

Romeo: A system for more flexible binding-safe programming*

PAUL STANSIFER and MITCHELL WAND

*College of Computer and Information Science,
Northeastern University Boston, Massachusetts, USA*
e-mail: pauls@ccs.neu.edu and wand@ccs.neu.edu

Abstract

Current systems for safely manipulating values containing names only support simple binding structures for those names. As a result, few tools exist to safely manipulate code in those languages for which name problems are the most challenging. We address this problem with Romeo, a language that respects α -equivalence on its values, and which has access to a rich specification language for binding, inspired by attribute grammars. Our work has the complex-binding support of David Herman's λ_m , but is a full-fledged binding-safe language like Pure FreshML.

1 Introduction

Name collision, the appearance of a name in a context where it means something other than what was intended, is usually thought of as a minor hazard of programming. But in the context of metaprogramming, the danger is significantly more serious, first because the metaprogrammer has no idea what names will appear in input syntax, and second because it is a natural programming practice to create copies of a single piece of syntax and expect each copy to behave independently, even when interleaved.

For example, suppose a metaprogrammer writes a function that, when given two lambda expressions, $(\text{lambda } (x_1) e_1)$ and $(\text{lambda } (x_2) e_2)$, constructs

$$(\text{lambda } (x_1) (\text{lambda } (x_2) (e_1 e_2))).$$

If such a function were passed (syntax for) two identity functions, it should produce (syntax for) the composition function:

$$(\text{lambda } (a) (\text{lambda } (b) (a b)))$$

But if both its arguments are the same identity function, shadowing will (in many metaprogramming systems) change the meaning to something else:

$$(\text{lambda } (c) (\text{lambda } (c) (c c)))$$

* This material is based on research sponsored by the Defense Advanced Research Projects Agency and the Air Force Research Laboratory under agreement number FA8750-10-2-0233. Any opinions, findings, conclusions or recommendations expressed herein are those of the authors, and do not necessarily reflect those of the US Government, DARPA, or the Air Force.

Bugs resulting from name collision can be difficult to track down and fix. If the user of a metaprogramming system encounters such a bug, it will manifest, at best, as a static error in *generated* code. The user must reason about the internals of the metaprogramming system to determine how they caused that code to be generated. In the worst case, it will manifest as a behavioral anomaly that is invisible to code inspection.

Unfortunately, metaprogrammers do not typically receive much help from their languages and tools; a recent survey of nine tools for - domain-specific language implementation found that eight of them were vulnerable to errors or incorrect behavior as a result of name collision (Erdweg *et al.*, 2014).

In order to solve this problem, it is necessary to codify the user's intentions. This is the responsibility of binding specifications, which indicate, for a particular term, how the set of names in scope in its subterms differ from the set of names in scope for it. Typically, each form in a language has a binding specification, and the binding behavior of a term is determined by looking up which form it corresponds to.

A system with binding specifications can be *binding-safe*. In binding-safe systems, α -equivalent inputs proceed to α -equivalent outputs (Herman & Wand, 2008). From the point of view of the user (either the programmer or the metaprogrammer), such a system behaves the same regardless of what names the user chooses. The primary benefit to the user is that name collision errors are impossible, but we have also discovered a secondary benefit: It is possible to write simpler code when terms can always be safely put into binding forms without fear of accidentally capturing free variables.

Two binding-safe metaprogramming systems are of particular interest to us: λ_m and Pure FreshML.

1.1 The λ_m calculus

David Herman's λ_m calculus (2010) is inspired by the world of Scheme macros.

Typically, Scheme is made up of a very small number of core forms, sufficiently expressive to write any program, but not user-friendly. For example, it might contain `lambda` but not `let`. The user (and the author of the standard library) uses metaprogramming to implement each user-friendly language form (macro) by describing how to translate (expand) it into core forms (or just simpler macros). The Scheme macro system orchestrates this expansion process.

Scheme predates the notion of binding safety, but Scheme programmers have been attacking the problem of name collision since at least 1986 (Kohlbecker *et al.*, 1986). The standard approach Schemers have developed is called "hygiene". Unlike binding safety, which is a property of the metaprogramming system, hygiene is a property of the macro *expander*, operating by annotating names on the input and output of the metaprogram that is a macro implementation, rather than affecting the execution of the metaprogram itself. Furthermore, hygiene does not require binding specifications.

One disadvantage of this "lightweightness" is that, without binding specifications to indicate the intended binding behavior of macros, it is hard to say what it means

for a system to be hygienic. In fact, it was not until 2015 that a notion of what it means for an implementation of hygiene to be correct was proposed (Adams, 2015). This makes it difficult to explain to the programmer how to understand hygiene, and makes it almost impossible to build any static guarantees off hygiene.

The λ_m -calculus addresses this by introducing binding specifications for macros, and then defining a binding-safe macro expander. One of the major challenges is that a single macro invocation may include arbitrarily many subterms, each with its own scope, and while each of those scopes may introduce a different set of names, the same name may be bound in many of them. λ_m 's binding specification system is powerful enough to express certain situations like this, including many standard Scheme macros.

The primary limitation of λ_m (from our point of view) is that it is merely a pattern-matching macro system. This means that the metaprogrammer can only specify an expected invocation syntax for a macro (with “holes” for parameters), and a template for transcription (with interpolations from those holes), as opposed to writing a function that could perform arbitrary computation. In Scheme parlance, this is like `syntax-rules` (minus Macro By Example's treatment of . . .), as opposed to `syntax-case`. This makes many common macros impossible to define.

Furthermore, macros are only one of many forms of metaprogramming that can benefit from binding safety.

1.2 Pure FreshML

Francois Pottier's Pure FreshML (2007a) is a core calculus of a functional language for metaprogramming, in which values are terms with binding structure. Programs in Pure FreshML are binding-safe. In contrast to systems that easily achieve binding safety by representing terms as functions (higher order abstract syntax) or by representing names canonically (de Bruijn indices), terms in Pure FreshML are ordinary trees that contain names, which seems to be the most ergonomic way to metaprogram.

The most important aspect of its enforcement of binding safety is that, whenever a binding term v is destructured (i.e. passes through Pure FreshML's pattern-matching construct), the name it binds (say, a) is automatically “freshened”, meaning that all occurrences of a in v are replaced by a new name. This prevents name collisions, but does not produce a fully binding-safe system.¹

Binding safety requires α -equivalent inputs to proceed to α -equivalent outputs, which is essentially a guarantee of determinism. The generation of fresh names is non-deterministic, so Pure FreshML must ensure that the output of its programs don't contain any of those fresh names (at least as free names; non-determinism in bound names is acceptable according to the definition of binding safety). For a freshened name to escape the context, it was introduced in is an error, so Pure FreshML has a static proof system to prove that such an escape will never occur.

¹ Pure FreshML is based on FreshML (Shinwell *et al.*, 2003), which stops here and achieves a less strong safety guarantee.

$CoreExpr ::= RAtom$	(variable reference)	var-ref
$Prod(CoreExpr, CoreExpr)$	(application)	apply
$Prod(BAtom, CoreExpr \downarrow 0)$	(abstraction)	lambda
$Expr ::= RAtom$	(variable reference)	—
$Prod(Expr, Expr)$	(application)	—
$Prod(BAtom, Expr \downarrow 0)$	(abstraction)	—
$Prod(LetStarClauses, Expr \downarrow 0)$	(sequential let)	let*
$LetStarClauses ::= Prod()$	(no clauses)	none
$Prod^{\uparrow 1 \triangleright 0} (Prod^{\uparrow 0} (BAtom, Expr), LetStarClauses \downarrow 0)$	(clause, and more clauses)	lsc-some

Fig. 1. Example types for two lambda calculi, one of which has the `let*` form. (The names from the right-hand column are used to identify injections in Figure 2.)

However, the only kind of binding form permitted in Pure FreshML is a pair of a name and a term, where the name is in scope for the term (in other words, the binding structure of a lambda term). This is partially addressed by *Czml* (Pottier, 2006), which improves on Pure FreshML’s expressivity somewhat. However, even *Czml*’s binding specifications are significantly limited. In particular, *Czml* terms are divided into *expressions* (which contain references) and *patterns* (which contain binders), while terms in λ_m can play both roles simultaneously.

1.3 Example

For example, consider the following use of the `let*` syntactic form in Scheme, which exhibits a complex binding structure:

```
(let* ((a 1)
      (b (+ a a))
      (c (* b 5)))
      (display c))
```

The `let*` form is defined to bind the names it introduces not only in the body, but also in the right-hand side of each subsequent arm. Thus, in the code above, all references to names are well-defined, and the value of `c` is 10. This behavior is similar to the behavior of `do` in Haskell and to telescopes in dependently typed languages.

We wish to “expand away” `let*`s. In Figure 1, we write some binding types using (a syntactically sugared version of) Romeo, and in Figure 2, we define a function (using those types) that translates expressions from the lambda calculus augmented with a `let*` construct into the plain lambda calculus. The α -equivalence-preserving nature of the translation is a property of Romeo, and the key is Romeo’s destructuring construct, **open** (similar to the “`val <x1>x2 = e`” construct in FreshML (Shinwell *et al.*, 2003) or the **case** construct in Pure FreshML (Pottier, 2007b)). For example, (**open let-star** (*lsc-e-body*) ...) behaves the same way as the (`match let-star [(list lsc e-body) ...]`) in Scheme, except that the value of *let-star* is α -converted in a way that avoids name collision (see E-OPEN-* in Section 4). This α -conversion is guided by the type of the value being destructured.

```

1 (define-fn (convert e:Expr) : CoreExpr
2   (case e
3     (var => (injvar-ref var))
4     (app => (open app (e1, e2) (injapply (prod convert(e1), convert(e2)))))
5     (lam => (open lam (bv, e-body) (injlambda (prod bv, convert(e-body) ↓0))))
6     (let-star =>
7       (open let-star (lsc, e-body)
8         (case lsc
9           (none => convert(e-body))
10          (lsc-some =>
11            (open lsc-some (bv, val-expr, lsc-rest)
12              (let e-rest be convert((injlet* (prod lsc-rest, e-body) ↓0))
13                in (injapply (prod (injlambda (prod bv, e-rest ↓0)),
14                            convert(val-expr))))))))))

```

Fig. 2. A Romeo-L function to expand away `let*`.

If we had wanted to translate `let` (which only differs from `let*` in its binding behavior) instead, the exact same code would also serve that purpose; all that would need to be changed would be the types (see Section 4.3.2). We will discuss this example in more detail: We will cover the types in Section 2.1, the behavior of the code in Section 4.3.1, and the static guarantees the code provides in Section 8.3.

1.4 Contributions

Our primary contribution is an extension of David Herman’s system for binding-safety in a pattern-matching macro system (Herman, 2010) to cover macros defined by procedures, and thus general meta-programming for terms with bindings. Our language is inspired by Pure FreshML (Pottier, 2007b).

Our system has the following features:

- Values in Romeo are “plain old data”: Atoms arranged in abstract syntax trees without binding information. Types provide the missing binding information.
- Romeo has an execution semantics which guarantees that instead of a name “escaping” the context in which it is defined, a `FAULT` is produced.
- Thanks to that guarantee, we prove a theorem showing that, in any execution, the dynamic environment can be replaced by one with α -equivalent values, and that execution will proceed to a value α -equivalent to what it otherwise would have.
- We provide a deduction system with which the programmer can establish that escape (and thus, `FAULT`) will never occur.

2 Binding language

2.1 Overview of binding types

Values in our system are plain old data, that is, S-expressions or something similar. We use binding types to specify the binding properties of these terms. Binding types augment a traditional context-free grammar with a single attribute (in the style of

an attribute grammar) that represents the flow of bindings from one subterm to another.

In Figure 1, the type definition of *CoreExpr* looks like a traditional grammar for the lambda calculus, with one major difference: The notation *CoreExpr* ↓ 0 indicates that the binder “exported” by the child in position 0 (the *BAtom*, which exports the name in that position) is to be in scope in the *CoreExpr* body of the lambda. To facilitate the matching of related names together, the type for name introductions (or binders), *BAtom*, is made distinct from the type of names that reference binders, *RAtom*.

Instead of a binary product (\times), our type system has a “wide” product ($\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$). Using a wide product increases the expressivity of our system because the position references on the right-hand-side of ↓ (and ↑) are indices into the product. Manipulating forms that introduce multiple different but related scopes was one of the primary challenges in proving the correctness of our system.

Our system observes the convention that all names bound in a particular term are bound in all subterms, unless overridden by a new binding for the same name.²

Things get more complicated when the binders are exported longer distances up the tree. Returning to Figure 1, consider *let**, in *Expr*. The “sequential let” line indicates that the binders exported by *LetStarClauses* are in scope in the body of the *let** expression. The grammar for *LetStarClauses* says that a set of *let**-clauses is either the empty list or the cons of a single clause and a *LetStarClauses*.

In the second production for *LetStarClauses*, the $\text{Prod}^{\uparrow 0}$ (*BAtom*, *Expr*) indicates that the first clause exports the binder from its position 0 (that is, the *BAtom*). (However, this name is not in scope in the *Expr*.) The *LetStarClauses* ↓ 0 indicates that the names exported from the first clause are in scope in the remainder of the clauses.

The ↑ (1 ▷ 0) indicates that the entire *LetStarClauses* exports all the binders exported by either the first clause or the rest of the clauses, and that names in the rest of the clauses override those from the first clause. If we had wanted to specify that all the binders in a *let** must be distinct, then we could have written ↑ (1 ⊔ 0), which behaves like ↑ (1 ▷ 0), except that duplicated atoms are an error.

Thus, our type system can be seen as an attribute grammar with a single attribute, whose values are sets of names representing bindings. These sets are synthesized from binders and values that export them, and inherited by every term underneath a term that imports one of them.

Our notations ↓, ↑, ▷, and ⊔ form an algebra of attributes; the tractability of this algebra is a key to many of our results. We call the terms in this language *binding combinators*.

² It is possible to imagine a system in which old names are removable (e.g. a construct (*unbind* *x e*), in which the name *x* is not a valid reference in *e*, even if it was outside that construct), but this does not appear to be a feature that users are clamoring for. (But see the end-of-scope operator described by Hendriks and van Oostrom (2003).)

$Expr ::= \dots$ (same as before)
 | $\text{Prod}(\text{BAtom}, \text{BAtom}, Expr \downarrow 0 \uplus 1, \text{BAtom}, Expr \downarrow 0 \uplus 3)$ (event handler)

Fig. 3. Typed production rule for event handler.

2.1.1 Example: multiple, partially shared bindings

For another example, imagine constructing a pair of event handlers, one of which handles mouse events and one of which handles keyboard events, but both of which need to know what GUI element is focused at the time of the event. This new form, defined in Figure 3, binds three atoms (the BAtoms, which are in positions 0, 1, and 3), one of which is bound in both subexpressions, and two of which are bound in only one of them. Here is a possible use of this new form:

```
(handler gui-elt
  mouse-evt (deal-with gui-elt mouse-evt)
  kbd-evt (put-tag gui-elt (text-of kbd-evt)))
```

And, here is an α -equivalent, but harder-to-read, version:

```
(handler a
  b (deal-with a b)
  b (put-tag a (text-of b)))
```

The scope of the first `b` is the `(deal-with ...)`, and the scope of the second one is the `(put-tag ...)`.

Regardless of whether they have the same names, the meanings of the two events must not be conflated, but in both subterms, the GUI element is the same. For this reason, the operations our system performs on products must handle binding by first identifying what names are exported by each child (e.g. a BAtom or a Prod with a non-empty \uparrow), and then determining which names are imported by which children. The latter is the responsibility of the \downarrow operator.

Our goal of supporting realistic concrete syntax is particularly relevant here. The user could have implemented the `handler` statement as a *function* with the following style of expected invocation:

```
(handler-fn (lambda (gui-elt mouse-evt)
  (deal-with gui-elt mouse-evt))
  (lambda (gui-elt kbd-evt)
  (put-tag gui-elt (text-of kbd-evt))))
```

If programmer convenience were irrelevant, languages would need no binding constructs other than `lambda`. However, programmer convenience is precisely the point of metaprogramming systems.

2.2 Binding types, in more detail

In this section, we introduce our actual language of binding types and the metalanguage we use to describe them.

$\tau \in \text{Type} ::=$	BAtom		RAtom
$a \in \text{Atom}$			$\tau + \tau$
$v \in \text{Value} ::= a$			$\text{Prod}^{\uparrow\beta}(\tau_i \downarrow \beta_i)_i$
	$\mathbf{inj0}(v)$		$\mu X.\tau$
	$\mathbf{inj1}(v)$		X
	$\mathbf{prod}(v_i)_i$		$\mathbf{nth}_i \tau$

Values are either atoms, left- or right- injections of values (to model sum types), or tuples of values. We write $\mathbf{prod}(v_i)_i$ for the tuple $(v_0 \dots v_n)$, for some n . We will use notation like this for sequence comprehensions throughout our presentation.

The basic types are BAtom (for binders) and RAtom (for references). These types tell us how to interpret atoms. By convention, BAtoms export themselves and RAtoms export nothing.

Tuples are interpreted by Prod types. The wide product type $\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_0 \downarrow \beta_0, \dots, \tau_n \downarrow \beta_n)$, which we denote by the comprehension $\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$, tells us how to interpret the value $\mathbf{prod}(v_i)_i$. The term β_i , constructed in our algebra of attributes, combines (a subset of) the binders exported by v_0, \dots, v_n to determine the local names bound in v_i . By convention, these names override those inherited from outside (above) $\mathbf{prod}(v_i)_i$. The binding combinator β_{ex} , similarly constructed in our algebra of attributes, combines the binders exported by v_0, \dots, v_n to determine the names exported as binders by the whole tuple $\mathbf{prod}(v_i)_i$. If \downarrow or \uparrow is omitted, the corresponding β defaults to \emptyset .

Instead of parsing S-expressions, we have explicit sum types and injections. A value $\mathbf{inj0}(v)$ (resp. $\mathbf{inj1}(v)$) is interpreted by the type $\tau_0 + \tau_1$, so that v is interpreted by τ_0 (resp. τ_1).

Last, we have recursive types $\mu X.\tau$ (where τ must be productive), to interpret a value v according to $\mu X.\tau$ is to interpret it according to $\tau[(\mu X.\tau)/X]$. To support it, we have type variables X , and type-level destructors \mathbf{nth}_j . We define the latter as $\mathbf{nth}_j \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i \triangleq \tau_j$ (reusing our existing type constructor as a way to write type tuples, the β_i s and β_{ex} are ignored). The \mathbf{nth}_j construct is uninteresting on its own, but it allows for the definition of mutually recursive types.

2.3 The algebra of binding combinators

Binding combinators are terms built from the following grammar:

$$\begin{aligned} \ell &\in \mathbb{N} \\ \beta \in \text{Beta} &::= \emptyset \mid \beta \uplus \beta \mid \beta \triangleright \beta \mid \ell. \end{aligned}$$

As discussed above, we use binding combinators to collect names from the sets exported by the subterms of a sequence $\mathbf{prod}(v_i)_i$. We will need to interpret these combinators as operating over both sets of names and substitutions (finite maps from names to names). As before, we make liberal use of comprehensions: We write

$\llbracket \beta \rrbracket (A_i)_i$ for $\llbracket \beta \rrbracket (A_0, \dots, A_n)$, etc. The interpretation is as follows:

$$\begin{aligned} \llbracket _ \rrbracket (_) &: \text{Beta} \times \overline{\text{AtomSet}} \rightarrow \text{AtomSet} \\ \llbracket \emptyset \rrbracket (A_i)_i &\triangleq \emptyset \\ \llbracket \ell \rrbracket (A_i)_i &\triangleq A_\ell \\ \llbracket \beta \triangleright \beta' \rrbracket (A_i)_i &\triangleq \llbracket \beta \rrbracket (A_i)_i \cup \llbracket \beta' \rrbracket (A_i)_i \\ \llbracket \beta \uplus \beta' \rrbracket (A_i)_i &\triangleq \llbracket \beta \rrbracket (A_i)_i \uplus \llbracket \beta' \rrbracket (A_i)_i. \end{aligned}$$

Here and elsewhere, we write \overline{X} to mean a sequence of X s.

In Romeo, constructing a value whose type contains a β that contains a \uplus that attempts to union two non-disjoint sets of names is an error. We omit checking for this error, as it is straightforward, and is merely provided for metaprogrammers to enforce their intended usage rules.

A substitution σ is a partial function from atoms to atoms. For the purposes of manipulating them, we represent a substitution as a set of ordered pairs of atoms. Our substitutions are naive, which is to say that they ignore binding structure and simply affect all names. We interpret β 's on substitutions as follows:

$$\begin{aligned} \llbracket _ \rrbracket (_) &: \text{Beta} \times \overline{\text{Subst}} \rightarrow \text{Subst} \\ \llbracket \emptyset \rrbracket (\sigma_i)_i &\triangleq \emptyset \\ \llbracket \ell \rrbracket (\sigma_i)_i &\triangleq \sigma_\ell \\ \llbracket \beta \triangleright \beta' \rrbracket (\sigma_i)_i &\triangleq \llbracket \beta \rrbracket (\sigma_i)_i \triangleright \llbracket \beta' \rrbracket (\sigma_i)_i \\ \llbracket \beta \uplus \beta' \rrbracket (\sigma_i)_i &\triangleq \llbracket \beta \rrbracket (\sigma_i)_i \uplus \llbracket \beta' \rrbracket (\sigma_i)_i. \end{aligned}$$

We define the ‘‘override’’ operation $\sigma \triangleright \sigma'$ as follows:

$$\sigma \triangleright \sigma' \triangleq \sigma \cup \left\{ \langle a_d, a_r \rangle \mid \langle a_d, a_r \rangle \in \sigma' \right. \\ \left. a_d \notin \text{dom}(\sigma) \right\}$$

The operation for combining disjoint substitutions $\sigma \uplus \sigma'$ is like \triangleright , except that it is undefined if the domains of the substitutions in question overlap.

Now, using $\llbracket \beta \rrbracket (A_i)_i$, we can compute the *exported binders* from any value. These are also called the free binders, because they are considered to be free names in their term (the expected theorems relating free names to α -equivalence only hold if we do this). The free binders of a value are determined using the value's type.

$$\begin{aligned} \text{fb}(\text{Prod}^{\uparrow \beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i, \mathbf{prod}(v_i)_i) &\triangleq \llbracket \beta_{\text{ex}} \rrbracket (\text{fb}(\tau_i, v_i))_i \\ \text{fb}(\tau_0 + \tau_1, \mathbf{inj0}(v)) &\triangleq \text{fb}(\tau_0, v) \\ \text{fb}(\tau_0 + \tau_1, \mathbf{inj1}(v)) &\triangleq \text{fb}(\tau_1, v) \\ \text{fb}(\mu X. \tau, v) &\triangleq \text{fb}(\tau[\mu X. \tau / X], v) \\ \text{fb}(\text{BAtom}, a) &\triangleq \{a\} \\ \text{fb}(\text{RAtom}, a) &\triangleq \emptyset \end{aligned}$$

There are several other useful quantities that we can compute using these combinators. First is the set of free references of a term:

$$\begin{aligned} \text{fr}(\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i, \mathbf{prod}(v_i)_i) &\triangleq \bigcup_i \left(\text{fr}(\tau_i, v_i) \setminus \llbracket \beta_i \rrbracket (\text{fb}(\tau_j, v_j))_j \right) \\ \text{fr}(\tau_0 + \tau_1, \mathbf{inj0}(v)) &\triangleq \text{fr}(\tau_0, v) \\ \text{fr}(\tau_0 + \tau_1, \mathbf{inj1}(v)) &\triangleq \text{fr}(\tau_1, v) \\ \text{fr}(\mu X. \tau, v) &\triangleq \text{fr}(\tau[\mu X. \tau/X], v) \\ \text{fr}(\text{BAtom}, a) &\triangleq \emptyset \\ \text{fr}(\text{RAtom}, a) &\triangleq \{a\} \end{aligned}$$

The set of free atoms of a term is the union of the free references and the exported (or free) binders:

$$\text{fa}(\tau, v) \triangleq \text{fr}(\tau, v) \cup \text{fb}(\tau, v)$$

The set of exposable atoms is the set of those non-free names in a value that will become free when that value is broken into subterms. These are the atoms which are on their “last chance” for renaming before they become free. This set, only defined on products, is equal to the union of the binders exported by each term in a sequence, less the terms that are exported to the outside:

$$\text{xa}(\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i, \mathbf{prod}(v_i)_i) \triangleq \left(\bigcup_i \text{fb}(\tau_i, v_i) \right) \setminus \text{fb}(\text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_j \downarrow \beta_j)_j, \mathbf{prod}(v_j)_j)$$

Finally, we can determine whether a particular index is in the support of (i.e. referred to by) β :

$$\begin{aligned} _ \hat{=} _ &\subseteq \mathbb{N} \times \text{Beta} \\ \ell \hat{=} \emptyset &\triangleq \text{false} \\ \ell \hat{=} \ell' &\triangleq (\ell = \ell') \\ \ell \hat{=} \beta \triangleright \beta' &\triangleq \ell \hat{=} \beta \text{ or } \ell \hat{=} \beta' \\ \ell \hat{=} \beta \uplus \beta' &\triangleq \ell \hat{=} \beta \text{ or } \ell \hat{=} \beta' \end{aligned}$$

3 α -equivalence

Our next task is to go from a binding type to a notion of α -equivalence on values described by that type. Because our binding types allow for buried binders (i.e. binders that may be an arbitrary depth from the form that binds them) to be exported, we define two values to be α -equivalent, if both:

- they export identical binders, and
- their non-exported binders can be renamed along with the names that reference them to make the terms identical.

$$_ =_{\alpha} _ : _ \subseteq \text{Value} \times \text{Value} \times \text{Type}$$

$$\frac{v =_{\text{B}} v' : \tau \quad v =_{\text{R}} v' : \tau}{v =_{\alpha} v' : \tau} \alpha\text{EQ}$$

We use $=_B$ (pronounced “binder-equivalent”) for the first relation and $=_R$ (pronounced “reference-equivalent”) for the second.

3.1 Binder equivalence

Two values are $=_B$ iff their exported (free) binders in the same positions are identical. Note that non-exported binders are irrelevant to $=_B$, but are important to the calculation of $=_R$.

Here and throughout, we omit the rules for injections and fixed points, which are trivial.

$$\begin{array}{c}
 _ =_B _ : _ \subseteq \text{Value} \times \text{Value} \times \text{Type} \\
 \\
 \frac{}{a =_B a : \text{BAtom}} \text{B}\alpha\text{-BATOM} \qquad \frac{}{a =_B a' : \text{RAtom}} \text{B}\alpha\text{-RATOM} \\
 \\
 \frac{\forall i \hat{\in} \beta_{\text{ex}}. v_i =_B v'_i : \tau_i}{\mathbf{prod}(v_i)_i =_B \mathbf{prod}(v'_i)_i : \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i} \text{B}\alpha\text{-PROD}
 \end{array}$$

For example,

$$\mathbf{prod}(a b c) =_B \mathbf{prod}(a b d) : \text{Prod}^{\uparrow 0 \triangleright 1}(\text{BAtom}, \text{BAtom}, \text{BAtom})$$

because non-exported atoms are ignored, but

$$\mathbf{prod}(a b c) \neq_B \mathbf{prod}(b a c) : \text{Prod}^{\uparrow 0 \triangleright 1}(\text{BAtom}, \text{BAtom}, \text{BAtom})$$

because exported names in each position must be the same.

3.2 Reference equivalence

Calculating $=_R$ is analogous to the conventional notion of α -equivalence, except that we need to extract and rename the bindings that are buried in subterms.

3.2.1 Joining the binders

The first step in the wide product case is to match binders in identical positions with each other, for which we must define the \bowtie operator (pronounced “join”). It walks through both values in lockstep, assigning a common fresh atom for each pair of corresponding binding atoms. The result is a pair of injective substitutions whose domains are equal to the set of exported binders (fb) of the values being joined.

To calculate \bowtie in the wide product case, we recursively generate such a pair of substitutions for each subterm of the products being compared, make sure that the generated names (the ranges of those substitutions) are disjoint, and then combine those substitutions with $\llbracket \beta_{\text{ex}} \rrbracket$. The resulting pair of substitutions is the output of the \bowtie relation.

$$_ \bowtie _ : _ \rightarrow _ \bowtie _ \subseteq \text{Value} \times \text{Value} \times \text{Type} \times \text{Subst} \times \text{Subst}$$

$$\frac{}{a \bowtie a' : \text{BAtom} \rightarrow \{\langle a, a_{\text{fresh}} \rangle\} \bowtie \{\langle a', a_{\text{fresh}} \rangle\}} \text{J-BATOM}$$

$$\frac{}{a \bowtie a' : \text{RAtom} \rightarrow \emptyset \bowtie \emptyset} \text{J-RATOM}$$

$$\frac{\forall i. v_i \bowtie v'_i : \tau \rightarrow \sigma_i \bowtie \sigma'_i \quad \forall i \neq j. \text{rng}(\sigma_i) \# \text{rng}(\sigma_j) \quad \sigma = \llbracket \beta_{\text{ex}} \rrbracket (\sigma_i)_i \quad \sigma' = \llbracket \beta_{\text{ex}} \rrbracket (\sigma'_i)_i}{\text{prod}(v_i)_i \bowtie \text{prod}(v'_i)_i : \text{Prod}^{\uparrow \beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i \rightarrow \sigma \bowtie \sigma'} \text{J-PROD}$$

Here and elsewhere, we use $\#$ to denote the disjointness operator over names, sets of names, and (the names in) values. It is naive, meaning that it entirely ignores binding structure. (Note that in J-PROD, we only examine the range of the substitutions without primes because the ranges of the two substitutions emitted by \bowtie are identical.)

For example, consider the two `let*` expressions we have previously discussed:

$$\begin{array}{ll} (\text{let*} ((a 1) & (\text{let*} ((d 1) \\ (b (+ a a)) & (d (+ d d)) \\ (c (* b 5))) & (d (* d 5))) \\ (\text{display } c)) & (\text{display } d)) \end{array}$$

The results of \bowtie on their subterms in position 1 (the `display` expressions) are \emptyset and \emptyset , because neither one has any free binders. Position 0 corresponds to the *LetStarClauses*, and is more interesting. \bowtie will non-deterministically generate three names, which we will choose to be `aa`, `bb`, and `cc`. Then, we will have $\sigma_0 = \{\langle a, aa \rangle, \langle b, bb \rangle, \langle c, cc \rangle\}$ and $\sigma'_0 = \{\langle d, cc \rangle\}$. The different ranges of these substitutions indicate that some names (the ones called `a` and `b` in the left-hand value) are shadowed and cannot be referred to at all by references on the right-hand side.

A more complete derivation is available in the supplementary material.

3.2.2 Comparison by substitution

Now, we can write the rules for $=_R$. At the type `RAtom`, the atoms being compared are necessarily free, and are reference equal iff they are identical. Symmetrically to $=_B$, any two atoms are $=_R$ at `BAtom`.

At a wide product, the information from performing \bowtie on the subterms pairwise is used to unify references that refer to binders in the same position. This is done as follows: For each pair of subterms v_i, v'_i , we use \bowtie to generate a pair of substitutions σ_i, σ'_i that rename the binders exported by v_i, v'_i to be identical.

Then, for each subterm v_i , we compute $\llbracket \beta_i \rrbracket (\sigma_j)_j(v_i)$ (and the symmetrical value for v'_i), adjusting all the imported names (as defined in β_i) so that, regardless of the difference between v_i and v'_i , references to binders in the same position become equal to a single new value (as generated by \bowtie).

j	v_j	v'_j	σ_j	σ'_j
0	<code>gui-elt</code>	\bowtie a	:BAtom	\rightarrow $\{\{\text{gui-elt,gg}\}\} \bowtie \{\{a,gg\}\}$
1	<code>mouse-evt</code>	\bowtie b	:BAtom	\rightarrow $\{\{\text{mouse-evt,mm}\}\} \bowtie \{\{b,mm\}\}$
2	<code>(handle-evt ...)</code>	\bowtie (handle-evt ...)	:Expr	\rightarrow $\emptyset \bowtie \emptyset$
3	<code>kbd-evt</code>	\bowtie b	:BAtom	\rightarrow $\{\{\text{kbd-evt,kk}\}\} \bowtie \{\{b,kk\}\}$
4	<code>(tag ...)</code>	\bowtie (tag ...)	:Expr	\rightarrow $\emptyset \bowtie \emptyset$

Fig. 4. Substitutions generated for the handler example.

Note that this substitution is naive (that is, it disregards types and therefore binding). Even though we are only interested in the substitution’s effect on free references, this naivete is acceptable because, first, $=_R$ does not examine free binders, and second (broadly speaking), the substitution of unbound names is harmless (see Lemma 7.31). Because our substitutions are naive, we require that each substitution’s range be disjoint from the values being examined (recall that \bowtie non-deterministically chooses the substitutions’ ranges). Without this requirement, we would have $(\text{let* } ((x \ 7)) \ x) =_R (\text{let* } ((y \ 7)) \ a) : Expr$, witnessed by $\sigma_0 = \{\{x, a\}\}$ and $\sigma'_0 = \{\{y, a\}\}$.

$$\begin{array}{c}
 _ =_R _ : _ \subseteq \text{Value} \times \text{Value} \times \text{Type} \\
 \\
 \frac{}{a =_R a' : \text{BAtom}} \text{R}\alpha\text{-BATOM} \qquad \frac{}{a =_R a : \text{RAtom}} \text{R}\alpha\text{-RATOM} \\
 \\
 \frac{\forall i. v_i \bowtie v'_i : \tau_i \rightarrow \sigma_i \bowtie \sigma'_i \quad \forall i, j. \text{rng}(\sigma_i), \text{rng}(\sigma'_i) \# v_j, v'_j \quad \forall i \neq j. \text{rng}(\sigma_i) \# \text{rng}(\sigma_j) \quad \forall i. \llbracket \beta_i \rrbracket (\sigma_j)_j(v_i) =_R \llbracket \beta_i \rrbracket (\sigma'_j)_j(v'_i) : \tau_i}{\mathbf{prod}(v_i)_i =_R \mathbf{prod}(v'_i)_i : \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)} \text{R}\alpha\text{-PROD}
 \end{array}$$

In our ongoing `let*` example, the appropriate substitution is a no-op on the *LetStarClauses*, which import nothing, but recursive application of $=_R$ will discover their shared binding structure and compare them as equal. On the other hand, the *Expr* bodies will be both transformed into `(display cc)`, which lacks binding structure, and is naively equal to itself. So, the two expressions are $=_R$. (They are also trivially $=_B$, and therefore $=_{\alpha}$.) A complete derivation of their reference-equivalence, including analyzing the *LetStarClauses* themselves, is available in the supplementary material.

An example that better demonstrates the complexities of renaming is the event handler example from Section 2.1.1. The result of invoking \bowtie on each pair of children, in order to compare the two versions for α -equivalence, is in Figure 4.

Because the corresponding subterms import nothing, $\beta_0 = \beta_1 = \beta_3 = \emptyset$. β_2 is $0 \uplus 1$ and β_4 is $0 \uplus 3$. The result of performing those substitutions is shown in Figure 5, establishing the relationship between `mouse-evt` and the first `b` and the relationship between `kbd-evt` and the second `b`.

$$\begin{aligned}
\llbracket 0 \uplus 1 \rrbracket (\sigma_j)_j(\text{handle-evt gui-elt mouse-evt}) &= (\text{handle-evt gg mm}) \\
\llbracket 0 \uplus 1 \rrbracket (\sigma'_j)_j(\text{handle-evt a b}) &= (\text{handle-evt gg mm}) \\
\llbracket 0 \uplus 3 \rrbracket (\sigma_j)_j(\text{put-tag gui-elt (text-of kbd-evt)}) &= (\text{put-tag gg (text-of kk)}) \\
\llbracket 0 \uplus 3 \rrbracket (\sigma'_j)_j(\text{put-tag a (text-of b)}) &= (\text{put-tag gg (text-of kk)})
\end{aligned}$$

Fig. 5. Result of substitution in the handler example (all other subterms are trivial).

4 Romeo

Romeo is a first-order, typed, side-effect-free language whose values are abstract syntax trees. It uses types to direct the interpretation of these trees as syntax trees with binding, and to direct the execution of expressions in a way that respects that binding structure. We divide the task of achieving safety into three parts:

- First, the execution semantics ensures that whenever the program causes a name to escape the context in which it is defined, a `FAULT` is produced.
- Second, based on the non-escape property, we prove that at any point in execution, the dynamic environment could be replaced by one with α -equivalent values, and execution would still proceed to a value α -equivalent to what it otherwise would have. Execution is deterministic up to α : that is, the non-deterministic choices that are made (e.g. for fresh identifiers) do not change the α -equivalence class of the result.
- Last, we provide a deduction system (see Section 8) to generate proof obligations which, if satisfied, guarantee that escape (and thus, `FAULT`) will never occur.

The syntax of Romeo is given as follows:

$$\begin{aligned}
p \in \text{Prog} &::= fD \dots e : \tau \\
fD \in \text{FnDef} &::= (\mathbf{define-fn} (f x : \tau \dots \mathbf{pre} C) : \tau \mathbf{e post} C) \\
e \in \text{Expr} &::= (f x \dots) \\
&| (\mathbf{fresh} x \mathbf{in} e) \\
&| (\mathbf{let} x \mathbf{where} C \mathbf{be} e \mathbf{in} e) \\
&| (\mathbf{case} x (x e) (x e)) \\
&| (\mathbf{open} x (x \dots) e) \\
&| (\mathbf{if} x \mathbf{equals} x e e) \\
&| e^{\text{qlit}} \\
e^{\text{qlit}} \in \text{QuasiLit} &::= x \\
&| (\mathbf{ref} x) \\
&| (\mathbf{inj}_0 e^{\text{qlit}} \tau) \\
&| (\mathbf{inj}_1 \tau e^{\text{qlit}}) \\
&| \left(\mathbf{prod}^{\uparrow\beta} \left(e_i^{\text{qlit}} \downarrow \beta_i \right)_i \right)
\end{aligned}$$

Here, C ranges over a language of invariants from which the proof obligations for static safety are constructed. Romeo's operational semantics does not refer to these invariants. This sublanguage is discussed in Section 8.

Typechecking is largely straightforward. In the body of **open**, the variables $x \dots$ are given the types of the subterms of the scrutinee x , and in the body of **fresh**, x is bound to a name that is distinct from all other names in the execution environment. In order to use that name as a reference, the **ref** form takes an argument of type **BAtom** and returns it as a **RAtom**. A complete definition of the typechecking judgment can be found in the supplementary material.

We annotate injections with the types of the arm-not-taken, and product constructors with their binding structure. This allows us to synthesize the types of expressions with a function $\text{typeof}(\Gamma, e)$ whose definition is routine.

To simplify the deduction system, we require variables in some places where expressions would be more natural (like function arguments or x_{obj} in **open**). As a result, programs are written in (roughly) Λ -normal form, naming intermediate results with **let**.

4.1 Operational semantics

We define Romeo’s execution in big-step style. An advantage, for our purposes, of the big-step style is that it allows us to simultaneously enforce constraints about the return values of and about the names generated by the **fresh** and **open** forms.

We begin with some auxiliary definitions that we will need

$$\begin{aligned} w \in \text{Result} & ::= v \mid \text{FAULT} \\ \rho \in \text{ValEnv} & ::= \epsilon \mid \rho[x \rightarrow v] \\ \Gamma \in \text{TypeEnv} & ::= \epsilon \mid \Gamma, x:\tau \\ \text{fa}_{\text{env}}(\Gamma, \rho) & = \bigcup_{x \in \text{dom}(\Gamma)} \text{fa}(\Gamma(x), \rho(x)) \end{aligned}$$

The form of the execution judgment is

$$\Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} w$$

The k argument indicates the number of execution steps taken to produce the result in question.

4.2 Execution rules

We can now give the rules for execution in Romeo. Rules that introduce names come in two forms, -OK, and -FAIL. In each case, the difference is that **FAULT** occurs in the -FAIL case. A **FAULT** indicates that a name has escaped the scope that created it (E-FRESH-*) or exposed it (E-OPEN-*). Much of the rest of the machinery in those rules is about ensuring that newly introduced names do not collide with each other or with names in the environment.

Some execution rules will depend on the type environment Γ . This is because the binding structures of values are represented in their types (τ), but not in their runtime representations (v). Therefore, type erasure is not possible — the meaning of values (and thus the behavior of those rules) depends on type information.

$$\frac{a \notin \text{fa}_{\text{env}}(\Gamma, \rho) \quad \Gamma, x:\text{BAtom} \vdash_{\text{exe}} \langle e, \rho[x \rightarrow a] \rangle \xrightarrow{k} w \quad \tau = \text{typeof}((\Gamma, x:\text{BAtom}), e) \quad w = \text{FAULT} \vee a \notin \text{fa}(\tau, w)}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{fresh} \ x \ \mathbf{in} \ e), \rho \rangle \xrightarrow{k+1} w} \text{E-FRESH-OK}$$

$$\frac{a \notin \text{fa}_{\text{env}}(\Gamma, \rho) \quad \Gamma, x:\text{BAtom} \vdash_{\text{exe}} \langle e, \rho[x \rightarrow a] \rangle \xrightarrow{k} w \quad \tau = \text{typeof}((\Gamma, x:\text{BAtom}), e) \quad w \neq \text{FAULT} \wedge a \in \text{fa}(\tau, w)}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{fresh} \ x \ \mathbf{in} \ e), \rho \rangle \xrightarrow{k+1} \text{FAULT}} \text{E-FRESH-FAIL}$$

We begin with the rules for evaluating **fresh** expressions. The rules require that the new name not occur in the environment ρ . Our determinacy theorems (Theorems 7.1 and 7.2) guarantee that the choice of the new name will not affect the result (up to α -equivalence). We have two versions of the rule: **FRESH-FAIL**, which returns **FAULT** when the new name appears free in the result of executing the body e , and **FRESH-OK**, which returns w when that is not the case.

During the execution of e , x is treated as a **BAtom**. It is convertible to a **RAtom** by the expression (**ref** x).

The hypothesis $\tau = \text{typeof}((\Gamma, x:\text{BAtom}), e)$ is needed to determine whether to produce **FAULT** or not. Determining the type, of course, is entirely static and could be pre-computed once rather than at each evaluation.

$$\frac{\rho(x_{\text{obj}}) =_{\alpha} \mathbf{prod}(v_i)_i : \tau_{\text{obj}} \quad \tau_{\text{obj}} = \Gamma(x_{\text{obj}}) \quad \tau_{\text{obj}} = \text{Prod}^{\hat{\alpha}\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i \quad \text{fa}_{\text{env}}(\Gamma, \rho) \vdash_{\text{suff-disj}} \mathbf{prod}(v_i)_i : \tau_{\text{obj}} \quad \Gamma, (x_i : \tau_i)_i \vdash_{\text{exe}} \langle e, \rho[x_i \rightarrow v_i]_i \rangle \xrightarrow{k} w \quad \tau = \text{typeof}(\Gamma, e) \quad w = \text{FAULT} \vee \text{xa}(\tau_{\text{obj}}, \mathbf{prod}(v_i)_i) \# \text{fa}(\tau, w)}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{open} \ x_{\text{obj}} \ ((x_i)_i) \ e), \rho \rangle \xrightarrow{k+1} w} \text{E-OPEN-OK}$$

$$\frac{\rho(x_{\text{obj}}) =_{\alpha} \mathbf{prod}(v_i)_i : \tau_{\text{obj}} \quad \tau_{\text{obj}} = \Gamma(x_{\text{obj}}) \quad \tau_{\text{obj}} = \text{Prod}^{\hat{\alpha}\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i \quad \text{fa}_{\text{env}}(\Gamma, \rho) \vdash_{\text{suff-disj}} \mathbf{prod}(v_i)_i : \tau_{\text{obj}} \quad \Gamma, (x_i : \tau_i)_i \vdash_{\text{exe}} \langle e, \rho[x_i \rightarrow v_i]_i \rangle \xrightarrow{k} w \quad \tau = \text{typeof}(\Gamma, e) \quad w \neq \text{FAULT} \wedge \neg(\text{xa}(\tau_{\text{obj}}, \mathbf{prod}(v_i)_i) \# \text{fa}(\tau, w))}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{open} \ x_{\text{obj}} \ ((x_i)_i) \ e), \rho \rangle \xrightarrow{k+1} \text{FAULT}} \text{E-OPEN-FAIL}$$

The next pair of rules destructure a product. Given a value $\rho(x_{\text{obj}}) = \mathbf{prod}(v_{\text{obj},0}, \dots, v_{\text{obj},n})$, the **open** expression chooses an α -variant $\mathbf{prod}(v_0, \dots, v_n)$ and binds the resulting pieces to the variables x_i . The free names in the α -variant must be distinct both from names in the environment ρ and from each other (to the extent that they are not actually related by binding). This is tested by a subsidiary judgment $\vdash_{\text{suff-disj}}$, discussed in Section 4.2.2. Choosing an α -variant that satisfies this requirement is the subject of Section 5.

As with **fresh**, we have two rules which branch on whether any of the new names appear free in the result of the body e . We test for escaped names by comparing the free atoms in the result with the exportable atoms of the renamed input. This suffices for safety (α -equivalence preservation) because the set of atoms free in the environment for e but not in the environment for $(\mathbf{open} \ x_{\text{obj}} \ ((x_i)_i) \ e)$ (the atoms

that “become free” in this execution step) is equal to the set $\text{xa}(\tau_{\text{obj}}, \mathbf{prod}(v_i)_i)$.

$$\frac{\Gamma \vdash_{\text{exe}} \langle e_{\text{val}}, \rho \rangle \xRightarrow{k_{\text{val}}} v_{\text{val}} \quad \tau_{\text{val}} = \text{typeof}(\Gamma, e_{\text{val}}) \quad \Gamma, x:\tau_{\text{val}} \vdash_{\text{exe}} \langle e_{\text{body}}, \rho [x \rightarrow v_{\text{val}}] \rangle \xRightarrow{k_{\text{body}}} w}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{let} \ x \ \mathbf{where} \ C \ \mathbf{be} \ e_{\text{val}} \ \mathbf{in} \ e_{\text{body}}), \rho \rangle \xRightarrow{k_{\text{val}}+k_{\text{body}}+1} w} \text{E-LET}$$

$$\frac{\Gamma \vdash_{\text{exe}} \langle e_{\text{val}}, \rho \rangle \xRightarrow{k} \text{FAULT}}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{let} \ x \ \mathbf{where} \ C \ \mathbf{be} \ e_{\text{val}} \ \mathbf{in} \ e_{\text{body}}), \rho \rangle \xRightarrow{k+1} \text{FAULT}} \text{E-LET-FAIL}$$

There are two evaluation rules for **let**, depending on whether calculating e_{val} faults. As noted above, the constraint C is ignored at run-time.

$$\frac{\text{body}(f) = e \quad \text{formals}(f) = (x_{\text{formal},i}:\tau_{\text{formal},i})_i \quad (x_{\text{formal},i}:\tau_{\text{formal},i})_i \vdash_{\text{exe}} \langle e, [x_{\text{formal},i} \rightarrow \rho_i(x_{\text{actual},i})]_i \rangle \xRightarrow{k} w}{\Gamma \vdash_{\text{exe}} \langle (f \ (x_{\text{actual},i})_i), \rho \rangle \xRightarrow{k+1} w} \text{E-CALL}$$

As in Pure FreshML (Pottier, 2007b), we assume that our expressions are evaluated in a context of function definitions, so that from a function name, we can retrieve the function’s formals and body. Since this context is constant throughout an execution, it is elided in the evaluation judgment.

$$\frac{\rho(x_1) = a \quad \rho(x_r) = b \quad a = b \quad \Gamma \vdash_{\text{exe}} \langle e_0, \rho \rangle \xRightarrow{k} w}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{if} \ x_1 \ \mathbf{equals} \ x_r \ e_0 \ e_1), \rho \rangle \xRightarrow{k+1} w} \text{E-IF-YES}$$

$$\frac{\rho(x_1) = a \quad \rho(x_r) = b \quad a \neq b \quad \Gamma \vdash_{\text{exe}} \langle e_1, \rho \rangle \xRightarrow{k} w}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{if} \ x_1 \ \mathbf{equals} \ x_r \ e_0 \ e_1), \rho \rangle \xRightarrow{k+1} w} \text{E-IF-NO}$$

$$\frac{\rho(x_{\text{obj}}) = \mathbf{inj0}(v_0) \quad \Gamma(x_{\text{obj}}) = \tau_0 + \tau_1 \quad \Gamma, x_0:\tau_0 \vdash_{\text{exe}} \langle e_0, \rho [x_0 \rightarrow v_0] \rangle \xRightarrow{k} w}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{case} \ x_{\text{obj}} \ (x_0 \ e_0) \ (x_1 \ e_1)), \rho \rangle \xRightarrow{k+1} w} \text{E-CASE-LEFT}$$

$$\frac{\rho(x_{\text{obj}}) = \mathbf{inj1}(v_1) \quad \Gamma(x_{\text{obj}}) = \tau_0 + \tau_1 \quad \Gamma, x_1:\tau_1 \vdash_{\text{exe}} \langle e_1, \rho [x_1 \rightarrow v_1] \rangle \xRightarrow{k} w}{\Gamma \vdash_{\text{exe}} \langle (\mathbf{case} \ x_{\text{obj}} \ (x_0 \ e_0) \ (x_1 \ e_1)), \rho \rangle \xRightarrow{k+1} w} \text{E-CASE-RIGHT}$$

The remainder of the rules are routine. For simplicity’s sake, the equality test construct works only on atoms. But in Romeo, for any type, it is straightforward to write an equality test out of **equals**, and that predicate will necessarily be an α -equivalence.

$$\frac{\epsilon \vdash_{\text{exe}} \langle e, \epsilon \rangle \xRightarrow{k} v}{\vdash_{\text{exe}} fD \dots e \xRightarrow{k+1} v} \text{E-PROG}$$

The E-PROG rule initiates evaluation of a program. It is notionally responsible for setting up the function context, which we omit from our notation, as it is otherwise constant.

4.2.1 Quasi-Literals

The last kind of Romeo expression is the *quasi-literals*, so called because they look like literal syntax for object-level syntax objects, except that they contain variable references (which denote values), not literal atoms. Of course, those variables may refer to atom values generated by **fresh**. Quasi-literals also contain some type information to make type synthesis possible.

$$\frac{v = \llbracket e^{\text{qlit}} \rrbracket_{\rho}}{\Gamma \vdash_{\text{exe}} \langle e^{\text{qlit}}, \rho \rangle \xRightarrow{1} v} \text{E-QLIT}$$

Their evaluation is routine, and is specified by the following rules:

$$\begin{aligned} \llbracket _ \rrbracket_{\rho} &: \text{QuasiLit} \times \text{ValEnv} \rightarrow \text{Value} \\ \llbracket x \rrbracket_{\rho} &\triangleq \rho(x) \\ \llbracket (\text{ref } x) \rrbracket_{\rho} &\triangleq \rho(x) \\ \llbracket (\text{inj}_0 e^{\text{qlit}} \tau) \rrbracket_{\rho} &\triangleq \text{inj}_0(\llbracket e^{\text{qlit}} \rrbracket_{\rho}) \\ \llbracket (\text{inj}_1 \tau e^{\text{qlit}}) \rrbracket_{\rho} &\triangleq \text{inj}_1(\llbracket e^{\text{qlit}} \rrbracket_{\rho}) \\ \llbracket (\text{prod}^{\uparrow \beta_{\text{exp}}} (e_i^{\text{qlit}} \downarrow \beta_i)_i) \rrbracket_{\rho} &\triangleq \text{prod} \left(\llbracket e_i^{\text{qlit}} \rrbracket_{\rho} \right)_i \end{aligned}$$

4.2.2 Sufficient disjointness

The requirement that evaluation be insensitive to α -equivalent inputs leads to strong requirements on the way that **open** destructures values. Consider the **let*** example from before

$$\begin{array}{ll} (\text{let* } ((a \ 1) & (\text{let* } ((d \ 1) \\ (b (+ a a)) & (d (+ d d)) \\ (c (* b 5))) & (d (* d 5))) \\ (\text{display } c)) & (\text{display } d)) \end{array}$$

These are α -equivalent, but if they were each destructured without renaming, we would have $d = d = d$, even though $a \neq b \neq c$, violating our goal of being indifferent to α -conversion. Therefore, E-OPEN potentially needs to freshen each binder to a distinct new name, e.g.

$$\begin{array}{l} (\text{let* } ((aa \ 1) \\ (bb (+ aa aa)) \\ (cc (* bb 5))) \\ (\text{display } cc)) \end{array}$$

The rule to ensure this is that, before destructuring, we must α -convert values so that the binders exposed by destructuring are disjoint from each other and from

any names that appear in the environment. This gives rise to the hypothesis

$$\text{fa}_{\text{env}}(\Gamma, \rho) \vdash_{\text{suff-disj}} \mathbf{prod}(v_i)_i : \tau_{\text{obj}}$$

in the E-OPEN- \star rules.

To calculate $\vdash_{\text{suff-disj}}$, we need the judgment $\vdash_{\text{bndrs-disj}} v : \tau$, which checks that the exported binders in v (as determined by the type τ) are disjoint from each other.

$$\frac{}{\vdash_{\text{bndrs-disj}} a : \text{BAtom}} \text{BD-BATOM} \qquad \frac{}{\vdash_{\text{bndrs-disj}} a : \text{RAtom}} \text{BD-RATOM}$$

$$\frac{\forall i, j \hat{\in} \beta_{\text{ex}}. i \neq j \Rightarrow \text{fb}(\tau_i, v_i) \# \text{fb}(\tau_j, v_j) \quad \forall i \hat{\in} \beta_{\text{ex}}. \vdash_{\text{bndrs-disj}} v_i : \tau_i}{\vdash_{\text{bndrs-disj}} \mathbf{prod}(v_i)_i : \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i} \text{BD-PROD}$$

We can now define $\vdash_{\text{suff-disj}}$, which checks the disjointness of *non*-exported subterms (because these are the binders that will become free after destructuring), and also that those names (calculable by xa) are disjoint from a set A of atoms (in practice, this is the set of free atoms in the environment).

$$\frac{\forall i \hat{\notin} \beta_{\text{ex}}. \vdash_{\text{bndrs-disj}} v_i : \tau_i \quad \forall i, j \hat{\notin} \beta_{\text{ex}}. i \neq j \Rightarrow \text{fb}(\tau_i, v_i) \# \text{fb}(\tau_j, v_j) \quad \text{xa}(\mathbf{prod}(v_i)_i, \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i) \# A}{A \vdash_{\text{suff-disj}} \mathbf{prod}(v_i)_i : \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i} \text{SUFF-DISJ}$$

4.3 Examples

4.3.1 Translation of $\text{let}\star$

For an example, we write code that translates between two languages: from the lambda calculus augmented with a $\text{let}\star$ construct into the plain lambda calculus.

Our code, in Figure 2, mentions types defined in Figure 1. It is written in Romeo-L (Muehlboeck, 2013), which is a friendlier front-end to Romeo. For our purposes, the important differences are that the arguments to function calls and the scrutinees of **open** and **case** may be arbitrary expressions (not just variable references), and that Romeo-L can infer the constraint C of **let**, so we may omit it. Furthermore, it will turn out (see Section 8.3) that our example needs no pre- or post-conditions from *convert* to show the absence of **FAULT**, so those constraints are also omitted.

Additionally, we have written our code using more readable n -way sum types. This means that our **case** construct can branch four ways depending on whether the *Expr* it examines is a variable reference, an application, a lambda abstraction, or a let-star statement, and that injections take (as a subscript) a description of the choice that they are constructing.

Lines 3–5 are straightforward traversal of the existing *Expr* forms that are already forms in the core language (but, since their subterms might not be, they still need to be converted by recursively invoking *convert*).

Lines 6–11 destructure $\text{let}\star$ forms and handle the trivial case, where the $\text{let}\star$ does not have any arms. In the case where $\text{let}\star$ has at least one arm, line 12 constructs a smaller $\text{let}\star$ with one fewer arm, and recursively converts it, calling

the result *e-rest*. Finally, lines 13–14 construct a beta-redex in the object language to bind the first arm’s name to its value expression in *e-rest*.

4.3.2 Translation of `let`

Consider again the example above. Suppose that we had implemented a normal `let` construct (where names from previous arms are not in scope for the later arms), with the type:

$$\begin{aligned} \text{LetClauses} ::= & \text{Prod} () \\ & | \text{Prod}^{\uparrow 1 \triangleright 0} (\text{Prod}^{\downarrow 0} (\text{BAtom}, \text{Expr}), \text{LetClauses}) \end{aligned}$$

The only difference, besides the name, is that the recursive *LetClauses* does not have a $\downarrow 0$. If we had wanted to change the code in Figure 2 to expand ordinary lets instead, the above change to the type of *Expr* is sufficient, and the otherwise identical code would respect *LetClause*’s binding behavior and correctly expand the `let` construct! This is a consequence of Theorem 7.1, which ensures that programs cannot observe anything about names except their binding structure, as defined by their binding specifications.

5 Freshening

Executing the E-OPEN- \star rules in a Romeo implementation requires the ability to take an arbitrary product value $\mathbf{prod}(v_i)_i$, and generate a new α -equivalent value $\mathbf{prod}(v'_i)_i$ such that $A \vdash_{\text{suff-disj}} \mathbf{prod}(v'_i)_i : \tau$ (for a particular A). We call this process “freshening”.

5.1 Approach

Conceptually, the $\vdash_{\text{suff-disj}}$ predicate (motivated and defined in Section 4.2.2) requires that all exposable binders in a value be mutually disjoint (and disjoint from A). This is easy to achieve: simply assign fresh names to each exposable binder. However, in order to produce a value that is also α -equivalent to the input, we must also rename all references that refer to those values.

The difficulty in this is best illustrated by an example. Suppose that we have started freshening, and we’ve chosen to rename the binder y to yy in the following binding form, and thus, to maintain α -equivalence, must rename all of its references:

```
(let* ((x (string-length y))
      (y 5)
      (z (+ y 1)))
      (+ y 2))
```

Because the binder y is both imported (into $(z (+ y 1))$) and exported (out to the `let*` and then imported into $(+ y 2)$), the substitution $\{\langle y, yy \rangle\}$ must be applied to multiple places. However, applying the substitution to the whole value is unacceptable because the free reference y in $(\text{string-length } y)$ must *not* be renamed.

Therefore, after generating a fresh replacement for y , we must use \downarrow (import) designations in the type of the value to determine where that particular y is in scope, and only rename references to y in those locations, in order to keep references in sync with the changes we made to the binders.

5.2 Reference renaming

While the freshening operation only affects *non-free* names in the value we are freshening, we will accomplish it by renaming the *free* names of its subterms (and sub-subterms, etc.). This is because those are the non-free names whose meaning is determined by imports from other subterms. To that end, we need an operation that applies a renaming only to the free references of a term. It needs to know the type of its value argument so that it can avoid touching binders and bound names. With that information, the implementation is straightforward.

$$\begin{aligned} _|_r(-:_) &: \text{Subst} \times \text{Value} \times \text{Type} \rightarrow \text{Value} \\ \sigma|_r(a:\text{BAtom}) &\triangleq a \\ \sigma|_r(a:\text{RAtom}) &\triangleq \sigma(a) \\ \sigma|_r(\mathbf{prod}(v_i) : \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i) &\triangleq \mathbf{prod}((\sigma \setminus A_i)|_r(v_i : \tau_i))_i \\ &\quad \text{where } \forall i. A_i \triangleq \llbracket \beta_i \rrbracket (\text{fb}(v_j, \tau_j))_j \end{aligned}$$

5.3 Implementation of freshening

Let v_{top} denote the term we are freshening, and v'_{top} denote the result of freshening.

We begin by defining freshen-subterm, which performs freshening on a subterm of v_{top} . It takes a value v (which is assumed to be a subterm of v_{top}), the type τ of that value, and a boolean parameter, *exported-from-v-top?*, which is true iff the exported binders from v are exported in an unbroken chain all the way out of v_{top} .

The function freshen-subterm applied to v returns a value v' and a substitution σ such that:

- $A \vdash_{\text{suff-disj}} v' : \tau$,³
- $v =_{\text{R}} v' : \tau$ (i.e. the structure of bound atoms and the free references are unchanged),
- $\sigma(v) =_{\text{B}} v' : \tau$ (i.e. the exported binders differ by σ),
- $\text{dom}(\sigma) = \text{fb}(v, \tau)$,⁴ and
- if *exported-from-v-top?* is true, $v =_{\text{B}} v' : \tau$ (i.e. the exported binders are identical, and thus $v =_{\alpha} v' : \tau$ and σ is a no-op),

³ For the sake of correctness, it is only necessary for v'_{top} (not all of its subterms, sub-subterms, etc.) to be sufficiently-disjoint. At the cost of implementation complexity, it is possible to optimize this algorithm so that it only freshens exposable names.

⁴ This means that, even when σ is a no-op substitution, it is still not an empty set. This is because of shadowing: If an import brought a particular name in twice, with a non-freshened version of it shadowing a freshened version of it, it is necessary that the non-freshened version correctly override the freshened version when the substitutions are combined by $\llbracket _ \rrbracket$.

- but if *exported-from-v-top?* is false, the exported binders of v' are fresh.

We omit the mechanics of threading environments of names through freshen-subterm, as they are routine. Informally speaking, “where aa is fresh” means that all generated names must be distinct from A and from each other.

The base cases of our recursive freshening function are all straightforward. References do not require any action at this level, but binders are freshened, provided they are not exported from v_{top} . This ensures that $v_{\text{top}} =_{\text{B}} v'_{\text{top}} : \tau_{\text{top}}$.

Either way, we also return a substitution that reflects the change (if any) in the exported binders of the result.

$$\begin{aligned} \text{freshen-subterm}(\cdot : \cdot, \cdot) & : \text{Value} \times \text{Type} \times \text{Bool} \rightarrow \text{Value} \times \text{Subst} \\ \text{freshen-subterm}(a : \text{RAtom}, \cdot) & \triangleq a, \emptyset \\ \text{freshen-subterm}(a : \text{BAtom}, \mathbf{true}) & \triangleq a, \{\langle a, a \rangle\} \\ \text{freshen-subterm}(a : \text{BAtom}, \mathbf{false}) & \triangleq aa, \{\langle a, aa \rangle\}, \text{ where } aa \text{ is fresh} \end{aligned}$$

Given a wide product $\mathbf{prod}(v_i)_i$, we first recursively call freshen-subterm on all of the subterms v_i , producing values v'_i whose binders have been renamed according to σ_i . Observe that v_i exports from v_{top} iff $\mathbf{prod}(v_i)_i$ does and v_{top} exports v_i .

Then, everywhere subterm j is imported, we rename free references according to σ_j . We do this by calculating, for each subterm i , the value $\llbracket \beta_i \rrbracket(\sigma_j)_j$.

Finally, we have to determine what substitution corresponds to the difference in exported binders between $\mathbf{prod}(v_i)_i$ and the value we return. This is calculated by $\llbracket \beta_{\text{ex}} \rrbracket(\sigma_j)_j$.

$$\begin{aligned} \text{freshen-subterm}(\mathbf{prod}(v_i)_i : \text{Prod}^{\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i, \text{exported-from-v-top?}) \\ \triangleq \mathbf{prod} \left(\llbracket \beta_i \rrbracket(\sigma_j)_j \Big|_{\Gamma} (v'_i : \tau_i) \right)_i, \quad \llbracket \beta_{\text{ex}} \rrbracket(\sigma_j)_j \\ \text{where } \forall j. (v'_j, \sigma_j) \triangleq \text{freshen-subterm}(v_j : \tau_j, \text{exported-from-v-top?} \wedge (i \hat{\in} \beta_{\text{ex}})) \end{aligned}$$

Using freshen-subterm to turn v_{top} to v'_{top} is now simple. We generate a sufficiently disjoint yet α -equivalent value to v_{top} by using freshen-subterm with *exported-from-v-top?* set to **true**, as everything exported from v_{top} is (tautologically) exported from v_{top} . We discard the resulting substitution (but by the invariants above, we know it to be a no-op whose domain is $\text{fb}(v_{\text{top}}, \tau_{\text{top}})$).

$$\begin{aligned} \text{freshen}(\cdot : \cdot) & : \text{Value} \times \text{Type} \rightarrow \text{Value} \\ \text{freshen}(v_{\text{top}} : \tau_{\text{top}}) & \triangleq v'_{\text{top}} \\ & \text{where } v'_{\text{top}}, - \triangleq \text{freshen-subterm}(v_{\text{top}} : \tau_{\text{top}}, \mathbf{true}) \end{aligned}$$

Thus, by the (conjectured) invariants above (and recalling that the informal treatment of freshness means that freshening doesn't explicitly examine A), we have:

Conjecture 5.1 (Freshening is always possible)

If $v' = \text{freshen}(v : \tau)$, then $v =_{\alpha} v' : \tau$ and $A \vdash_{\text{suff-disj}} v' : \tau$.

This shows that it is possible to construct a value $\mathbf{prod}(v_i)_i$ that satisfies the $\vdash_{\text{suff-disj}}$ and $=_{\alpha}$ premises of E-OPEN-*, and thus, it is possible to execute Romeo programs.

6 Romeo respects α -equivalence: a guide to the proofs

We are now ready to prove our main theorem: that Romeo respects α -equivalence. Romeo is non-deterministic in its choice of names in the E-FRESH- \star and E-OPEN- \star rules. This complicates the proof of respecting α -equivalence. In a system with simpler binding structures, like Pure FreshML, the proof would go as follows:

- Define all the states of the machine in Section 4 to be α -equivalence classes.
- Rely on the freshness condition for binders (Pouillard & Pottier, 2010) to show that each manipulation on machine states (defined in terms of α -equivalence-class representatives) respects α -equivalence.

In such a system, we would have needed to prove only Lemma 8.1, as does Pottier (2007a).

Unfortunately, we were unable to usefully model our complex binding structures, especially buried bindings, in nominal logic. Therefore, we were forced to proceed from first principles.

We show two results: First, that if the evaluations of an expression in two α -equivalent environments both terminate, then their results are α -equivalent, and second, if the evaluation of an expression in one of two α -equivalent environments yields a result, then evaluating the expression in the other one must yield at least one α -equivalent result as well. Since we are using big-step semantics, we cannot talk directly about non-termination.

For each of these theorems, the vast majority of the complexity is contained in the cases for E-FRESH- \star and E-OPEN- \star .

Since we must account for faulting, we extend the definition of α -equivalence to assert that $\text{FAULT} =_{\alpha} \text{FAULT}$.

Complete proofs of all the theorems and lemmas we mention in this paper are available in the supplementary material. Here, we discuss the most interesting proofs. It is typical for reasoning of this complexity to be done in a mechanical reasoning system like Coq, but because we initially underestimated its size and scope, our proof is entirely handwritten.

6.1 Executions on α -equivalent environments yield α -equivalent results

The first, and hardest, part of proving Romeo's soundness is Theorem 7.1, which shows that two terminating executions of α -equivalent environments yield α -equivalent result values.

The major problem in the proof is that the two executions will potentially generate different fresh names in E-FRESH- \star and E-OPEN- \star . Hence, even though the environments start out α -equivalent, they will not stay α -equivalent throughout execution. For example, a **fresh** statement non-deterministically introduces a new name into the environment. Therefore, we must generalize our induction hypothesis to account for the ways in which ρ and ρ' diverge from α -equivalence.

We account for this divergence with two injective substitutions that unify the names the two executions introduce. So our induction hypothesis says that if $\sigma \circ \rho =_{\alpha}$

$\sigma' \circ \rho' : \Gamma$ for a pair of injective substitutions σ and σ' , then the results will be α -equivalent, modulo the same transformation (i.e. $\sigma(w) =_{\alpha} \sigma'(w') : \tau$).

Consider the case of E-FRESH-*. As we enter the scope of the new name it generates, σ and σ' are extended to map the new names to a common fresh name (lines 7.49–7.50 in the proof in Section 7.3), for the sake of the induction hypothesis.

The induction hypothesis tells us that the results (of the recursive evaluation of e) are equivalent modulo the *extended* substitutions (line 7.55). The result of the **fresh** evaluation step is either the same as that of the recursive step, or **FAULT**. We first show that the *original* substitutions suffice to α -equate those two values (lines 7.61–7.62), and then that one side faults if and only if the other side does (lines 7.57–7.58).

The E-OPEN-* case proceeds with a similar structure.⁵ However, in this case, we are not generating a single pair of new names, but unpacking a pair of values, which potentially contain many names. The crucial lemma to handle this (Lemma 7.42, used on lines 7.74–7.77, described in Section 6.1.1) states that the $\vdash_{\text{suff-disj}}$ predicate in Romeo’s execution rules is strong enough that the technique from the E-FRESH-* rules works for **open**, even though the various subterms of the value being destructured have potentially different scopes. In a sense, the lemma is where the complex binding structures meet the binding-safe programming in our system, and it is discussed in more detail in Section 6.1.1. After the induction hypothesis, E-OPEN-* proceeds like E-FRESH-*.

The E-IF-* case, though simple, is crucial, because it shows that our induction hypothesis is strong enough to guarantee that a comparison between two names in ρ will always have the same result as a comparison between two names in ρ' . This is where the injectivity of the substitutions σ and σ' is used.

6.1.1 One substitution, not one per subterm, suffices for opening

Lemma 7.42 is a crucial part of our project to extend binding-safe programming to support complex binding structures. A binding form in our system may possess many scopes, with different meanings for the same names. However, to programmatically manipulate such a binding form, it must be destructured, which dumps all of those differently scoped subterms into the same execution environment.

This lemma is our most interesting technical trick, showing that, if $A \vdash_{\text{suff-disj}} \mathbf{prod}(v_i)_i : \tau$ holds (where A is the set of atoms in the environment), destructuring is safe. For the purposes of Theorem 7.1, “destructuring is safe” means that a single pair of substitutions suffices to bring all the subterms “back” into α -equivalence (in other words, it will not lose information by causing a name collision).

⁵ In fact, **open** can be used to replace **fresh** (whenever at least one binding form exists in the environment). We keep both constructs because we found it easiest to understand the E-OPEN-* cases of our proofs a generalization of the E-FRESH-* cases.

The proof of this lemma starts on line 7.13, which generates, not a single pair of substitutions, but one pair for each subterm of the pair of values. Lines 7.16–7.17 adjust those substitutions to avoid free names in the environment, and lines 7.18–7.26 restrict those substitutions by removing any exported names, and show that the resulting substitutions (interpreted by β_i) still can achieve $=_R$ of pairs of subterms of the values.

Lines 7.27–7.29 show that all those substitutions can be combined to form a single pair of substitutions. Examining one side of that pair, lines 7.30–7.37 show that that single substitution is an adequate replacement for any of the substitutions generated on line 7.22.

Lines 7.40–7.41 show that, for exported subterms, their exclusion from the substitutions means that they remain $=_B$ even after being substituted.

Lines 7.42–7.43 show that, for subterms that are not exported, the overall pair of substitutions makes them $=_B$.

6.1.2 α -equivalence and free names are connected

Lemma 7.30 (along with its cases Lemmas 7.28 and 7.29, and its corollary, Lemma 7.31) is used extensively in our proof, as it links the concepts of free names and α -equivalence. It states that a name is free if and only if a (naive) substitution of that name leads to a non- α -equivalent value.

The case for binders (Lemma 7.29) is straightforward, but the case for references (Lemma 7.28) is much more complex. Its sublemma, lines 7.1–7.6, applies the induction hypothesis in the case where the name being substituted is not imported by the overall wide product in the current subterm. The rest of it is a case analysis between (a) a being a free reference in the value, and therefore also a free reference in one of its subterms (lines 7.7–7.8), (b) a being non-free, which turns into a case analysis for each subterm: (b.i) a not being imported (lines 7.11–7.12), and (b.ii) a being imported. This final sub-case does not use the sublemma, but instead shows that the substitutions we generated in line 7.10 prove reference-equivalence of the whole value.

6.2 Termination is representation-oblivious

Theorem 7.1 leaves open the possibility that for a ρ whose execution leads to a result w , execution of some α -variant ρ' on the same program might fail to terminate. However, Theorem 7.2 says that if one α -variant terminates (either with a value or `FAULT`), then the execution of every α -variant will terminate (and, by Theorem 7.1, when it does, the value will be α -equivalent to the result of the original). Recall that we cannot talk directly about non-termination in our big-step semantics.

Broadly speaking, our approach to proving Theorem 7.2 is that, for every choice of fresh name in the original computation, we choose the same name in the other one. This preserves α -equivalence of ρ and ρ' for the induction hypothesis.

7 Soundness of execution for α -equivalence: lemmas and key proofs

7.1 General definitions

7.1.1 Naive operations

First, we show some basic properties regarding our naïve operations. These follow from straightforward inductive arguments.

Lemma 7.1 (Substitutions never increase disjointness)

If $\sigma(A) \# \sigma(A')$, then $A \# A'$.

$$\begin{array}{ll} \text{supp}(_) : \text{Value} \rightarrow \text{AtomSet} & \text{supp}(_) : \text{Subst} \rightarrow \text{AtomSet} \\ \text{supp}(a) \triangleq \{a\} & \text{supp}(\sigma) \triangleq \text{dom}(\sigma) \cup \text{rng}(\sigma) \\ \text{supp}(\mathbf{inj0}(v)) \triangleq \text{supp}(v) & \\ \text{supp}(\mathbf{inj1}(v)) \triangleq \text{supp}(v) & \text{supp}(_) : \text{ValEnv} \rightarrow \text{AtomSet} \\ \text{supp}(\mathbf{prod}(v_i)_i) \triangleq \bigcup_i \text{supp}(v_i) & \text{supp}(\rho) \triangleq \bigcup_{x \in \rho} \text{supp}(\rho(x)) \end{array}$$

Lemma 7.2 (Support contains all relevant atoms)

$$\forall v, \tau. \quad \begin{array}{l} \text{fa}(\tau, v) \subseteq \text{supp}(v) \text{ and } \text{fb}(\tau, v) \subseteq \text{supp}(v) \\ \text{and } \text{fr}(\tau, v) \subseteq \text{supp}(v) \text{ and } \text{xa}(\tau, v) \subseteq \text{supp}(v) \end{array}$$

As noted in Section 2.3, we represent a substitution as a set of ordered pairs.

Lemma 7.3 (Substitutions can sometimes commute)

If $\text{dom}(\sigma) \# \text{dom}(\sigma')$ and $\text{dom}(\sigma) \# \text{rng}(\sigma')$ and $\text{dom}(\sigma') \# \text{rng}(\sigma)$, then $\sigma(\sigma'(v)) = \sigma'(\sigma(v))$.

7.1.2 Semantics of β

Here, we list some properties of our $\llbracket \beta \rrbracket (_)$ operation, especially the way that it relates to the substitutions or sets of atoms that it takes as input. All of these proofs are easy inductive arguments on β .

Lemma 7.4 (β of σ is consistent with its inputs)

$$\text{dom}(\llbracket \beta \rrbracket (\sigma_i)_i) = \bigcup_{i \hat{\in} \beta} \text{dom}(\sigma_i) \text{ and } \llbracket \beta \rrbracket (\sigma_i)_i \subseteq \bigcup_{i \hat{\in} \beta} \sigma_i$$

Lemma 7.5 (β of A is consistent with its inputs)

$$\llbracket \beta \rrbracket (A_i)_i = \bigcup_{i \hat{\in} \beta} A_i$$

Lemma 7.6 (β of σ with disjoint domains does not shadow anything.)

$$(\forall i, j \hat{\in} \beta. i \neq j \Rightarrow \text{dom}(\sigma_i) \# \text{dom}(\sigma_j)) \text{ iff } \llbracket \beta \rrbracket (\sigma_i)_i = \bigcup_{i \hat{\in} \beta} \sigma_i$$

Lemma 7.7 (β of σ ignores the σ s not mentioned by β)

$$\text{If } \forall i \hat{\in} \beta. \sigma_i = \sigma'_i, \text{ then } \llbracket \beta \rrbracket (\sigma_i)_i = \llbracket \beta \rrbracket (\sigma'_i)_i$$

Lemma 7.8 (β of injective, mutually-range-disjoint σ s has an injective result)

$$\text{If } \forall i. \text{inj}(\sigma_i), \text{ and } \forall i \neq j. \text{rng}(\sigma_i) \# \text{rng}(\sigma_j), \text{ then } \text{inj}(\llbracket \beta \rrbracket (\sigma_i)_i).$$

Lemma 7.9 (β of a set respects renamings of that set)

$$\llbracket \beta \rrbracket (A_i [a'/a])_i = \llbracket \beta \rrbracket (A_i)_i [a'/a]$$

Lemma 7.10 (β commutes with dom)

$$\text{dom}(\llbracket \beta \rrbracket (\sigma_i)_i) = \llbracket \beta \rrbracket (\text{dom}(\sigma_i))_i$$

7.2 Basic α -equivalence lemmas

Now, we show some basic properties related to α -equivalence and \bowtie . First, we need to define $=_{\text{shape}}$. Attempting to compare two values for α -equivalence is only interesting if they have the same shape.

$$v =_{\text{shape}} v' \subseteq \text{Value} \times \text{Value}$$

$$\frac{}{a =_{\text{shape}} a'} \qquad \frac{v =_{\text{shape}} v'}{\mathbf{inj0}(v) =_{\text{shape}} \mathbf{inj0}(v')} \qquad \frac{v =_{\text{shape}} v'}{\mathbf{inj1}(v) =_{\text{shape}} \mathbf{inj1}(v')}$$

$$\frac{\forall i. v_i =_{\text{shape}} v'_i}{\mathbf{prod}(v)_i =_{\text{shape}} \mathbf{prod}(v')_i}$$

Lemma 7.11 ($=_{\text{shape}}$ is reflexive, preserved under renaming, and implied by $=_{\alpha}$)

$$\forall v. v =_{\text{shape}} v \text{ and } \forall v, a, a'. v =_{\text{shape}} v [a'/a] \text{ and } \forall \tau. \forall v =_{\alpha} v' : \tau. v =_{\text{shape}} v'$$

The \bowtie operator helps us relate exported binders in corresponding positions in different values. It is often applied to the respective subterms of two values whose α -equivalence is in question.

Lemma 7.12 (\bowtie cannot fail, and outputs arbitrary σ s)

If $v =_{\text{shape}} v'$ and $v : \tau$ and A is finite, then $\exists \sigma, \sigma'. v \bowtie v' : \tau \rightarrow \sigma \bowtie \sigma'$ and $\mathbf{inj}(\sigma)$ and $\mathbf{inj}(\sigma')$ and $\mathbf{rng}(\sigma) \# A$ and $\mathbf{rng}(\sigma') \# A$

Lemma 7.13 (\bowtie is about binders)

If $v \bowtie v' : \tau \rightarrow \sigma \bowtie \sigma'$, then $\mathbf{dom}(\sigma) = \mathbf{fb}(\tau, v)$ and $\mathbf{dom}(\sigma') = \mathbf{fb}(\tau, v')$

Next, we show that α -equivalence is an equivalence. Showing transitivity is most complex because it depends on properties of substitutions, but even it is relatively straightforward.

Lemma 7.14 (Reflexivity of α -equivalence)

$$v =_{\mathbf{R}} v : \tau \text{ and } v =_{\mathbf{B}} v : \tau \text{ and } v =_{\alpha} v : \tau$$

Lemma 7.15 (Symmetry of α -equivalence)

$$v =_{\alpha} v' : \tau \text{ implies } v' =_{\alpha} v : \tau.$$

Lemma 7.16 (Transitivity of α -equivalence)

$$v =_{\alpha} v' : \tau \text{ and } v' =_{\alpha} v'' : \tau \text{ implies } v =_{\alpha} v'' : \tau.$$

We need an operation on substitutions that can cancel out single-atom renamings on values (see Lemma 7.18). We call it “ \otimes ”:

$$\text{If } a' \notin \mathbf{dom}(\sigma), \sigma \otimes [a'/a] \triangleq \left\{ \langle a_{\mathbf{d-old}} [a'/a], a_{\mathbf{r-old}} \rangle \mid \langle a_{\mathbf{d-old}}, a_{\mathbf{r-old}} \rangle \in \sigma \right\}$$

Lemma 7.17 (\otimes can be a no-op)

$$\text{If } a \notin \mathbf{dom}(\sigma), \text{ then } \sigma \otimes [a'/a] = \sigma.$$

Lemma 7.18 (\otimes can undo substitutions)

If $a \in \text{dom}(\sigma)$ and $a' \# v$, then $\sigma \otimes [a'/a](v [a'/a]) = \sigma(v)$.

Lemma 7.19 (\otimes adjusts domains)

$a_d [a'/a] \in \text{dom}(\sigma \otimes [a'/a])$ iff $a_d \in \text{dom}(\sigma)$

Lemma 7.20 (\otimes can distribute over \triangleright)

If $a' \notin \text{dom}(\sigma) \cup \text{dom}(\sigma')$, then $(\sigma \triangleright \sigma') \otimes [a'/a] = (\sigma \otimes [a'/a]) \triangleright (\sigma' \otimes [a'/a])$

Lemma 7.21 (\otimes commutes with β)

$\llbracket \beta \rrbracket (\sigma_i \otimes [a'/a_i])_i = \llbracket \beta \rrbracket (\sigma_i)_i \otimes [a'/a]$

Lemma 7.22 (\bowtie reflects substitutions)

If $v [a'/a] \bowtie v : \tau \rightarrow \sigma \bowtie \sigma'$ and $\text{rng}(\sigma) \# v [a'/a], v$, then $\sigma = \sigma' \otimes [a'/a]$

Lemma 7.23 (\bowtie reflects substitution, generalized)

If $v \bowtie v' : \tau \rightarrow \sigma \bowtie \sigma'$, then $v [a'/a] \bowtie v' : \tau \rightarrow \sigma \otimes [a'/a] \bowtie \sigma'$

Lemma 7.24 (Substitution commutes with free binders)

$\text{fb}(\tau, \sigma(v)) = \sigma(\text{fb}(\tau, v))$

Lemma 7.25 (Substitution commutes with free references)

If $\text{rng}(\sigma) \# v$ and $\text{inj}(\sigma)$, then $\text{fr}(\tau, \sigma(v)) = \sigma(\text{fr}(\tau, v))$

Lemma 7.26 (Substitution commutes with free atoms)

If $\text{rng}(\sigma) \# v$ and $\text{inj}(\sigma)$, then $\text{fa}(\tau, \sigma(v)) = \sigma(\text{fa}(\tau, v))$

Lemma 7.27 (Substitution commutes with exposable atoms)

If $\text{rng}(\sigma) \# \mathbf{prod}(v_i)_i$ and $\text{inj}(\sigma)$, then

$\text{xa}(\text{Prod}^{\hat{\beta}^{\text{ex}}}(\tau_i \downarrow v_i)_i, \sigma(\mathbf{prod}(v_i)_i)) = \sigma(\text{xa}(\text{Prod}^{\hat{\beta}^{\text{ex}}}(\tau_i \downarrow v_i)_i, \mathbf{prod}(v_i)_i))$

Now, we need to relate free atoms and α -equivalence. This proof is more involved, so we set it out in more detail:

Lemma 7.28 (Renaming references respects reference-equivalence iff they are not free)

Suppose $a' \# v$. Then, $a \notin \text{fr}(\tau, v)$ if and only if $v [a'/a] =_R v : \tau$.

Proof

By induction on the size of v (all cases are trivial except for products and atoms):

1. Suppose $v = a_v$ and $\tau = \text{RAtom}$.

$$\text{fr}(\text{RAtom}, a_v) = \{a_v\}$$

By def. of fr.

$$a \notin \text{fr}(\text{RAtom}, a_v) \text{ iff } v [a'/a] =_R a_v : \text{RAtom}$$

By $R\alpha$ -RATOM.

2. Suppose $v = a_v$ and $\tau = \text{BAtom}$.

$$\text{fr}(\text{BAtom}, a_v) = \emptyset$$

By def. of fr.

$$a \notin \text{fr}(\text{BAtom}, a_v) \text{ iff } v [a'/a] =_R a_v : \text{BAtom}$$

Trivially, by $R\alpha$ -BATOM.

3. Suppose $v = \mathbf{prod}(v_i)_i$ and $\tau = \mathbf{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$.

First, we define (motivated by R α -PROD):

$$P((\sigma_i)_i, (\sigma'_i)_i) \triangleq \forall i. \begin{array}{l} v_{\text{sub}} [a'/a] \bowtie v_{\text{sub},i} : \tau_{\text{sub},i} \rightarrow \sigma_i \bowtie \sigma'_i \\ \text{and } \text{rng}(\sigma_i), \text{rng}(\sigma'_i) \# v [a'/a], v \end{array}$$

First, we prove a sub-lemma that applies the induction hypothesis in the cases where a is not imported into a particular subterm.

Assume for the sake of argument:

$$P((\sigma_i)_i, (\sigma'_i)_i) \tag{7.1}$$

Fix an i . We know,

$$\begin{array}{ll} \sigma_i = \sigma'_i \otimes [a'/a] & \text{By Lemma 7.22.} \\ \llbracket \beta_i \rrbracket (\sigma_j)_j = \llbracket \beta_i \rrbracket (\sigma'_j)_j \otimes [a'/a] & \text{By Lemma 7.21.} \end{array} \tag{7.2}$$

And

$$\mathbf{fb}(\tau_{\text{sub},i}, v_{\text{sub},i}) = \text{dom}(\sigma'_i) \tag{7.3}$$

By Lemma 7.13.

Further, assume for the sake of argument:

$$a \notin \llbracket \beta_i \rrbracket (\mathbf{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}))_j \tag{7.3}$$

$$a \notin \llbracket \beta_i \rrbracket (\text{dom}(\sigma'_j))_j \tag{7.4}$$

By Lemma 7.13.

$$a \notin \text{dom}(\llbracket \beta_i \rrbracket (\sigma'_j)_j) \tag{7.5}$$

By Lemma 7.10.

$$\llbracket \beta_i \rrbracket (\sigma_j)_j = \llbracket \beta_i \rrbracket (\sigma'_j)_j \tag{7.6}$$

By 7.2 and Lemma 7.17.

$$\left(\llbracket \beta_i \rrbracket (\sigma_j)_j \right) (v_i) = \left(\llbracket \beta_i \rrbracket (\sigma'_j)_j \right) (v_i) \tag{7.7}$$

Trivially.

$$\begin{array}{l} \left(\llbracket \beta_i \rrbracket (\sigma_j)_j (v_i) \right) [a'/a] =_{\mathbf{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \\ \text{iff } a \notin \text{fr}(\tau_i, v_i) \end{array} \tag{7.4}$$

By IH.

$$a \# v \text{ or } a \notin \text{rng}(\sigma_i) \tag{7.5}$$

By 7.1.

$$a' \notin \text{dom}(\sigma_i) \tag{7.6}$$

Because $a' \# v$.

$$\left(\llbracket \beta_i \rrbracket (\sigma_j)_j (v_i) \right) [a'/a] = \llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) \tag{7.7}$$

By 7.5.

$$\begin{array}{l} \llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) =_{\mathbf{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \\ \text{iff } a \notin \text{fr}(\tau_i, v_i) \end{array} \tag{7.8}$$

By 7.4.

So

$$\forall i. \left(\begin{array}{l} P \left((\sigma_j)_j, (\sigma'_j)_j \right) \\ \text{and } a \notin \llbracket \beta_i \rrbracket \left(\text{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}) \right)_j \\ \Rightarrow \llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) =_{\text{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \\ \text{iff } a \notin \text{fr}(\tau_i, v_i) \end{array} \right) \quad \begin{array}{l} \text{Unassuming 7.1 and} \\ 7.3. \end{array} \quad (7.6)$$

Now, we actually show the product case, which has another two layers of case analysis:

a. Suppose that: $a \in \text{fr}(\tau, v)$

$$\exists i. a \in \text{fr}(\tau_i, v_i) \text{ and } a \notin \llbracket \beta_i \rrbracket \left(\text{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}) \right)_j \quad \text{By def. of fr.} \quad (7.7)$$

$$P \left((\sigma_j)_j, (\sigma'_j)_j \right) \quad \text{By Lemma 7.12.}$$

But

$$\exists i. \llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) \neq_{\text{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \quad \text{By 7.6 and 7.7.}$$

So

$$a \in \text{fr}(\tau, v) \Rightarrow v [a'/a] \neq_{\text{R}} v : \tau \quad \begin{array}{l} \text{By R}\alpha\text{-PROD and} \\ \text{supposition.} \end{array} \quad (7.8)$$

b. Suppose that: $a \notin \text{fr}(\tau, v)$

$$\forall i. a \notin \text{fr}(\tau_i, v_i) \text{ or } a \in \llbracket \beta_i \rrbracket \left(\text{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}) \right)_j \quad \text{By def. of fr.} \quad (7.9)$$

We can have $\forall j. \exists \sigma_j, \sigma'_j. :$

$$P \left((\sigma_j)_j, (\sigma'_j)_j \right) \quad \begin{array}{l} \text{By Lemmas 7.12 and} \\ 7.11. \end{array} \quad (7.10)$$

Fix an i .

i. Suppose that $a \notin \llbracket \beta_i \rrbracket \left(\text{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}) \right)_j :$

$$a \notin \text{fr}(\tau_i, v_i) \quad \text{By definition of fr.} \quad (7.11)$$

$$\llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) =_{\text{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \quad \text{By 7.6 and 7.10.} \quad (7.12)$$

ii. Suppose that $a \in \llbracket \beta_i \rrbracket \left(\text{fb}(\tau_{\text{sub},j}, v_{\text{sub},j}) \right)_j.$

In this case, we do not need the IH, because the substitutions generated by \bowtie already account for the renaming:

$$\llbracket \beta_i \rrbracket (\sigma_j)_j = \llbracket \beta_i \rrbracket (\sigma'_j)_j \otimes [a'/a] \quad \begin{array}{l} \text{By 7.10 and Lemmas} \\ 7.22 \text{ and } 7.21. \end{array}$$

$$\llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) = \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) \quad \text{By Lemma 7.18.}$$

$$\llbracket \beta_i \rrbracket (\sigma_j)_j (v_i [a'/a]) =_{\text{R}} \llbracket \beta_i \rrbracket (\sigma'_j)_j (v_i) : \tau_i \quad \text{By Lemma 7.14.}$$

So

$$a \notin \text{fr}(\tau, v) \Rightarrow v [a'/a] =_{\text{R}} v : \tau \quad \text{By 7.12.}$$

And so

$$a \in \text{fr}(\tau, v) \text{ iff } v [a_v/a] =_{\alpha} a_v : \tau \quad \text{By 7.8.}$$

□

Lemma 7.29 (Renaming binders respects binder equivalence iff they are not free)

Suppose $a' \# v$. Then, $a \notin \text{fb}(\tau, v)$ iff $v [a'/a] =_{\text{B}} v : \tau$.

Proof

By induction on v (all cases are trivial except for products and atoms):

1. Suppose $v = a_v$ and $\tau = \text{BAtom}$:

$$\text{fb}(\text{BAtom}, v) = \{a_v\} \quad \text{By def. of fb.}$$

$$a \in \{a_v\} \text{ iff } v [a'/a] =_{\text{B}} v : \tau \quad \text{By B}\alpha\text{-BATOM.}$$

2. Suppose $v = a_v$ and $\tau = \text{RAtom}$:

$$\text{fb}(\text{BAtom}, a_v) = \emptyset \quad \text{By def. of fb.}$$

$$a \notin \text{fb}(\text{BAtom}, a_v) \text{ iff } v [a'/a] =_{\text{B}} a_v : \text{BAtom} \quad \text{Trivially, by B}\alpha\text{-BATOM.}$$

3. Suppose $v = \mathbf{prod}(v_i)_i$ and $\tau = \text{Prod}^{\uparrow\beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$:

$$a \in \text{fb}(\tau, v) \text{ iff } a \in \llbracket \beta_{\text{ex}} \rrbracket (\text{fb}(\tau_i, v_i))_i \quad \text{By def. of fb.}$$

$$\text{iff } \exists i \hat{\in} \beta_{\text{ex}}. a \in \text{fb}(\tau_i, v_i) \quad \text{By a straightforward induction on } \beta.$$

$$\text{iff } \exists i \hat{\in} \beta_{\text{ex}}. v_i [a'/a] \neq_{\text{B}} v_i : \tau_i \quad \text{By IH.}$$

$$\text{iff } v [a'/a] \neq_{\text{B}} v : \tau \quad \text{By B}\alpha\text{-PROD.}$$

□

Lemma 7.30 (Renaming atoms respect α -equivalence iff they are not free)

Suppose $a' \# v$. Then, $a \notin \text{fa}(\tau, v)$ if and only if $v [a'/a] =_{\alpha} v : \tau$.

Proof

Follows from Lemmas 7.28 and 7.29 and αEQ .

□

Now, we prove some corollaries of those lemmas.

Lemma 7.31 (Non-free atoms can be mass-renamed)

Suppose $\text{inj}(\sigma)$ and $\text{rng}(\sigma) \# v$ and $\text{dom}(\sigma) \# \text{rng}(\sigma)$.

$$\text{dom}(\sigma) \# \text{fb}(\tau, v) \Rightarrow \sigma(v) =_{\text{B}} v : \tau$$

$$\text{dom}(\sigma) \# \text{fr}(\tau, v) \Rightarrow \sigma(v) =_{\text{R}} v : \tau$$

$$\text{dom}(\sigma) \# \text{fa}(\tau, v) \Rightarrow \sigma(v) =_{\alpha} v : \tau$$

Lemma 7.32 (Good renamings preserve binder equivalence)

Suppose $a' \# v, v'$. Then, $v =_{\text{B}} v' : \tau \Rightarrow v [a'/a] =_{\text{B}} v' [a'/a] : \tau$

Lemma 7.33 (Good renamings preserve reference equivalence)

Suppose $a' \# v, v'$. Then, $v =_{\text{R}} v' : \tau \Rightarrow v [a'/a] =_{\text{R}} v' [a'/a] : \tau$

Lemma 7.34 (Good renamings preserve α -equivalence)

Suppose $a' \# v, v'$. Then, $v =_{\alpha} v' : \tau \Rightarrow v [a'/a] =_{\alpha} v' [a'/a] : \tau$

Lemma 7.35 (Good substitutions preserve α -equivalence)

Suppose $\text{rng}(\sigma) \# v, v'$ and $\text{dom}(\sigma) \# \text{rng}(\sigma)$ and $\text{inj}(\sigma)$.

$$v =_{\text{B}} v' : \tau \Rightarrow \sigma(v) =_{\text{B}} \sigma(v') : \tau$$

$$v =_{\text{R}} v' : \tau \Rightarrow \sigma(v) =_{\text{R}} \sigma(v') : \tau$$

$$v =_{\alpha} v' : \tau \Rightarrow \sigma(v) =_{\alpha} \sigma(v') : \tau$$

Lemma 7.36 (Free atoms are the same for α -equivalent values)

$$v =_{\text{B}} v' : \tau \Rightarrow \text{fb}(\tau, v) = \text{fb}(\tau, v')$$

$$v =_{\text{R}} v' : \tau \Rightarrow \text{fr}(\tau, v) = \text{fr}(\tau, v')$$

$$v =_{\alpha} v' : \tau \Rightarrow \text{fa}(\tau, v) = \text{fa}(\tau, v')$$

The following are helper lemmas for Lemma 7.42.

Lemma 7.37 (Binder-equivalent values don't need their free binders renamed)

Suppose $v \bowtie v' : \tau \rightarrow \sigma \bowtie \sigma'$. If $v =_{\text{B}} v' : \tau$, then $\sigma = \sigma'$.

Lemma 7.38 (Identical renamings can be removed, retaining reference equality)

If $\sigma(v) =_{\text{R}} \sigma'(v') : \tau$ and $\sigma(a) = a' = \sigma'(a)$ and $a' \# v, v'$,

then $(\sigma \setminus \{a\})(v) =_{\text{R}} (\sigma' \setminus \{a\})(v') : \tau$

Lemma 7.39 (Identical renamings can be removed, retaining binder equality)

If $\sigma(v) =_{\text{B}} \sigma'(v') : \tau$ and $\sigma(a) = a' = \sigma'(a)$ and $a' \# v, v'$,

then $(\sigma \setminus \{a\})(v) =_{\text{B}} (\sigma' \setminus \{a\})(v') : \tau$

Lemma 7.40 (Disjointness of unrelated binders)

Let $v = \mathbf{prod}(v_i)_i$ and $\tau = \text{Prod}^{\hat{\beta}_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$ and $\sigma_{\text{all}} = \bigcup_i \sigma'_i$.

Suppose the following: $\forall i. \sigma'_i \subseteq \sigma_i$ and $\text{dom}(\sigma_{\text{all}}) \# \text{fr}(\tau, v)$ and $\forall i \neq j. \text{dom}(\sigma_i) \# \text{dom}(\sigma_j)$ and $\forall i. \text{fb}(\tau_i, v_i) = \text{dom}(\sigma_i)$.

Then, $\forall i. \text{dom}(\sigma_{\text{all}}) \setminus \text{dom}(\llbracket \beta_i \rrbracket(\sigma'_j)_j) \# \text{fr}(\tau_i, v_i)$.

Lemma 7.41 (\bowtie makes pairs of binder-disjoint values binder-equal)

If $\vdash_{\text{bndrs-disj}} v : \tau$ and $\vdash_{\text{bndrs-disj}} v' : \tau$, then

$$v \bowtie v' : \tau \rightarrow \sigma \bowtie \sigma' \Rightarrow \sigma(v) =_{\text{B}} \sigma'(v') : \tau.$$

The purpose of the following lemma is to show that any pair of α -equivalent, sufficiently disjoint values are suitable for unpacking into a pair of environments that contain the free names A . By this, we mean that a single pair of substitutions can make their *subterms* α -equivalent, not one pair per pair of subterms.

Lemma 7.42 (Sufficiently disjoint α -equivalent values can be opened)

Let $\tau_{\text{obj}} = \text{Prod}^{\hat{\beta}_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$ and $v = \mathbf{prod}(v_i)_i$ and $v' = \mathbf{prod}(v'_i)_i$.

Suppose $\text{fa}(\tau_{\text{obj}}, v) \vdash_{\text{suff-disj}} v : \tau_{\text{obj}}$ and $\text{fa}(\tau_{\text{obj}}, v') \vdash_{\text{suff-disj}} v' : \tau_{\text{obj}}$ and A is finite.

If $v =_{\alpha} v' : \tau_{\text{obj}}$, then $\exists \sigma, \sigma'. \text{dom}(\sigma) \subseteq \text{xa}(\tau_{\text{obj}}, v)$ and $\text{dom}(\sigma') \subseteq \text{xa}(\tau_{\text{obj}}, v')$ and $\forall i. \sigma(v_i) =_{\alpha} \sigma'(v'_i) : \tau_i$ and $\text{rng}(\sigma) \# A$ and $\text{rng}(\sigma') \# A$ and $\text{inj}(\sigma)$ and $\text{inj}(\sigma')$

Proof

$$\forall i. v_i \bowtie v'_i : \tau_i \rightarrow \sigma_{\text{orig},i} \bowtie \sigma'_{\text{orig},i} \quad \text{By R}\alpha\text{-PROD.} \quad (7.13)$$

$$\forall i, j. \text{rng}(\sigma_{\text{orig},i}), \text{rng}(\sigma'_{\text{orig},i}) \# v_j, v'_j \quad \text{Likewise.} \quad (7.14)$$

$$\forall i. \llbracket \beta_i \rrbracket (\sigma_{\text{orig},j})_j(v_i) =_{\text{R}} \llbracket \beta_i \rrbracket (\sigma'_{\text{orig},j})_j(v'_i) : \tau_i \quad \text{Likewise.} \quad (7.15)$$

To avoid collisions, choose σ_{disj} :

$$\text{dom}(\sigma_{\text{disj}}) = \bigcup_i \text{rng}(\sigma_{\text{orig},i}) \cup \text{rng}(\sigma'_{\text{orig},i}) \quad (7.16)$$

$$\text{inj}(\sigma_{\text{disj}}) \text{ and } \forall i. \text{rng}(\sigma_{\text{disj}}) \# v_i, v'_i \text{ and } \text{rng}(\sigma_{\text{disj}}) \# A$$

To keep exported binders fixed, define

$$\forall i. \sigma_i \triangleq \sigma_{\text{disj}} \circ \sigma_{\text{orig},i} \text{ and } \forall i. \sigma'_i \triangleq \sigma_{\text{disj}} \circ \sigma'_{\text{orig},i} \quad (7.17)$$

Define

$$\sigma_{\text{local},i} \triangleq \begin{cases} \emptyset & i \hat{\in} \beta_{\text{ex}} \\ \sigma_i & i \hat{\notin} \beta_{\text{ex}} \end{cases} \text{ and } \sigma'_{\text{local},i} \triangleq \begin{cases} \emptyset & i \hat{\in} \beta_{\text{ex}} \\ \sigma'_i & i \hat{\notin} \beta_{\text{ex}} \end{cases} \quad (7.18)$$

$$(7.19)$$

Now,

$$\forall i, j. \begin{array}{l} \text{rng}(\sigma_{\text{local},i}), \text{rng}(\sigma'_{\text{local},i}) \# v_j, v'_j, A \\ \text{and } \text{inj}(\sigma_{\text{local},i}) \end{array} \quad \text{By 7.14 and 7.16.} \quad (7.20)$$

$$(7.21)$$

For convenience, define

$$\sigma_{\text{imp},i} \triangleq \llbracket \beta_i \rrbracket (\sigma_{\text{local},j})_j \text{ and } \sigma'_{\text{imp},i} \triangleq \llbracket \beta_i \rrbracket (\sigma'_{\text{local},j})_j \quad (7.22)$$

$$\forall i \hat{\notin} \beta_{\text{ex}}. \sigma_{\text{imp},i}(v_i) =_{\text{R}} \sigma'_{\text{imp},i}(v'_i) : \tau_i \quad \text{By 7.15.} \quad (7.23)$$

$$\text{dom}(\sigma_{\text{local},i}) \subseteq \text{fb}(\tau_i, v_i) \text{ and } \text{dom}(\sigma'_{\text{local},i}) \subseteq \text{fb}(\tau_i, v'_i) \quad \text{By 7.13 and Lemma 7.13.} \quad (7.24)$$

But

$$v =_{\text{B}} v' : \tau_{\text{obj}} \quad \text{Because } v =_{\alpha} v' : \tau_{\text{obj}}.$$

$$\forall i \hat{\in} \beta_{\text{ex}}. v_i =_{\text{B}} v'_i : \tau_i \quad \text{By def. of } =_{\text{B}}.$$

$$\forall i \hat{\in} \beta_{\text{ex}}. \sigma_i = \sigma'_i \quad \text{By Lemma 7.37.} \quad (7.25)$$

So it doesn't matter whether a subterm is exported:

$$\forall i. \sigma_{\text{imp},i}(v_i) =_{\text{R}} \sigma'_{\text{imp},i}(v'_i) : \tau_i \quad \text{By 7.23, 7.18, and 7.20} \\ \text{and using Lemma 7.38} \\ \text{once for each} \quad (7.26)$$

$a \in \text{dom}(\sigma_i)$, where
 $i \hat{\in} \beta_{\text{ex}}$.

Now,

$$\forall i, j \in \beta_{\text{ex}}. i \neq j \Rightarrow \text{dom}(\sigma_{\text{local},i}) \# \text{dom}(\sigma_{\text{local},j}) \quad \text{By def. of } \vdash_{\text{suff-disj}} \text{ and} \quad (7.27)$$

7.24.

And

$$\forall i \neq j. \text{dom}(\sigma_{\text{local},i}) \# \text{dom}(\sigma_{\text{local},j}) \quad \text{By 7.18.} \quad (7.28)$$

So, define a single substitution:

$$\sigma \triangleq \bigsqcup_i \sigma_{\text{local},i} \tag{7.29}$$

$$\text{dom}(\sigma) \# \text{rng}(\sigma) \tag{By 7.20.} \tag{7.30}$$

But:

$$\text{dom}(\sigma) \# \text{fb}(\tau_{\text{obj}}, v) \tag{By def. of fb.} \tag{7.31}$$

$$\text{dom}(\sigma) \subseteq \text{xa}(\tau_{\text{obj}}, v) \tag{By def. of xa.} \tag{7.32}$$

$$\text{dom}(\sigma) \# \text{fa}(\tau_{\text{obj}}, v) \tag{By def. of } \vdash_{\text{suff-disj}} \tag{7.33}$$

$$\text{dom}(\sigma) \# \text{fr}(\tau_{\text{obj}}, v) \tag{By def. of fa.}$$

$$\forall i. \text{dom}(\sigma) \setminus \text{dom}(\sigma_{\text{imp},i}) \# \text{fr}(\tau_i, v_i) \tag{By 7.28 on Lemma 7.40.} \tag{7.33}$$

And,

$$\forall i. \text{dom}(\sigma) \setminus \text{dom}(\sigma_{\text{imp},i}) \# \bigcup_i \text{rng}(\sigma_i) \tag{By 7.20 and 7.24 and Lemma 7.2.} \tag{7.34}$$

But

$$\forall i. \text{fr}(\tau_i, \sigma_{\text{imp},i}(v_i)) = \sigma_{\text{imp},i}(\text{fr}(\tau_i, v_i)) \tag{By 7.20 and Lemma 7.25.}$$

$$\subseteq (\text{fr}(\tau_i, v_i) \setminus \text{dom}(\sigma_{\text{imp},i})) \cup \text{rng}(\sigma_{\text{imp},i}) \tag{By set arithmetic.}$$

$$\subseteq \text{fr}(\tau_i, v_i) \cup \text{rng}(\sigma_{\text{imp},i}) \tag{By set arithmetic.}$$

$$\subseteq \text{fr}(\tau_i, v_i) \cup \bigcup_j \text{rng}(\sigma_j) \tag{By Lemma 7.5.}$$

$$\# \text{dom}(\sigma) \setminus \text{dom}(\sigma_{\text{imp},i}) \tag{By 7.33 and 7.34.}$$

$$\text{fr}(\tau_i, \sigma_{\text{imp},i}(v_i)) \# \text{dom}(\sigma \setminus \sigma_{\text{imp},i}) \tag{By set arithmetic.} \tag{7.35}$$

$$\sigma_{\text{imp},i} \subseteq \sigma \tag{By Lemma 7.4.}$$

$$\forall i. \sigma_{\text{imp},i}(v_i) =_{\text{R}} (\sigma \setminus \sigma_{\text{imp},i})(\sigma_{\text{imp},i}(v_i)) : \tau_i \tag{By 7.35 and 7.30 on Lemma 7.31.}$$

$$\forall i. \quad = \sigma(v_i) \tag{By 7.30.} \tag{7.37}$$

And,

$$\sigma' \triangleq \bigsqcup_i \sigma'_{\text{local},i} \text{ and } \forall i. \sigma'_{\text{imp},i}(v'_i) =_{\text{R}} \sigma'(v'_i) : \tau_i \tag{By a similar argument.}$$

$$\forall i. \sigma(v_i) =_{\text{R}} \sigma'(v'_i) : \tau_i \tag{By 7.37 and 7.26 and Lemmas 7.28 and 7.16.} \tag{7.39}$$

Also,

$$\forall i \hat{\in} \beta_{\text{ex}}. \text{dom}(\sigma) \# \text{fb}(\tau_i, v_i) \text{ and } \text{dom}(\sigma') \# \text{fb}(\tau_i, v'_i) \tag{By 7.31 and def. of fb and Lemma 7.5.} \tag{7.40}$$

$$\forall i \hat{\in} \beta_{\text{ex}}. \sigma(v_i) =_{\text{B}} \sigma'(v'_i) : \tau_i \tag{By 7.20 and 7.30 on Lemma 7.31.} \tag{7.41}$$

And,

$$\forall i \hat{\notin} \beta_{\text{ex}}. \sigma_{\text{orig},i}(v_i) =_{\text{B}} \sigma'_{\text{orig},i}(v'_i) : \tau_i \quad \begin{array}{l} \text{By def. of } \vdash_{\text{suff-disj}} \text{ and} \\ 7.13 \text{ on Lemma 7.41.} \end{array} \quad (7.42)$$

$$\forall i \hat{\notin} \beta_{\text{ex}}. \sigma_i(v_i) =_{\text{B}} \sigma'_i(v'_i) : \tau_i \quad \begin{array}{l} \text{By 7.16 on Lemma} \\ 7.31. \end{array}$$

$$\forall i \hat{\notin} \beta_{\text{ex}}. \sigma_{\text{local},i}(v_i) =_{\text{B}} \sigma'_{\text{local},i}(v'_i) : \tau_i \quad \begin{array}{l} \text{Because} \\ \forall i \hat{\notin} \beta_{\text{ex}}. \sigma_{\text{local},i} = \sigma_i \\ \text{and } \sigma'_{\text{local},i} = \sigma'_i. \end{array}$$

$$\forall i \hat{\notin} \beta_{\text{ex}}. \sigma(v_i) =_{\text{B}} \sigma'(v'_i) : \tau_i \quad \begin{array}{l} \text{By 7.16 on Lemma} \\ 7.31. \end{array} \quad (7.43)$$

Finally,

$$\forall i. \sigma(v_i) =_{\alpha} \sigma'(v'_i) : \tau_i \quad \text{By 7.39 and 7.41.}$$

Furthermore, we have:

$$\text{dom}(\sigma) \subseteq \text{xa}(\tau_{\text{obj}}, v) \text{ and } \text{dom}(\sigma') \subseteq \text{xa}(\tau_{\text{obj}}, v') \quad \text{By 7.32.}$$

$$\text{rng}(\sigma) \# A \text{ and } \text{rng}(\sigma') \# A \quad \text{By 7.16.}$$

$$\text{inj}(\sigma) \text{ and } \text{inj}(\sigma') \quad \text{By 7.13 and 7.16.}$$

□

Lemma 7.43 (New, disjoint, binders can be generated)

Given $\mathbf{prod}(v_i)_i : \text{Prod}(\tau_i \downarrow \beta_i)_i$, and finite $A: \forall i. \exists \sigma_i$, such that $v_i \bowtie v_i : \tau_i \rightarrow \sigma_i \bowtie \sigma_i$ and $\text{inj}(\sigma_i)$ and $\text{rng}(\sigma_i) \# A$ and $\forall j \neq i. \text{rng}(\sigma_i) \# \text{rng}(\sigma_j)$ and $\forall j. \text{rng}(\sigma_i) \# \text{dom}(\sigma_j)$.

Lemma 7.44 (Products with binder-equivalent components are binder-equivalent)

If $\forall i. v_i =_{\text{B}} v'_i : \tau_i$, then $\mathbf{prod}(v_i)_i =_{\text{B}} \mathbf{prod}(v'_i)_i : \text{Prod}^{\uparrow \beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$.

Lemma 7.45 (Products with α -equivalent components are α -equivalent)

If $\forall i. v_i =_{\alpha} v'_i : \tau_i$, then $\mathbf{prod}(v_i)_i =_{\alpha} \mathbf{prod}(v'_i)_i : \text{Prod}^{\uparrow \beta_{\text{ex}}}(\tau_i \downarrow \beta_i)_i$.

Lemma 7.46 (Substitutions don't affect FAULT's uniqueness)

$\sigma(w) = \text{FAULT}$ iff $w = \text{FAULT}$

Lemma 7.47 (Good substitutions are undoable)

If $\text{inj}(\sigma)$ and $\text{rng}(\sigma) \# v$, then $v = \text{inv}(\sigma)(\sigma(v))$ and $\text{inj}(\text{inv}(\sigma))$ and $\text{rng}(\text{inv}(\sigma)) \# \sigma(v)$.

Lemma 7.48 (Good substitutions reflect $\vdash_{\text{suff-disj}}$)

If $\text{inj}(\sigma)$ and $\text{rng}(\sigma) \# v$ and $\sigma(A) \vdash_{\text{suff-disj}} \sigma(v) : \tau$, then $A \vdash_{\text{suff-disj}} v : \tau$

Lemma 7.49 (Good substitutions preserve $\vdash_{\text{suff-disj}}$)

If $\text{inj}(\sigma)$ and $\text{rng}(\sigma) \# A, v$ and $A \vdash_{\text{suff-disj}} v : \tau$, then $\sigma(A) \vdash_{\text{suff-disj}} \sigma(v) : \tau$.

Lemma 7.50 (Range update can be expressed in terms of composition)

If $\text{rng}(\sigma) \# v$, then for all σ_{avoid} , we have $\{\langle a, \sigma_{\text{avoid}}(a') \rangle \mid \langle a, a' \rangle \in \sigma\}(v) = \sigma_{\text{avoid}}(\sigma(v))$

Lemma 7.51 (Simultaneous range updates respect α -equivalence)

If $\text{rng}(\sigma) \# v$ and $\text{rng}(\sigma') \# v'$ and $\text{inj}(\sigma_{\text{avoid}})$ and $\text{rng}(\sigma_{\text{avoid}}) \# \sigma(v), \sigma'(v')$ and $\text{dom}(\sigma_{\text{avoid}}) \# \text{rng}(\sigma_{\text{avoid}})$, then

$$\begin{aligned} \left\{ \langle a, \sigma_{\text{avoid}}(a') \rangle \mid \langle a, a' \rangle \in \sigma \right\}(v) &=_{\alpha} \left\{ \langle a, \sigma_{\text{avoid}}(a') \rangle \mid \langle a, a' \rangle \in \sigma' \right\}(v') \\ &: \tau \text{ iff } \sigma(v) =_{\alpha} \sigma'(v') : \tau \end{aligned}$$

7.3 Execution soundness itself

Theorem 7.1 (Determinism up to α -equivalence, termination-insensitive version)

If $\tau = \text{typeof}(\Gamma, e)$
 and $\rho =_{\alpha} \rho' : \Gamma$
 and $\Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} w$
 and $\Gamma \vdash_{\text{exe}} \langle e, \rho' \rangle \xRightarrow{k'} w'$
 then $w =_{\alpha} w' : \tau$

Proof

By induction on k , and case analysis on e in

If $\Gamma \vdash_{\text{type}} e : \tau$
 and $\sigma(\rho) =_{\alpha} \sigma'(\rho') : \Gamma$
 and $\Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} w$
 and $\Gamma \vdash_{\text{exe}} \langle e, \rho' \rangle \xRightarrow{k'} w'$
 and $\text{inj}(\sigma)$ and $\text{rng}(\sigma) \# \rho, w$
 and $\text{inj}(\sigma')$ and $\text{rng}(\sigma') \# \rho', w'$
 then $\sigma(w) =_{\alpha} \sigma'(w') : \tau$

- Case E-FRESH- \star :

First,

$$\begin{array}{l} \Gamma \vdash_{\text{exe}} \langle (\mathbf{fresh} \ x \ \mathbf{in} \ e), \rho \rangle \xRightarrow{k+1} w_{\text{out}} \\ \Gamma \vdash_{\text{exe}} \langle (\mathbf{fresh} \ x \ \mathbf{in} \ e), \rho' \rangle \xRightarrow{k'+1} w'_{\text{out}} \end{array} \quad \text{By hypothesis.} \quad (7.44)$$

Define

$$\Gamma_{\text{IH}} \triangleq \Gamma, x:\text{BAtom}$$

We know,

$$\exists a. \Gamma_{\text{IH}} \vdash_{\text{exe}} \langle e, \rho [x \rightarrow a] \rangle \xRightarrow{k} w \quad \text{By inv. of E-FRESH-}\star. \quad (7.45)$$

$$\exists a'. \Gamma_{\text{IH}} \vdash_{\text{exe}} \langle e, \rho' [x \rightarrow a'] \rangle \xRightarrow{k'} w' \quad \text{By inv. of E-FRESH-}\star. \quad (7.46)$$

$\text{inj}(\sigma)$ and $\text{inj}(\sigma')$

$\text{rng}(\sigma) \# \rho, w, a$

$\text{rng}(\sigma') \# \rho', w', a'$

$\sigma(\rho) =_{\alpha} \sigma'(\rho') : \Gamma$

By hypothesis and inv. of E-FRESH- \star . (7.47)

By hypothesis. (7.48)

Choose

$$a_{\text{fr}} \# \sigma, \sigma', w, w', \rho, \rho', a, a' \quad (7.49)$$

Define

$$\sigma_{\text{IH}} \triangleq \{ \langle \sigma(a), a_{\text{fr}} \rangle \} \circ \sigma \quad (7.50)$$

$$\sigma'_{\text{IH}} \triangleq \{ \langle \sigma'(a'), a_{\text{fr}} \rangle \} \circ \sigma'$$

$$\sigma_{\text{IH}}(a) = \sigma'_{\text{IH}}(a') \quad \text{By straightforward substitution.} \quad (7.51)$$

And,

$$\begin{array}{ll}
 a \notin \text{fa}_{\text{env}}(\Gamma, \rho) \text{ and } a' \notin \text{fa}_{\text{env}}(\Gamma, \rho') & \text{By inv. of E-FRESH-}\star. \\
 \sigma(a) \notin \text{fa}_{\text{env}}(\Gamma, \sigma(\rho)) & \text{By 7.47 and Lemma} \\
 \sigma'(a') \notin \text{fa}_{\text{env}}(\Gamma, \sigma(\rho')) & \text{7.26.} \\
 \sigma_{\text{IH}}(\rho) =_{\alpha} \sigma'_{\text{IH}}(\rho') : \Gamma & \text{By 7.49 on Lemma} \\
 & \text{7.30.} \\
 \sigma_{\text{IH}}(\rho [x \rightarrow a]) =_{\alpha} \sigma'_{\text{IH}}(\rho' [x \rightarrow a']) : \Gamma_{\text{IH}} & \text{By 7.51.} \quad (7.52)
 \end{array}$$

And,

$$\text{inj}(\sigma_{\text{IH}}) \text{ and } \text{inj}(\sigma'_{\text{IH}}) \quad \text{By 7.47 and 7.49.} \quad (7.53)$$

$$\begin{array}{ll}
 \text{rng}(\sigma_{\text{IH}}) \# \rho [x \rightarrow a], w & \\
 \text{rng}(\sigma'_{\text{IH}}) \# \rho' [x \rightarrow a'], w' & \text{By 7.47 and 7.49.} \quad (7.54)
 \end{array}$$

So

$$\sigma_{\text{IH}}(w) =_{\alpha} \sigma'_{\text{IH}}(w') : \tau \quad \text{By 7.52 and 7.53 on} \quad (7.55) \\ \text{IH.}$$

Now, we need to show that the results are α -equivalent under the original substitution.

1. Suppose that $w = \text{FAULT}$ or $w' = \text{FAULT}$:

$$\begin{array}{ll}
 w = w' = \text{FAULT} & \text{By 7.55 on Lemma} \quad (7.56) \\
 & \text{7.46.} \\
 w_{\text{out}} = \text{FAULT and } w'_{\text{out}} = \text{FAULT} & \text{By 7.45 and 7.46 and} \\
 & \text{E-FRESH-}\star.
 \end{array}$$

2. Suppose that $w \neq \text{FAULT}$ and $w' \neq \text{FAULT}$:

$$\begin{array}{ll}
 \text{fa}(\tau, \sigma_{\text{IH}}(w)) = \text{fa}(\tau, \sigma'_{\text{IH}}(w')) & \text{By 7.55 on Lemma} \quad (7.57) \\
 & \text{7.36.} \\
 \sigma_{\text{IH}}(\text{fa}(\tau, w)) = \sigma'_{\text{IH}}(\text{fa}(\tau, w')) & \text{By 7.53 and 7.54 on} \\
 & \text{Lemma 7.26.} \\
 \sigma_{\text{IH}}(a) \notin \text{fa}(\tau, \sigma_{\text{IH}}(w)) & \\
 \text{iff } \sigma'_{\text{IH}}(a') \notin \text{fa}(\tau, \sigma'_{\text{IH}}(w')) & \text{By 7.51 and 7.53.} \\
 a \notin \text{fa}(\tau, w) \text{ iff } a' \notin \text{fa}(\tau, w') & \text{By 7.53.} \quad (7.58)
 \end{array}$$

a. So, suppose that both names are not free:

$$a \notin \text{fa}(\tau, w) \text{ and } a' \notin \text{fa}(\tau, w') \quad (7.59)$$

$$w_{\text{out}} = w \text{ and } w'_{\text{out}} = w' \quad \text{By 7.45 and 7.46 and} \quad (7.60) \\ \text{E-FRESH-}\star.$$

And also,

$$\begin{array}{ll}
 \sigma_{\#}(a) \notin \text{fa}(\tau, \sigma_{\#}(w)) & \\
 \text{and } \sigma'_{\#}(a') \notin \text{fa}(\tau, \sigma'_{\#}(w')) & \text{By 7.59 and 7.47 and} \quad (7.61) \\
 & \text{Lemma 7.26.}
 \end{array}$$

$$\sigma_{\#}(w) =_{\alpha} \sigma'_{\#}(w') : \tau \quad \text{By 7.55 and 7.49 on} \\ \text{Lemma 7.30.}$$

$$\begin{array}{ll}
\sigma(w) =_{\alpha} \sigma'(w') : \tau & \text{By 7.49 on Lemma 7.51.} \\
\sigma(w_{\text{out}}) =_{\alpha} \sigma'(w'_{\text{out}}) : \tau & \text{By 7.60.} \quad (7.62)
\end{array}$$

b. Alternatively, suppose $a \in \text{fa}(\tau, w)$ and $a' \in \text{fa}(\tau, w')$:

$$w_{\text{out}} = \text{FAULT and } w'_{\text{out}} = \text{FAULT} \quad \text{By 7.45 and 7.46 and E-FRESH-}^{\star}$$

- *Case E-OPEN- \star* : Let $v_{\text{obj}} \triangleq \mathbf{prod}(v_i)_i$ and $v'_{\text{obj}} \triangleq \mathbf{prod}(v'_i)_i$ and $e_{\text{all}} \triangleq (\mathbf{open} \ x_{\text{obj}} \ ((x_i)_i) \ e)$.

First,

$$\begin{array}{ll}
\Gamma \vdash_{\text{exe}} \langle e_{\text{all}}, \rho \rangle \xRightarrow{k+1} w_{\text{out}} & \text{By hypothesis.} \\
\text{and } \Gamma \vdash_{\text{exe}} \langle e_{\text{all}}, \rho' \rangle \xRightarrow{k'+1} w'_{\text{out}} & (7.63)
\end{array}$$

Define

$$\Gamma_{\text{IH}} \triangleq \Gamma, (x_i : \tau_i)_i$$

Now,

$$\forall i. \exists v_i. \Gamma_{\text{IH}} \vdash_{\text{exe}} \langle e, \rho [x_i \rightarrow v_i]_i \rangle \xRightarrow{k} w \quad \text{By inv. of E-OPEN-}\star \text{ on 7.63.} \quad (7.64)$$

$$\text{fa}_{\text{env}}(\Gamma, \rho) \vdash_{\text{suff-disj}} v_{\text{obj}} : \tau_{\text{obj}} \quad \text{Likewise.} \quad (7.65)$$

$$\forall i. \exists v_i. \Gamma_{\text{IH}} \vdash_{\text{exe}} \langle e, \rho' [x_i \rightarrow v'_i]_i \rangle \xRightarrow{k'} w' \quad \text{Likewise.} \quad (7.66)$$

$$\text{fa}_{\text{env}}(\Gamma, \rho') \vdash_{\text{suff-disj}} v'_{\text{obj}} : \tau_{\text{obj}} \quad \text{Likewise.} \quad (7.67)$$

$$v_{\text{obj}} \triangleq \mathbf{prod}(v_i)_i =_{\alpha} \rho(x_{\text{obj}}) : \tau_{\text{obj}} \quad \text{Likewise.} \quad (7.68)$$

$$v'_{\text{obj}} \triangleq \mathbf{prod}(v'_i)_i =_{\alpha} \rho'(x_{\text{obj}}) : \tau_{\text{obj}} \quad \text{Likewise.} \quad (7.68)$$

$$\sigma(\rho) =_{\alpha} \sigma'(\rho') : \Gamma \quad \text{By hypothesis.} \quad (7.69)$$

And,

$$\sigma(v_{\text{obj}}) =_{\alpha} \sigma(\rho(x_{\text{obj}})) : \tau_{\text{obj}} \quad \text{By 7.68 and Lemma 7.35.} \quad (7.70)$$

$$=_{\alpha} \sigma'(\rho'(x_{\text{obj}})) : \tau_{\text{obj}} \quad \text{By 7.69.}$$

$$=_{\alpha} \sigma'(v'_{\text{obj}}) : \tau_{\text{obj}} \quad \text{By 7.68 and Lemma 7.35.} \quad (7.71)$$

$$\sigma(v_{\text{obj}}) =_{\alpha} \sigma'(v'_{\text{obj}}) : \tau_{\text{obj}} \quad \text{By 7.70–7.71 and Lemma 7.16.} \quad (7.72)$$

And,

$$\begin{array}{ll}
\sigma(\text{fa}_{\text{env}}(\Gamma, \rho)) \vdash_{\text{suff-disj}} \sigma(v_{\text{obj}}) : \tau_{\text{obj}} & \text{By 7.65 and 7.67 and Lemma 7.49.} \\
\sigma'(\text{fa}_{\text{env}}(\Gamma, \rho')) \vdash_{\text{suff-disj}} \sigma'(v'_{\text{obj}}) : \tau_{\text{obj}} & (7.73)
\end{array}$$

$$\sigma(\text{fa}(\tau_{\text{obj}}, v_{\text{obj}})) \vdash_{\text{suff-disj}} \sigma(v_{\text{obj}}) : \tau_{\text{obj}} \quad \text{By 7.70 and 7.71 and Lemma 7.36.}$$

$$\sigma'(\text{fa}(\tau_{\text{obj}}, v_{\text{obj}})) \vdash_{\text{suff-disj}} \sigma'(v'_{\text{obj}}) : \tau_{\text{obj}} \quad \text{By 7.70 and 7.71 and Lemma 7.36.}$$

$\exists \sigma_{\text{op}}, \sigma'_{\text{op}}.$

$$\begin{aligned} \text{dom}(\sigma_{\text{op}}) &\subseteq \text{xa}(\tau_{\text{obj}}, \sigma(v_{\text{obj}})) && \text{By 7.65 on Lemma} \\ \text{dom}(\sigma'_{\text{op}}) &\subseteq \text{xa}(\tau_{\text{obj}}, \sigma'(v'_{\text{obj}})) && 7.42. \end{aligned} \quad (7.74)$$

Such that

$$\forall i. \sigma_{\text{op}}(\sigma(v_i)) =_{\alpha} \sigma'_{\text{op}}(\sigma'(v'_i)) : \tau \quad \text{Likewise.} \quad (7.75)$$

$$\text{rng}(\sigma_{\text{op}}), \text{rng}(\sigma'_{\text{op}}) \# \sigma, \sigma', \rho, \rho', w, w', v_{\text{obj}}, v'_{\text{obj}} \quad \text{Likewise.} \quad (7.76)$$

$$\text{inj}(\sigma_{\text{op}}) \text{ and } \text{inj}(\sigma'_{\text{op}}) \quad \text{Likewise.} \quad (7.77)$$

And,

$$\begin{aligned} \text{dom}(\sigma_{\text{op}}) \# \text{rng}(\sigma_{\text{op}}) &&& \text{By 7.74 and 7.76 and} \\ \text{dom}(\sigma'_{\text{op}}) \# \text{rng}(\sigma'_{\text{op}}) &&& \text{Lemma 7.2.} \end{aligned} \quad (7.78)$$

Define

$$\sigma_{\text{IH}} \triangleq \sigma_{\text{op}} \circ \sigma \text{ and } \sigma'_{\text{IH}} \triangleq \sigma'_{\text{op}} \circ \sigma'$$

Now,

$$\begin{aligned} \text{xa}(\tau_{\text{obj}}, \sigma(v_{\text{obj}})) \# \text{fa}_{\text{env}}(\Gamma, \sigma(\rho)) &&& \text{By 7.73 and def. of} \\ \text{xa}(\tau_{\text{obj}}, \sigma'(v'_{\text{obj}})) \# \text{fa}_{\text{env}}(\Gamma, \sigma'(\rho')) &&& \vdash_{\text{suff-disj}} \text{ and Lemma} \end{aligned} \quad (7.79)$$

$$\sigma_{\text{IH}}(\rho) =_{\alpha} \sigma'_{\text{IH}}(\rho') : \Gamma \quad \text{By 7.69, 7.74, and 7.78 on Lemma 7.31.}$$

$$\sigma_{\text{IH}}(\rho[x_i \rightarrow v_i]_i) =_{\alpha} \sigma'_{\text{IH}}(\rho'[x_i \rightarrow v'_i]_i) : \Gamma_{\text{IH}} \quad \text{By 7.75.} \quad (7.80)$$

And,

$$\text{inj}(\sigma_{\text{IH}}) \text{ and } \text{inj}(\sigma'_{\text{IH}}) \quad \text{Because } \text{inj}(\sigma) \text{ and } \text{inj}(\sigma') \text{ and by 7.76 and 7.77.} \quad (7.81)$$

And,

$$\begin{aligned} \text{rng}(\sigma_{\text{IH}}) \# \rho[x_i \rightarrow v_i]_i &&& \text{Because } \text{rng}(\sigma) \# \rho \\ \text{rng}(\sigma'_{\text{IH}}) \# \rho'[x_i \rightarrow v'_i]_i &&& \text{and } \text{rng}(\sigma') \# \rho' \text{ and} \\ &&& \text{by 7.76.} \end{aligned}$$

So:

$$\sigma_{\text{IH}}(w) =_{\alpha} \sigma'_{\text{IH}}(w') : \tau \quad \text{By 7.80 and 7.81 on IH.} \quad (7.82)$$

Case 1:

$$w = \text{FAULT} \text{ or } w' = \text{FAULT}$$

$$w = w' = \text{FAULT} \quad \text{By 7.82 on Lemma 7.46.} \quad (7.83)$$

$$w_{\text{out}} = w'_{\text{out}} = \text{FAULT}$$

By 7.64 and 7.66 and E-OPEN- \star .

Case 2:

$$w \neq \text{FAULT} \text{ and } w' \neq \text{FAULT} \quad (7.84)$$

We know,

$$\sigma_{IH}(v_{obj}) =_{\alpha} \sigma_{IH}(v'_{obj}) : \tau_{obj}$$

$$xa(\tau_{obj}, \sigma_{IH}(v_{obj})) = xa(\tau_{obj}, \sigma'_{IH}(v'_{obj}))$$

$$\sigma_{IH}(xa(\tau_{obj}, v_{obj})) = \sigma'_{IH}(xa(\tau_{obj}, v'_{obj}))$$

And,

$$fa(\tau, \sigma_{IH}(w)) = fa(\tau, \sigma'_{IH}(w'))$$

$$\sigma_{IH}(fa(\tau, w)) = \sigma'_{IH}(fa(\tau, w'))$$

$$\begin{aligned} \sigma_{IH}(xa(\tau_{obj}, v_{obj})) \# \sigma_{IH}(fa(\tau, w)) \\ \text{iff } \sigma'_{IH}(xa(\tau_{obj}, v'_{obj})) \# \sigma'_{IH}(fa(\tau, w')) \end{aligned}$$

$$\begin{aligned} xa(\tau_{obj}, v_{obj}) \# fa(\tau, w) \\ \text{iff } xa(\tau_{obj}, v'_{obj}) \# fa(\tau, w') \end{aligned}$$

By 7.75 on Lemma 7.45.
By def. of xa and 7.75 and Lemma 7.36. (7.85)
By 7.76 and 7.81 and Lemma 7.27. (7.86)

By 7.82 on Lemma 7.36.
By 7.81 on Lemma 7.26.
By 7.86.

By Lemma 7.1.

So, case 2.1:

$$\begin{aligned} xa(\tau_{obj}, v_{obj}) \# fa(\tau, w) \\ \text{and } xa(\tau_{obj}, v'_{obj}) \# fa(\tau, w') \end{aligned} \tag{7.87}$$

$$w_{out} = w \text{ and } w'_{out} = w' \tag{7.88}$$

And also,

$$\begin{aligned} \sigma(xa(\tau_{obj}, v_{obj})) \# \sigma(fa(\tau, w)) \\ \sigma'(xa(\tau_{obj}, v'_{obj})) \# \sigma'(fa(\tau, w')) \end{aligned}$$

$$\begin{aligned} \text{dom}(\sigma_{op}) \# \sigma(fa(\tau, w)) \\ \text{dom}(\sigma'_{op}) \# \sigma'(fa(\tau, w')) \end{aligned}$$

$$\begin{aligned} \text{dom}(\sigma_{op}) \# fa(\tau, \sigma(w)) \\ \text{dom}(\sigma'_{op}) \# fa(\tau, \sigma'(w')) \end{aligned}$$

$$\begin{aligned} \sigma_{IH}(w) =_{\alpha} \sigma(w) : \tau \\ \sigma'_{IH}(w') =_{\alpha} \sigma'(w') : \tau \end{aligned}$$

$$\sigma(w) =_{\alpha} \sigma'(w') : \tau$$

$$\sigma(w_{out}) =_{\alpha} \sigma'(w'_{out}) : \tau$$

Or, case 2.2:

$$w_{out} = \text{FAULT and } w'_{out} = \text{FAULT}$$

By 7.64 and 7.66 and E-OPEN-OK. (7.88)

By 7.87 and because inj(σ).

By 7.74 and Lemma 7.27.

By Lemma 7.26.

By 7.77 and 7.78 on Lemma 7.31.

By 7.82 and Lemma 7.16.

By 7.88.

By 7.64 and 7.66 and E-OPEN-OK.

• *Case E-IF-★:*

First,

$$\begin{aligned} \rho(x_0) = a \text{ and } \rho'(x_0) = a' \\ \rho(x_0) = b \text{ and } \rho'(x_0) = b' \end{aligned}$$

By inv. of E-IF-★.

$$\sigma(a) =_{\alpha} \sigma'(a') : \tau' \text{ and } \sigma(b) =_{\alpha} \sigma'(b') : \tau'$$

Because
 $\sigma(\rho) =_{\alpha} \sigma'(\rho') : \Gamma.$

$$\sigma(a) = \sigma'(a') \text{ and } \sigma(b) = \sigma'(b')$$

By R α -RATOM or
B α -BATOM.

$$\sigma(a) = \sigma(b) \text{ iff } \sigma'(a') = \sigma'(b')$$

Trivially.

$$a = b \text{ iff } a' = b'$$

Because inj(σ).

Without loss of generality, assume

$$a = b \tag{7.89}$$

Now,

$$\Gamma \vdash_{\text{exe}} \langle e_0, \rho \rangle \xRightarrow{k} w$$

By inv. of E-IF-★. (7.90)

$$\Gamma \vdash_{\text{exe}} \langle e_0, \rho' \rangle \xRightarrow{k'} w'$$

Likewise.

$$w =_{\alpha} w' : \tau$$

By 7.89 and 7.90 and
because
 $\sigma(\rho) =_{\alpha} \sigma'(\rho') : \Gamma$ on
IH.

The cases for E-CALL, E-LET-★, E-CASE-★, and E-QLIT are available in the supplementary material. □

Theorem 7.2 (α -equivalent environments have equivalent termination behavior)

$$\begin{aligned} \text{If } & \tau = \text{typeof}(\Gamma, e) \\ & \text{and } \rho =_{\alpha} \rho' : \Gamma \\ & \text{and } \Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} w \\ \text{then } & \exists w'. \Gamma \vdash_{\text{exe}} \langle e, \rho' \rangle \xRightarrow{k} w' \text{ and } w =_{\alpha} w' : \tau \end{aligned}$$

Proof

By induction on k , and case analysis on e :

• *Case E-FRESH-OK: (fresh x in e)*

It suffices to show lines 7.95 and 7.94.

First,

$$\Gamma, x:\text{BAtom} \vdash_{\text{exe}} \langle e, \rho [x \rightarrow a] \rangle \xRightarrow{k} w \tag{7.91}$$

By inv. of
E-FRESH-OK.

$$a \notin \text{fa}_{\text{env}}(\Gamma, \rho) \text{ and } w = \text{FAULT} \vee a \notin \text{fa}(\tau, w) \tag{7.92}$$

Likewise.

But

$$\rho [x \rightarrow a] =_{\alpha} \rho' [x \rightarrow a] : \Gamma, x:\text{BAtom}$$

Because $\rho =_{\alpha} \rho' : \Gamma.$

$$\exists w'. \Gamma, x:\text{BAtom} \vdash_{\text{exe}} \langle e, \rho' [x \rightarrow a] \rangle \xRightarrow{k} w' \tag{7.93}$$

By 7.91 and IH.

$$w =_{\alpha} w' : \tau \quad \text{Likewise.} \quad (7.94)$$

But

$$a \notin \text{fa}_{\text{env}}(\Gamma, \rho') \text{ and } w' = \text{FAULT} \vee a \notin \text{fa}(\tau, w') \quad \text{By 7.92 on Lemma 7.36.}$$

$$\Gamma \vdash_{\text{exe}} \langle (\text{fresh } x \text{ in } e), \rho' \rangle \xRightarrow{k+1} w' \quad \text{By E-FRESH-OK.} \quad (7.95)$$

- *Case E-FRESH-FAIL*: This case proceeds almost identically to E-FRESH-OK.
- *Case E-OPEN-OK*: (**open** x $((x_i)_i)$ e)

It suffices to show lines 7.102 and 7.100.

First,

$$\Gamma, (x_i : \tau_i)_i \vdash_{\text{exe}} \langle e, \rho [x_i \rightarrow w_i]_i \rangle \xRightarrow{k} w \quad \text{By inv. of E-OPEN-OK.} \quad (7.96)$$

$$\text{fa}_{\text{env}}(\Gamma, \rho) \vdash_{\text{suff-disj}} \mathbf{prod}(w_i)_i : \tau_{\text{obj}} \quad \text{Likewise.} \quad (7.97)$$

$$w = \text{FAULT} \vee \text{xa}(\tau_{\text{obj}}, \mathbf{prod}(w_i)_i) \# \text{fa}(\tau, w) \quad \text{Likewise.} \quad (7.98)$$

But

$$\rho [x_i \rightarrow w_i]_i =_{\alpha} \rho' [x_i \rightarrow w_i]_i : \Gamma, (x_i : \tau_i)_i \quad \text{Because } \rho =_{\alpha} \rho' : \Gamma \text{ and by Lemma 7.14.}$$

$$\exists w'. \Gamma', (x_i : \tau_i)_i \vdash_{\text{exe}} \langle e, \rho' [x_i \rightarrow w_i]_i \rangle \xRightarrow{k} w' \quad \text{By 7.96 and IH.} \quad (7.99)$$

$$w =_{\alpha} w' : \tau \quad \text{Likewise.} \quad (7.100)$$

But

$$\begin{aligned} \text{fa}_{\text{env}}(\Gamma, \rho) &= \text{fa}_{\text{env}}(\Gamma, \rho') && \text{Because } \rho =_{\alpha} \rho' : \Gamma \text{ and by Lemma 7.36.} \\ \text{fa}_{\text{env}}(\Gamma, \rho') &\vdash_{\text{suff-disj}} \mathbf{prod}(w_i)_i : \tau_{\text{obj}} && \text{By 7.97.} \end{aligned} \quad (7.101)$$

And,

$$\text{fa}(\tau, w) = \text{fa}(\tau, w') \quad \text{By 7.100 on Lemma 7.36.}$$

And,

$$w' = \text{FAULT} \vee \text{xa}(\tau_{\text{obj}}, \mathbf{prod}(w'_i)_i) \# \text{fa}(\tau, w') \quad \text{By 7.98 and 7.100.}$$

$$\Gamma \vdash_{\text{exe}} \langle (\text{open } x ((x_i)_i) e), \rho' \rangle \xRightarrow{k+1} w' \quad \text{By E-OPEN-OK on 7.101 and 7.99.} \quad (7.102)$$

- *Case E-OPEN-FAIL*: This case proceeds almost identically to E-OPEN-OK.

The cases for E-CALL, E-IF- \star , E-QLIT, E-CASE- \star , and E-QLIT are available in the supplementary material. \square

8 Checking binding safety statically

This section describes the Romeo deduction system. The purpose of this deduction system is to generate constraints (proof obligations) which, if satisfied, guarantee that escape, and therefore FAULT, will never occur (see Theorem 8.1).

The proof system's judgment is of the form $\Gamma \vdash_{\text{proof}} \{H\} e \{P\}$. Like the execution semantics, it is type-dependent, and for a similar reason: types control which atoms will be bound or free, and thus whether operations are valid or not.

We use H , P , and C to range over constraints. Typically, H , the hypothesis, contains facts (about the atoms in the environment) that are true by construction. P , the post-condition, contains predicates that describe the connection between atoms in the environment and atoms in the output. C is used for other constraints.

The obligations emitted by the deduction system must be satisfied by showing them true for all ρ compatible with Γ ; in practice, we do this with an SMT (satisfiability modulo theories) Solver. (For example, Romeo-L (Muehlboeck, 2013) uses Z3 (de Moura & Bjørner, 2008). We discuss this in more detail in Section 9.8.)

We begin by giving the syntax of constraints.

$$\begin{aligned}
 z \in \text{ConstrSetVar} &::= x \mid \cdot \\
 s \in \text{SetDesc} &::= \emptyset \\
 &\mid s \cup s \\
 &\mid s \cap s \\
 &\mid sf(z) \\
 &\mid \mathcal{F}_e(\Gamma) \\
 &\mid \mathcal{X}_z(x_i)_i \\
 sf \in \text{SetFn} &::= \mathcal{F} \mid \mathcal{F}_r \mid \mathcal{F}_b \\
 H, P, C \in \text{Constraint} &::= C \wedge C \\
 &\mid s = s \\
 &\mid s \neq s \\
 &\mid s \# s \\
 &\mid s \subseteq s \\
 &\mid z \cong e^{\text{qlit}} \\
 &\mid \mathbf{true} \\
 \Gamma_{\text{dot}} \in \text{TypeEnvWithDot} &::= \epsilon \mid \Gamma_{\text{dot}}, z:\tau
 \end{aligned}$$

Formulas are constructed from variables z , which range over program variables x , and \cdot (which refers to the output value of the current expression). Set-valued terms are constructed from the free names (\mathcal{F}), free references (\mathcal{F}_r), free binders (\mathcal{F}_b), and exposable names (\mathcal{X}) of values (here, the corresponding types can be retrieved from the environment), and the free names (\mathcal{F}_e) of environments (which is syntax sugar for a union of \mathcal{F} s), and then by the standard set constructors (\cup and \cap). Atomic formulas denote equality, inequality, etc., of sets, and \cong denotes that two values have the same free binders and references. Finally, constraints are conjunctions of atomic formulas.

The syntax for the \mathcal{X} set function is unusual; it consumes the subterms of its “argument” instead of the argument itself, and also a type indicating its binding structure. (The set functions that take a single variable as an argument do not need this, because a type environment is available.) This is because, unlike \mathcal{F} , \mathcal{F}_r , and \mathcal{F}_b , the result of \mathcal{X} depends on bound names of its argument. However, it only depends on the free names of the *subterms* of its argument. This way, two values that are \cong have identical behavior under all the set functions. Furthermore, this syntax fits the way that \mathcal{X} is used in P-OPEN.

We use quasi-literals to describe values with variable interpolation. Because the type environment Γ is present, the type annotations of quasi-literals are redundant,

$$\begin{array}{c}
\frac{\text{typeof}(\Gamma, e^{\text{qlit}}) = \tau \quad \Gamma, ::\tau \models H \wedge (\cdot \cong e^{\text{qlit}}) \Rightarrow P}{\Gamma \vdash_{\text{proof}} \{H\} e^{\text{qlit}} \{P\}} \text{P-QLIT} \\
\\
\frac{\begin{array}{l} \text{retype}(f) = \tau \quad \text{formals}(f) = (x_{\text{formal},i})_i \\ \text{argtype}(f) = (\tau_{\text{formal},i})_i \quad \Gamma \models H \Rightarrow \text{pre}(f) [x_{\text{actual},i}/x_{\text{formal},i}]_i \\ \Gamma, ::\tau \models H \wedge \mathcal{F}(\cdot) \subseteq \mathcal{F}_e((x_{\text{actual},i}:\tau_{\text{formal},i})_i) \wedge \text{post}(f) [x_{\text{actual},i}/x_{\text{formal},i}]_i \Rightarrow P \end{array}}{\Gamma \vdash_{\text{proof}} \{H\} (f (x_{\text{actual},i})_i) \{P\}} \text{P-CALL} \\
\\
\frac{x \text{ fresh for } \Gamma, H, P \quad \Gamma, x:\text{BAtom} \vdash_{\text{proof}} \{H \wedge \mathcal{F}(x) \# \mathcal{F}_e(\Gamma)\} e \{P \wedge \mathcal{F}(x) \# \mathcal{F}(\cdot)\}}{\Gamma \vdash_{\text{proof}} \{H\} (\text{fresh } x \text{ in } e) \{P\}} \text{P-FRESH} \\
\\
\frac{\forall i. x_i \text{ is fresh for } \Gamma, H, P \quad \Gamma(x_{\text{obj}}) = \tau = \text{Prod}^{\uparrow\beta_e} (\tau_i \downarrow \beta_i)_i}{\Gamma, (x_i:\tau_i)_i \vdash_{\text{proof}} \{H \wedge \mathcal{X}_\tau(x_i) \# \mathcal{F}_e(\Gamma) \wedge x_{\text{obj}} \cong (\text{prod}^{\uparrow\beta_e} (x_i \downarrow \beta_i)_i)\} e \{P \wedge \mathcal{X}_\tau(x_i) \# \mathcal{F}(\cdot)\}} \text{P-OPEN} \\
\\
\frac{\begin{array}{l} x \text{ fresh for } \Gamma, H, P, C \quad \Gamma \vdash_{\text{proof}} \{H\} e_{\text{val}} \{C\} \\ \text{typeof}(\Gamma, e_{\text{val}}) = \tau_{\text{val}} \quad \Gamma, x:\tau_{\text{val}} \vdash_{\text{proof}} \{H \wedge C[x/\cdot] \wedge \mathcal{F}(x) \subseteq \mathcal{F}_e(\Gamma)\} e_{\text{body}} \{P\} \end{array}}{\Gamma \vdash_{\text{proof}} \{H\} (\text{let } x \text{ where } C \text{ be } e_{\text{val}} \text{ in } e_{\text{body}}) \{P\}} \text{P-LET} \\
\\
\frac{\begin{array}{l} \Gamma(x) = \tau_0 + \tau_1 \quad \Gamma, x_0:\tau_0 \vdash_{\text{proof}} \{H \wedge x \cong (\text{inj}_0 x_0 \tau_1)\} e_0 \{P\} \\ \Gamma, x_1:\tau_1 \vdash_{\text{proof}} \{H \wedge x \cong (\text{inj}_1 \tau_0 x_1)\} e_1 \{P\} \end{array}}{\Gamma \vdash_{\text{proof}} \{H\} (\text{case } x (x_0 e_0) (x_1 e_1)) \{P\}} \text{P-CASE} \\
\\
\frac{\Gamma \vdash_{\text{proof}} \{H \wedge \mathcal{F}(x_0) = \mathcal{F}(x_1)\} e_0 \{P\} \quad \Gamma \vdash_{\text{proof}} \{H \wedge \mathcal{F}(x_0) \# \mathcal{F}(x_1)\} e_1 \{P\}}{\Gamma \vdash_{\text{proof}} \{H\} (\text{if } x_0 \text{ equals } x_1 e_0 e_1) \{P\}} \text{P-IFEQ} \\
\\
\frac{(x_i:\tau_i)_i \vdash_{\text{proof}} \{C_0\} e \{C_1\}}{\vdash_{\text{proof}} (\text{define-fn } (f (x_i:\tau_i)_i \text{ pre } C_0) : \tau_0 e \text{ post } C_1) \text{ ok}} \text{P-FNDEF} \\
\\
\frac{\forall i. \vdash fD_i \text{ ok} \quad \varepsilon \vdash_{\text{proof}} \{\text{true}\} e \{\text{true}\}}{\vdash_{\text{proof}} (fD_i)_i e \text{ ok}} \text{P-PROG}
\end{array}$$

Fig. 6. Verification rules for the deduction system.

but for economy of abstraction, we elected to reuse an existing concept instead of creating a new one.

In general, our rules, in Figure 6, are patterned after those in Pottier (2007b), using the type information in Γ to collect information about values. This subsumes Pottier's Δ . Most rules discharge their proof obligations by delegating them to proof obligations on subexpressions. The base cases of this recursion are P-CALL and P-QLIT, which describe proof obligations of the form $\Gamma_{\text{dot}} \models H \Rightarrow P$. P-QLIT has only one obligation, which is to ensure that the result it produces obeys whatever constraints were imposed in P , given that the environment satisfies the assumptions in H . P-CALL has two obligations; first, that the invoked function's pre-condition is true (given H), and second that the resulting value satisfies the constraints in P (given H and the post-condition of the function).

As one might expect, the key rules are P-FRESH and P-OPEN, whose definitions are closely connected to E-FRESH-OK and E-OPEN-OK.

Proving the theorems in Section 6 required our language to have two important properties: that (a) no name can escape the context that exposed it, except as a bound name, and (b) no name is exposed from two different binding relationships in the same environment (thereby purporting to show equality between two names that are not related by binding). For the purposes of dynamically respecting α -equivalence, property (a) was enforced by detecting such a situation and emitting `FAULT` instead, and property (b) was established by the constraints imposed on the names exposed in `E-FRESH-*` and `E-OPEN-*`.

Now, for the purposes of the deduction system, property (a) appears in the *post-condition* of both `P-FRESH` and `P-OPEN`, as an obligation to prove that the exposed free names are disjoint from the free names of the result value (spelled “.”), because the purpose of the deduction system is to prevent `FAULTS`. On the other hand, property (b) is a guarantee provided by the language dynamics, and therefore appears in the *hypothesis* of both rules, saying that the exposed names are guaranteed to be disjoint from the environment so far.

The rules `P-OPEN` and `P-CASE` each add additional information to their hypotheses using \cong . This information conveys the relationship between the atoms in the scrutinee (x_{obj} and x_0 , respectively) and the atoms in its component(s).

8.1 Odds and ends

In the `let` expression, the body subexpression has the same result as the expression as a whole, but the value subexpression does not. Therefore, in `P-LET`, the condition C (whose \cdot refers to the value subexpression) must be adjusted for use as a hypothesis for the body subexpression. Fortunately, the name x refers to the value subexpression in question, so a simple $[x/\cdot]$ substitution suffices. H may be used unchanged by both subexpressions because it will contain no references to \cdot .

A similar issue occurs in `P-CALL`. The pre- and post-conditions of the function (not to be confused with the expression’s logical post-condition P) are expressed relative to the formal parameters, which are meaningless out of context. Because the actual arguments to a function invocation are all required to be variable references (rather than allowing them to be whole subexpressions), the solution is again simple: A simultaneous substitution from the formals to the actuals suffices to make the pre- and post-conditions meaningful in the caller’s context.

Shadowing among Romeo program variables is incompatible with the deduction system, because obligations must be able to refer to (and distinguish) everything in Γ by name. This gives rise to the requirement that certain x ’s be fresh for Γ , H , and P ; this requirement is easily satisfied by a simple renaming pass prior to type and proof checking.

`P-CALL`’s body hypothesis contains $\mathcal{F}(\cdot) \subseteq \mathcal{F}_e((x_{\text{actual},i} : \tau_{\text{formal},i})_i)$, a term representing extra information as a consequence of Lemma 8.1, which states that the free atoms in the result of any expression are a subset of the free atoms in the environment in which it is evaluated. A similar term appears in the hypothesis for `P-LET`’s body subexpression.

Finally, in P-IFEQ, the information resulting from the comparison can be expressed in our predicate language; in the branch in which the two atoms are equal, we note that their free atom sets (known to be singletons) are equal, and in the other branch, we note that their free atoms sets are disjoint.

8.2 Soundness of the deduction system

The soundness of the deduction system is expressed in the following theorem.

Theorem 8.1 (Soundness of the deduction system)

$$\begin{array}{l} \text{If} \\ \quad \tau = \text{typeof}(\Gamma, e) \\ \quad \text{and } \Gamma \vdash_{\text{proof}} \{H\} e \{P\} \\ \quad \text{and } \Gamma \vdash_{\text{type-env}} \rho \\ \quad \text{and } \Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} w \\ \text{then} \quad w \neq \text{FAULT} \text{ and } \Gamma, \cdot; \tau; \rho [\cdot \rightarrow w] \models H \Rightarrow P \end{array}$$

Proof

By induction on k . □

The most important lemma for proving Theorem 8.1 is “No Names Made Up”, which is similar to Pottier’s No Atoms Made Up lemma (2007a).

Lemma 8.1 (No names made up)

$$\begin{array}{l} \text{If} \\ \quad \tau = \text{typeof}(\Gamma, e) \\ \quad \text{and } \Gamma \vdash_{\text{type-env}} \rho \\ \quad \text{and } \Gamma \vdash_{\text{exe}} \langle e, \rho \rangle \xRightarrow{k} v \\ \text{then} \quad \text{fa}(\tau, v) \subseteq \text{fa}_{\text{env}}(\Gamma, \rho) \end{array}$$

Proof

By induction on k . □

The proofs of Lemma 8.1 and Theorem 8.1 can be found in the supplementary material.

8.3 Example

The Romeo-L code in Figure 2 contains a number of **opens**, each of which potentially can produce a **FAULT**. However, our deduction shows that **FAULT** will never happen. A complete derivation is too large to include here, but we will informally look at two examples.

First, on line 5, we are opening up a lambda abstraction. This “exposes” the lambda’s binder (binding it to the variable bv). Fortunately, $(\text{inj}_{\text{lambda}}(\text{prod } bv, \text{convert}(e\text{-body}) \downarrow 0))$, the body of the **open**, does not have bv free: its left-hand child, bv , is an unexported binder, and its right-hand child imports it, meaning that bv can’t be free in it either, regardless of what convert returns.

Second, on line 7, we open up the whole `let*` form, exposing all of the names that `lsc` exports. We must show that those names do not escape this context as free variables. When `lsc-some` is destructured, we know from its type that `bv` and `lsc-rest`, together, export that same set of names.

The value returned from the **open** is an application, constructed on lines 13–14. We first examine its left-hand side, which is a lambda binding `bv` in `e-rest`. Therefore, by a similar argument to the one above, `bv` is no problem, and we only need to show that the names exported by `lsc-rest` are bound in `e-rest`. Fortunately, `e-rest` is a `let*` construct (line 12), defined to bind the names exported by `lsc-rest` in `e-body`, which is exactly what we needed. Therefore, the left-hand side of the function application contains no free references that could cause a `FAULT`.

Now, we look at the right-hand-side of that application, which we generate by calling `convert(val-expr)`. By Lemma 8.1, we know that `convert` produces a value whose free names are a subset of its argument. How do we know what names are free in `val-expr`? We know that, as an expression, it exports nothing, and so has no free binders. Any free references in it would have also been free in `lsc-some` (because it binds no names in the scope of its value expression), and therefore free in `let-star` itself. But `let-star` is part of the environment in which it was opened (on line 7), so, by the freshness of newly exposed names, the names we are worried about must be fresh for `val-expr`.

A similar argument can be used to verify the safety of the other **opens**. In this example, the programmer didn't need to supply any constraints to justify the function calls. In general, constraints are necessary for the same reasons as in Pure FreshML (Pottier, 2007b), and the same examples apply.

9 Related work

9.1 Statically specified binding in template macros

The work of Herman and Wand (Herman & Wand, 2008; Herman, 2010) introduced the idea of a static binding specification for a template-based macro system (like Scheme's `syntax-rules`). Herman defined a language for binding specifications, and gave an algorithm for deciding whether a pattern-and-template macro was consistent with its binding specification. In practice, however, the complex macros in a language like Scheme are often not expressible in a pattern-matching system. Romeo provides a path for extending this macro system to a procedurally based one, like Scheme's `syntax-case`.

Although our binding annotation system is very similar in power to Herman's, we have made some changes in representation. The most noticeable of these is that where Herman and Wand use addresses into binary trees of values, we use indices into wide products. However, this does not affect the representational power of the system.

9.2 Pure FreshML

The second key source for this work is Pure FreshML (Pottier, 2007b). Both Pure FreshML and Romeo are first-order, side-effect-free languages in which a runtime system ensures that introduced names do not escape their scope, and both provide a

proof system that generates proof obligations which, if true, guarantee statically that no faults will occur. One difference, important for application to macro-expanders, is that Romeo manipulates plain S-expression-like data, guided by types, whereas Pure FreshML saves type information in values. Our presentation of the language and semantics are somewhat different: for example, we have separate constructs for destructuring products (**open**) and destructuring sum types (**case**).

Pure FreshML leaves the actual language of binding specifications underdetermined. All the formal development is done in a simple system, roughly equivalent to the λ -calculus, but one of the key examples, normalization by evaluation, is done using the more expressive system of C ω ml (Pottier, 2006). Still, both of these systems are too weak to express complex binding constructs. For example, neither can express the natural syntax of the `let*` construct.

Pure FreshML's safety guarantees are proved true in a nominal logic in which values are entire α -equivalence classes. However, we know of no nominal logic powerful enough to represent the complex binding structures available in Romeo, so our proof of correctness must do the work of a nominal logic "by hand".

9.3 Ott

Ott (Sewell *et al.*, 2010) is a system for metaprogramming that accepts binding specifications with a syntax and semantics similar to ours. However, Ott's goals are significantly different. Instead of providing a complete, name-aware programming system, Ott generates code for use in a theorem-prover, including definitions of types and a capture-avoiding substitution function. Ott supports a number of theorem-provers and a number of representations for the terms in them. Additionally, it can export boilerplate code for capture-avoiding substitution in OCaml.

Ott's binding specifications are strictly more expressive than ours: Effectively, they allow for a single value to export multiple sets of names (these sets are designated by "auxiliary functions"), which can be bound separately.

In order to support theorem-provers, Ott includes a definition for α -equivalence between its "concrete abstract syntax trees" (the equivalent of our values v). This definition is based on a partial equivalence relation which relates two tree positions in a term if they are connected by binding (i.e. they would have to be renamed together). From this intra-tree relation, it is fairly straightforward to extract a notion of α -equivalence: two trees are α -equivalent if both (1) their free names match, and (2) the partial equivalence relations representing their bound names are identical. It is not clear whether this definition would lead to simpler proofs of Theorems like 7.1 and 7.2.

9.4 Hygienic macro systems in Scheme

The goals of our work have a great deal in common with the goals of hygienic macro systems, like those used in Scheme.

Despite extensive use of hygienic macro systems, what it meant for them to be "hygienic" has resisted definition until recently. The algorithm described by Dybvig

(1992), which is the basis for hygienic macro expansion in Scheme, seems “correct” in the sense that it tends to behave consistently with the intuition of its users.

More recently, a reformulation by Matthew Flatt (2016) called “sets of scopes” provides largely the same behavior with a simpler semantics and in an easier-to-use fashion.

The original attempts at specification, given by Clinger (1991) and Dybvig (1992), are phrased in terms of the bindings inserted or introduced by the macro. However, given a macro definition (say, in template style), there is no obvious way to tell what those bindings should be without reference to the expansion algorithm. So this specification is circular. Recent work by Adams (2015) presents a plausible definition that examines the support of unexpanded code, while being agnostic as to its binding structure. However, we feel that the use of static binding specifications (introduced for this purpose by Herman (2010)) provides a more straightforward rigorous basis for hygiene, at least in situations in which binding specifications are appropriate.

Secondly, the Dybvig algorithm offers no *static* guarantees. If the macro designer makes a mistake, it will only be discovered after the macro is expanded (and probably after it is used, perhaps by an innocent end-user). By contrast, our system offers a static guarantee: If the macro definition binds a name incorrectly, the error will be detected at macro-definition time by the deduction system, not as a runtime fault.

9.5 Binding in theorem-proving systems

Ever since the POPLMark Challenge (Aydemir *et al.*, 2005), there has been a large interest in coding terms with bindings in various proof assistants (Urban, 2008; Pouillard & Pottier, 2010; Pollack *et al.*, 2012). These works have differing goals than ours; they are primarily concerned with proving facts about programs, while we are aiming at a usable meta-programming system. They also generally depend on representing abstract syntax trees in a pre-existing theorem-proving framework like Coq or Agda, whereas we are concerned with the complications of concrete syntax (even in an S-expression based language).

9.6 Unbound

Unbound (Weirich *et al.*, 2011) is a Haskell library for safely manipulating syntax that operates on many of the same principles as the FreshML family that Romeo comes from. In particular, Unbound’s **unbind** function corresponds roughly to our **open** form, and **bind** $p\ e$ corresponds to a quasiliteral of the form $(\mathbf{prod}\ e_p^{\text{qlit}}, e_e^{\text{qlit}} \downarrow 0)$.

Although it has a number of internal differences (such as the locally nameless representation of syntax), the programming model is fairly similar. The primary difference noticeable by the user is Unbound’s binding specification system, which is similar to C ζ ml (Pottier, 2006), but extended to the point that it supports constructs like `let*`. Like C ζ ml, it is oriented toward abstract syntax, not concrete syntax like Romeo.

9.7 SoundX

In general, it is impossible to typecheck an expression without knowing its type environment. Historically, this has been the reason that languages with both a type system and a macro system perform typechecking after macro expansion. The unfortunate consequence of this is that the user is presented type errors in terms of code that they didn't even write.

SoundX (Lorenzen & Erdweg, 2016) is a step toward a solution to this problem, demonstrating extensions to language syntax which are less powerful than macros, but whose expansions are guaranteed to be type-safe. We believe that the major remaining steps between SoundX and true typed macros are:

- Extensions defined by arbitrary code, instead of pattern matching. This is essentially the difference between λ_m (Herman, 2010) and Romeo. Proof that syntax extensions preserve the validity of type judgments must be providable by the typesystem, which might be very complex.
- Locally scoped extensions. This is potentially complex, but it may be able to at least follow the example from hygienic locally scoped macros in Scheme.
- User-friendly binding safety. SoundX extensions are statically rejected if they might perform accidental capture, requiring manual freshening and contortions around capturing constructs, and restricting what names the extension's user may choose.⁶ We have shown that it is instead possible to automatically rename to avoid collisions, which is far more convenient to the user.

We believe that typed macros will be a major user-friendliness improvement for the interaction between types and macros.

9.8 Romeo-L

As we have described it, writing programs in Romeo is tedious. Programs must be written in an ANF-like style. For example, the arguments to function calls must be variables, not more complicated expressions.

A second problem is that we have not yet described how the truth of $\Gamma \models H \Rightarrow P$ is to be determined, and, if it fails, how the programmer is supposed to figure out how to fix it.

These problems are addressed in Muehlboeck's master's thesis (2013), which presents a more usable front-end for Romeo, called Romeo-L. Romeo-L programs are written in a more natural dialect, and are automatically translated into the core Romeo we have described here. This translation introduces **let** expressions to avoid the need to program in A-normal form, and for each **let** it infers a constraint C . Therefore, the user need only supply constraints for function definitions. In practice, this means a vastly reduced annotation burden.

⁶ Although it is not immediately obvious, SoundX does possess binding specifications. This is by virtue of its type rules. For example, typing a λ expression in the environment Γ requires typing its body in the environment $\Gamma, x : \tau$, where x is the name bound by the lambda, and τ is its type. This indicates that x is bound in the λ 's body.

Romeo-L also includes a connection to the Z3 SMT solver (de Moura & Bjørner, 2008), which is able to check statements of the form $\Gamma \models H \Rightarrow P$, completing the automated checking of the deduction system. The translation takes the constraints H and P , and (in a sense) partially evaluates them until they are expressed only in terms of the free binders and free references of program variables that are never destructured (if a variable is destructured, the sets of atoms relevant to it can be expressed in terms of the variables it is destructured into). The sets of atoms that remain are uninterpreted, except to add size constraint approximations (when it is possible to determine that the number of elements of a set is 0, 1, ≤ 1 , or ≥ 1). At that point, the constraints are directly expressible as SMT problems using combinatory array logic as the concrete theory.

Furthermore, Romeo-L can translate counterexamples provided by Z3 into sets of names, so that the user can understand them, and it can explain how they violate a constraint either written by the user, or implicit in the rules for **fresh** or **open**.

Acknowledgments

We would like to thank Fabian Muehlboeck for his work in testing and validating Romeo, and Dionna Glaze for her collaboration in hammering out the formal properties of our binding system. We would also like to thank the anonymous reviewers for their comments to improve this paper.

References

- Adams, M. D. (2015) Towards the essence of hygiene. In Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. New York, NY, USA: ACM, pp. 457–469.
- Aydemir, B. E., Bohannon, A., Fairbairn, M., Foster, J. N., Pierce, B. C., Sewell, P., Vytiniotis, D., Washburn, G., Weirich, S. & Zdancewic, S. (2005). Mechanized metatheory for the masses: The PoplMark challenge. In Proceedings of the 18th International Conference on Theorem Proving in Higher Order Logics. Berlin, Heidelberg: Springer-Verlag, pp. 50–65.
- Clinger, W. & Rees, J. (1991). Macros that work. In Proceedings of the 18th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. New York, NY, USA: ACM, pp. 155–162.
- de Moura, L. & Bjørner, N. (2008). Z3: An efficient SMT solver. In Tools and Algorithms for the Construction and Analysis of Systems. Lecture Notes in Computer Science, vol. 4963, pp. 337–340. Berlin Heidelberg: Springer.
- Dybvig, R. K., Hieb, R. & Bruggeman, C. (1992). Syntactic abstraction in scheme. *LISP Symb. Comput.* 5(4), 295–326.
- Erdweg, S., van der Storm, T. & Dai, Y. (2014). Capture-avoiding and hygienic program transformations. In *ECOOP 2014 Object-Oriented Programming*, Jones, R. (ed), Lecture Notes in Computer Science, vol. 8586, pp. 489–514. Berlin Heidelberg: Springer.
- Flatt, M. (2016). Binding as sets of scopes. *SIGPLAN Not.* 51(1), 705–717.
- Hendriks, D. & van Oostrom, V. (2003). Adbmal. In *Automated Deduction CADE-19*, Baader, F. (ed), Lecture Notes in Computer Science, vol. 2741. Berlin Heidelberg: Springer, pp. 136–150.
- Herman, D. (2010). *A Theory of Typed Hygienic Macros*. PhD Thesis, Boston, MA, USA: Northeastern University.

- Herman, D. & Wand, M. (2008). A theory of hygienic macros. In *Proceedings of the Theory and Practice of Software, 17th European Conference on Programming Languages and Systems*. Berlin, Heidelberg: Springer-Verlag, pp. 48–62.
- Kohlbecker, E., Friedman, D. P., Felleisen, M. & Duba, B. (1986). Hygienic macro expansion. In *Proceedings of the 1986 ACM Conference on Lisp and Functional Programming, LFP '86*. New York, NY, USA: ACM, pp. 151–161.
- Lorenzen, F. & Erdweg, S. (2016). Sound type-dependent syntactic language extension. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016*. New York, NY, USA: ACM.
- Muehlboeck, F. (2013). *Checking Binding Hygiene Statically*. Master's Thesis, Boston, MA, USA: Northeastern University.
- Pollack, R., Sato, M. & Ricciotti, W. (2012). A canonical locally named representation of binding. *J. Autom. Reason.* **49**(2), 185–207.
- Pottier, F. (2006). An overview of C_zml. *Electron. Notes Theor. Comput. Sci.* **148**(2), 27–52.
- Pottier, F. (2007a). *Static Name Control for FreshML*. Available at: <http://gallium.inria.fr/~fpottier/publis/fpottier-pure-freshml-long.pdf>. [4 July, 2016]
- Pottier, F. (2007b). Static name control for FreshML. In *Proceedings of the 22nd Annual IEEE Symposium on Logic in Computer Science*. Washington, DC, USA: IEEE Computer Society, pp. 356–365.
- Pouillard, N. & Pottier, F. (2010). A fresh look at programming with names and binders. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming*. New York, NY, USA: ACM, pp. 217–228.
- Sewell, P., Nardelli, F. Z., Owens, S., Peskine, G., Ridge, T., Sarkar, S. & Strnia, R. (2010). Ott: Effective tool support for the working semanticist. *J. Funct. Program.* **20**, 71–122.
- Shinwell, M. R., Pitts, A. M. & Gabbay, M. J. (2003). FreshML: Programming with binders made simple. In *Proceedings of the 8th ACM SIGPLAN International Conference on Functional Programming, ICFP '03*, pp. 263–274. New York, NY, USA: ACM.
- Urban, C. (2008). Nominal techniques in Isabelle/HOL. *J. Autom. Reason.* **40**(4), 327–356.
- Weirich, S., Yorgey, B. A. & Sheard, T. (2011). Binders unbound. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming*. New York, NY, USA: ACM, pp. 333–345.