \rightarrow \alpha) \times (\beta \times \alpha \rightarrow \text{unit})$ . We use pattern matching notation here merely for clarity and brevity (imagine replacing occurrences of  $r$  and  $x$  in the third function with  $\text{fst}$  and  $\text{snd}$  projections, respectively, of the argument).

This example generalizes Koutavas and Wand's original version in two ways. First, we parameterize over the cell content type  $\alpha$ , which can of course be instantiated with an arbitrary higher type, thus exercising our ability to handle higher-order stored values. Second, instead of just implementing a single object, our example actually models a *class*, where  $\beta$  represents the abstract class type, and the first function acts as a constructor for creating new cell objects. (A subsequent paper by Koutavas and Wand also considers a class-based version of their original example [13], but it is modeled with a Java-like nominal type system, not with existential types.)

Similar to [12], we want to prove this canonical cell implementation equivalent to one using two alternating slots:

$$e' = \Lambda\alpha. \text{pack } (\text{ref int} \times (\text{ref } \alpha \times \text{ref } \alpha)), \\ \langle \lambda x : \alpha. \langle \text{ref } 1, (\text{ref } x, \text{ref } x) \rangle, \\ \lambda \langle r_0, \langle r_1, r_2 \rangle \rangle : (\text{ref int} \times (\text{ref } \alpha \times \text{ref } \alpha)). \\ \text{if } !r_0 = 1 \text{ then } !r_1 \text{ else } !r_2, \\ \lambda \langle \langle r_0, \langle r_1, r_2 \rangle \rangle, x \rangle : (\text{ref int} \times (\text{ref } \alpha \times \text{ref } \alpha)) \times \alpha. \\ \text{if } !r_0 = 1 \text{ then } (r_0 := 2; r_2 := x) \\ \text{else } (r_0 := 1; r_1 := x) \rangle \text{ as } \sigma$$

The island for this proof has to relate the objects created on each side. More precisely, it relates the associated locations, in the form of tuples  $\langle l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle \rangle$ , where  $l$  is the reference used by  $e$ , while the others belong to the respective object from  $e'$ . As in Section 5.2, we abuse term-level tuples to express the relation in the population  $V$ . Assuming we are given a relational interpretation  $\chi_\alpha \in \text{Type}[\tau_\alpha, \tau'_\alpha]$  for the type parameter  $\alpha$ , we define:

$$w_{p+1} = (\eta_{k_0}^\emptyset, \mathcal{L}_{k_0}) \\ \eta_k^V = (\psi_k^V, V, \Sigma_V, \Sigma'_V) \\ \psi_k^V = \{(j, W, s, s') \in \text{StoreAtom}_k \mid \\ \forall \langle l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle \rangle \in V. \exists i \in \{1, 2\}. \\ s'(l'_0) = i \wedge (j, W, s(l), s'(l'_i)) \in \chi_\alpha\} \\ \Sigma_V = \{l : \tau_\alpha \mid \langle l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle \rangle \in V\} \\ \Sigma'_V = \{l'_0 : \text{int}, l'_1 : \tau'_\alpha, l'_2 : \tau'_\alpha \mid \langle l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle \rangle \in V\} \\ \mathcal{L}_k = \{(j, \eta_j^V) \in \text{LawAtom}_k\} \\ \chi_\beta = \{(j, W, l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle) \in \\ \text{Atom}[\text{ref } \tau_\alpha, \text{ref int} \times (\text{ref } \tau'_\alpha \times \text{ref } \tau'_\alpha)]_{k_0} \mid \\ \langle l, \langle l'_0, \langle l'_1, l'_2 \rangle \rangle \rangle \in W[p+1].V\}$$

Here, we assume that  $V$  mentions any location at most once. The store relation  $\psi_k^V$  ensures that, for each pair of objects recorded in  $V$ , the content of its reference  $l$  in  $e$  is related to the proper slot  $l'_1$  or  $l'_2$  in  $e'$ , depending on the current flag value stored in its  $l'_0$ . Note how  $\psi_k^V$  relies crucially on the presence of the world parameter  $W$ .

## 5.5 Callback with Lock

The proofs for the examples presented so far do not use step indices in an interesting way. While it does not actually involve existential quantification, the last of our examples demonstrates an unexpected case where the steps come in handy. Our example is similar to the reentrant callback example of Banerjee and Naumann [7].

Consider the following object encoding of higher-order type  $\tau = ((\text{unit} \rightarrow \text{unit}) \rightarrow \text{unit}) \times (\text{unit} \rightarrow \text{int})$ :

$$e = C[f(); x := !x + 1]$$

where

$$C = \text{let } x = \text{ref } 0 \text{ in } \langle \lambda f : \text{unit} \rightarrow \text{unit}. [\cdot], \lambda z : \text{unit}. !x \rangle$$

It implements a counter object with two methods: an increment function, and a get function inquiring the current counter value. A special feature of this object is that its increment method takes a callback argument that is invoked before the counter actually is incremented.

Consider the following alternative implementation for this object, in which  $x$  is dereferenced *before* the callback:

$$e' = C[\text{let } n = !x \text{ in } f(); x := n + 1]$$

One might naïvely assume that the two versions are equivalent, because  $x$  is not globally visible. But  $f$  might perform arbitrary operations, including recursively calling the increment function! In this case,  $x$  may be modified between read and write access in  $e'$ .

Such reentrance can be prevented by adding a lock:

$$C = \text{let } b = \text{ref true in} \\ \text{let } x = \text{ref } 0 \text{ in} \\ \langle \lambda f : \text{unit} \rightarrow \text{unit}. \\ (\text{if } !b \text{ then } (b := \text{false}; [\cdot]; b := \text{true}) \text{ else } ()) \\ \lambda z : \text{unit}. !x \rangle$$

Note that it is still possible for  $f$  to invoke the get function, which just *reads* the current  $x$ .

With  $C$  reimplemented using a lock,  $e$  and  $e'$  are now contextually equivalent. But how do we go about actually proving this? To show the two increment functions equivalent, we need to establish that  $f$  cannot modify  $x$ . But how can we set up an island that ensures that? After all, the island's law certainly *has to* allow updates to  $x$  in *general*. How can we formulate a law that allows the store to change, but still can *temporarily* prohibit it?

Steps to the rescue! When proving that the two increment functions are related, we assume that one terminates with  $j$  steps. Assuming  $b$  is set to `true` (i.e., assuming that  $x$  is "unlocked"), we can partition the reduction sequence for its execution into 3 phases of length  $j_1 + j_2 + j_3 = j$ , where  $j_2$  spans the steps spent in the call to  $f$ . These  $j_2$  steps are the time window in which  $x$  is not allowed to change. So the idea is to define a law that allows setting up time windows of this kind, during which  $!x$  is constant.

The following island definition does the trick:

$$w_{p+1} = (\eta_{k_0}^{\{(k_0, k_0, 0)\}}, \mathcal{L}_{k_0}) \\ \eta_k^V = (\psi_k^{\text{min}(V)}, V, \{l_b : \text{bool}, l_x : \text{int}\}, \{l'_b : \text{bool}, l'_x : \text{int}\}) \\ \psi_k^{\langle k_1, k_2, n \rangle} = \{(j, W, s, s') \in \text{StoreAtom}_k \mid \\ (k_1 \geq j \wedge s(l_b) = s'(l'_b) \wedge s(l_x) = s'(l'_x)) \wedge \\ (k_2 > j \vee (s(l_b) = \text{false} \wedge s(l_x) = n))\} \\ \mathcal{L}_k = \{(j, \eta_j^V) \in \text{LawAtom}_k \mid \\ V = \{\langle k_1, k'_1, n_1 \rangle, \dots, \langle k_m, k'_m, n_m \rangle\} \wedge \\ k_1 \geq k'_1 > k_2 \geq \dots \geq k'_m > k_m \geq k'_m\}$$

Each window is represented by a triple  $\langle k_1, k_2, n \rangle$  in  $V$  (assuming the obvious encoding of triples using pairs), with  $k_1$  and  $k_2$  giving its first (upper) and last (lower) step, and  $n$  being the value to which  $x$  is fixed during the window. The side condition in  $\mathcal{L}_k$  ensures that windows do not overlap. Consequently, there is always a unique lowest (newest) window  $\text{min}(V) = \langle k_1, k_2, n \rangle$ , i.e., the one with the least first projection (the  $k_1$ ). The store relation  $\psi$  is set up such that it requires  $!x$  to be the  $n$  taken from the newest window, *whenever the step level  $j$  has not yet passed that window's lower bound, i.e. whenever  $j \geq k_2$* . The definition also disallows windows that start in the future (where  $j > k_1$ ).

To prove equivalence of the increment functions, starting at step  $k$  with  $s(l_x) = s'(l'_x) = n$  and  $s(l_b) = s'(l'_b) = \text{true}$  (the interesting case), we proceed in  $j_1$  steps to set  $b$  to **false**, and then add a new lowest window  $\langle k - j_1, k - j_1 - j_2 - 1, n \rangle$  to the population of the  $(p+1)$ -th island. Next, we know  $f()$  returns after exactly  $j_2$  steps in some future world  $W$ , and the stores  $s$  and  $s'$  that it returns must be related by  $W$  at step  $m = k - j_1 - j_2$ , which means that  $(m - 1, [W]_{m-1}, s, s') \in W[p + 1].\psi$ . Since the step level  $m - 1$  is still in the range of the window we installed, we know that  $f()$  could not have added an even lower window to the population of the  $(p+1)$ -th island (as the law disallows adding windows that start in the future). Thus, we know that  $W[p + 1].\psi = \psi_m^{\langle m + j_2, m - 1, n \rangle}$ , and consequently it is still the case that  $s(l_x) = s'(l'_x) = n$  and  $s(l_b) = s'(l'_b) = \text{false}$ . That is, thanks to our use of the lock, the call to  $f()$  could not have affected our local state.

## 5.6 Deferred Divergence

To conclude, here is one interesting example, suggested to us by Hongseok Yang, that our method *cannot* handle:<sup>2</sup>

```
e1 = λf : unit → unit. f (λz : unit. diverge)
e2 = λf : unit → unit.
    let x = ref 0 in let y = ref 0 in
    f (λz : unit. if !x = 0 then y := 1 else diverge);
    if !y = 0 then x := 1 else diverge
```

Here,  $f$  may either call its argument directly, in which case the computation clearly diverges (in  $e_2$  this happens eventually because  $y$  is set to 1), or it may store its argument in some `ref` cell. In the latter case, any subsequent call to the stored argument by the program context will also cause divergence (in the case of  $e_2$ , because  $x$  will be 1 at that point). Only if neither  $f$  nor the context ever tries to call  $f$ 's argument may the computation terminate.

For us to prove  $e_1$  and  $e_2$  equivalent, we would need some way of relating the two arguments to  $f$ . Initially, however, when the arguments are invoked, one terminates and the other does not, so it is not obvious how to relate them. In fact, they are only related under the knowledge of what  $e_1$  and  $e_2$  will do after the call to  $f$ . This suggests to us that one way to handle such an example might be to define a relation on terms coupled with their continuations.

## 6. Related and Future Work

There is a vast body of work on methods for reasoning about local state and abstract data types. For space reasons, we only cite a representative fraction of the most closely related recent work, but there are numerous other papers that we consider relevant as well. We encourage the reader to let us know what citations we have omitted that they feel are critically important.

**Logical Relations** Our work continues (and, to an extent, synthesizes) two lines of recent work: one on using logical relations to reason about type abstraction in more realistic languages, the other on using logical relations to reason about local state.

Concerning the former, Pitts [17] provides an excellent overview, although it is now slightly out-of-date — in the last few years, several different logical relations approaches have been proposed for handling general recursive (as well as polymorphic) types [14, 2, 10], which Pitts considers an open problem. Much of the work on this topic is concerned with logical relations that are both sound and complete with respect to contextual equivalence. Completeness is useful for establishing various *extensionality properties* at different types, e.g., that two values of type  $\forall\alpha.\tau$  are contextually

<sup>2</sup> A similar example is discussed in Benton and Leperchey [8], at the end of their section 5. However, the two terms in their example are not actually equivalent in our language, because we have higher-order store.

equivalent iff their instantiations at any particular type  $\tau'$  are equivalent. In general, however, just because a method is complete with respect to contextual equivalence does not mean that it is *effective* in proving all contextual equivalences. In fact, Pitts gives a representation independence example for which existing techniques are “effectively” incomplete.<sup>3</sup>

For a logical relation to be complete it must typically be what Pitts terms “equivalence-respecting.” There are different ways to achieve this condition, such as  $\top\top$ -closure [17], biorthogonality [14], or working with contextual equivalence classes of terms [10]. Pitts’  $\top\top$ -closure neatly combines the equivalence-respecting property together with *admissibility* (or *continuity*, necessary for handling recursive functions) into one package.

We build on the work of Ahmed [2] on step-indexed logical relations for recursive and quantified types. One advantage of the step-indexed approach is that admissibility comes “for free,” in the sense that it is built directly into the model. By only ever reasoning about finite approximations of the logical relation  $(\mathcal{V}_n \llbracket \tau \rrbracket) \rho$ , we avoid the need to ever prove admissibility. (In other words, an inadmissible relation is indistinguishable from an admissible one if one only ever examines its step-indexed approximations.) Of course, the price one pays for this is that one is forced to use stepwise reasoning *everywhere*, so admissibility is not really “free” after all. To ameliorate this burden, we are currently investigating techniques for proving logical approximation in our model *without* stepwise reasoning. As we showed in Section 5.5, though, sometimes the presence of the step indices can be helpful.

Like Ahmed’s previous work, our logical relation is sound, but not complete, with respect to contextual equivalence. (Hers is complete except for the case of existential types.<sup>4</sup>) While our method cannot in its current form prove extensionality properties of contextual equivalence, it is still useful for proving representation independence results, which is our primary focus. Recent work by Ahmed and Blume [3] involves a variant of [2] that *is* complete with respect to contextual equivalence, where completeness is obtained by essentially Church-encoding the logical interpretation of existentials (this is roughly similar to what  $\top\top$ -closure does, too). We are currently attempting to develop a complete version of our method, using a similar approach to Ahmed and Blume.

Concerning the second line of work — logical relations for reasoning about local state — most of the recent previous work we know of employs possible-worlds models of the sort we discussed in Section 3.2, so we refer the reader to that earlier section for a detailed comparison with previous approaches [18, 19, 8, 9].

Recently, Acar *et al.* [1] gave an untyped relational step-indexed model for reasoning about the consistency of imperative self-adjusting computation. The model employs a simple notion of possible worlds for reasoning about accessibility of memory locations, but the worlds are restricted to only allow partial bijections between locations.

**Bisimulations** For reasoning about contextual equivalences (involving either type abstraction or local state), one of the most successful alternatives to logical relations is the coinductive technique of *bisimulations*. Sumii and Pierce define bisimulations for an untyped language with a dynamic sealing operator [24], as well as an extension of System F with general recursive types [25]. Koutavas and Wand [12] adapt the Sumii-Pierce technique to han-

<sup>3</sup> As it turns out, Pitts’ example is provable quite easily by a *transitive* combination of logical relations proofs ([tti-c.org/dreyer/pitts.txt](http://tti-c.org/dreyer/pitts.txt)). Dreyer has suggested a harder example, mentioned in Sumii and Pierce [25] (page 25), for which there is not even any known “brute-force” proof.

<sup>4</sup> The published conference version of her paper claims full completeness, but the proof contains a technical flaw uncovered by the second author. The extended version of her paper corrects the error [2].

de an untyped higher-order language with general references; in the process, they improve on Sumii-Pierce’s treatment of contextual equivalences involving higher-order functions. Interestingly, the Koutavas-Wand technique involves the use of inductive stepwise reasoning when showing that two functions are in the bisimulation, thus leading Koutavas to dub his technique *inductive bisimulations*. Subsequently, Sangiorgi *et al.* [23] propose *environmental bisimulations*, which generalize Sumii and Pierce’s previous work to an untyped framework subsuming that of Koutavas-Wand’s, but in a way that does not appear to require any stepwise reasoning. While all of these bisimulation approaches are sound and complete with respect to contextual equivalence, none handles a language with both existential type abstraction and mutable state.

There are many similarities between bisimulations and logical relations, although a precise comparison of the techniques remains elusive (and an extremely interesting direction for future work). With bisimulations, one defines the relational interpretations of abstract types, or the invariants about local state, *up front*, as part of a relation also containing the terms one wishes to prove contextually equivalent, and then one proceeds to show that the relation one has defined is in fact a bisimulation. With logical relations, the proof proceeds backward in a structured way from the goal of showing two terms logically equivalent, and the invariants about type representations or local state are chosen in mid-proof. It is arguably easier to sketch a bisimulation proof (by just stating the bisimulation), whereas the islands and  $\chi$  definitions in our proof sketches must be stated *in medias res*. On the other hand, our islands and  $\chi$ ’s are more minimal than bisimulations, which must often explicitly include a number of redundant intermediate proof steps.

The Sumii-Pierce-Koutavas-Wand-Sangiorgi-Kobayashi-Sumii approach is roughly to define bisimulations as sets of relations, with each relation tied to a particular *environment*, *e.g.*, a type interpretation, a pair of stores, etc. Various “up-to” techniques are used to make bisimulations as small as possible. This approach seems conceptually similar to possible-worlds semantics, but the exact relationship is unclear, and we plan to explore the connection further in future work.

**Separation Logic** To reason about imperative programs in a localized manner, O’Hearn, Reynolds *et al.* introduced *separation logic* [21] as an extension to Hoare logic. Separation logic has been enormously influential in the last few years, but it has not to our knowledge been used to reason about higher-order typed functional languages with type abstraction and higher-order store. Notably, however, the desire to scale separation logic to reason about a functional programming language has led to Hoare Type Theory (HTT) [16]. HTT is a dependently typed system where computations are assigned a monadic type in the style of a Hoare triple. Under this approach, programs generally have to pass around explicit proof objects to establish properties. Currently, HTT only handles strong update (where a location’s type can vary over time), not ML-style references with weak update (and thus stronger invariants).

**Relational Reasoning About Classes** There is a large body of work on reasoning techniques for object-oriented languages. For example, Banerjee and Naumann [6] present a denotational method for proving representation independence for a Java-like language. Koutavas and Wand [13] have adapted their bisimulation approach to a subset of Java. The languages considered in these works do not provide generativity and first-class existential types, but rather tie encapsulation to static class definitions. We believe that the generativity of existential quantification and the separation enforced by possible-island semantics are closely related to various notions of ownership and ownership types, but we leave the investigation of this correspondence to future work.

## References

- [1] Umur A. Acar, Amal Ahmed, and Matthias Blume. Imperative self-adjusting computation. In *POPL*, 2008.
- [2] Amal Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *ESOP*, 2006. Extended/corrected version of this paper available as Harvard University TR-01-06.
- [3] Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *ICFP*, 2008.
- [4] Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent representation independence (technical appendix), 2008. Available at: <http://ttic.uchicago.edu/~amal/papers/sdri/>.
- [5] Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *Transactions on Programming Languages and Systems*, 23(5):657–683, 2001.
- [6] Anindya Banerjee and David A. Naumann. Ownership confinement ensures representation independence in object-oriented programs. *Journal of the ACM*, 52(6):894–960, 2005.
- [7] Anindya Banerjee and David A. Naumann. State based ownership, reentrance, and encapsulation. In *ECOOP*, 2005.
- [8] Nick Benton and Benjamin Leperchey. Relational reasoning in a nominal semantics for storage. In *TLCA*, 2005.
- [9] Nina Bohr and Lars Birkedal. Relational reasoning for recursive types and references. In *APLAS*, 2006.
- [10] Karl Cray and Robert Harper. Syntactic logical relations for polymorphic and recursive types. In *Computation, Meaning and Logic: Articles dedicated to Gordon Plotkin*. 2007.
- [11] Derek Dreyer, Karl Cray, and Robert Harper. A type system for higher-order modules. In *POPL*, 2003.
- [12] Vasileios Koutavas and Mitchell Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL*, 2006.
- [13] Vasileios Koutavas and Mitchell Wand. Reasoning about class behavior. In *FOOLWOOD*, 2007.
- [14] Paul-André Melliès and Jérôme Vouillon. Recursive polymorphic types and parametricity in an operational framework. In *LICS*, 2005.
- [15] John C. Mitchell. Representation independence and data abstraction. In *POPL*, 1986.
- [16] Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. Hoare type theory, polymorphism and separation. In *ICFP*, 2006.
- [17] Andrew Pitts. Typed operational reasoning. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, pages 245–289. The MIT Press, 2005.
- [18] Andrew Pitts and Ian Stark. Operational reasoning for functions with local state. In *HOOTS*, 1998.
- [19] Uday Reddy and Hongseok Yang. Correctness of data representations involving heap data structures. In *ESOP*, 2003.
- [20] John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing*, 1983.
- [21] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.
- [22] Claudio V. Russo. Non-dependent types for Standard ML modules. In *PPDP*, 1999.
- [23] Davide Sangiorgi, Naoki Kobayashi, and Eijiro Sumii. Environmental bisimulations for higher-order languages. In *LICS*, 2007.
- [24] Eijiro Sumii and Benjamin Pierce. A bisimulation for dynamic sealing. *Theoretical Computer Science*, 375(1–3):161–192, 2007.
- [25] Eijiro Sumii and Benjamin Pierce. A bisimulation for type abstraction and recursion. *Journal of the ACM*, 54(5):1–43, 2007.
- [26] Janis Voigtländer and Patricia Johann. Selective strictness and parametricity in structural operational semantics, inequationally. *Theoretical Computer Science*, 388(1–3):290–318, 2007.