

Types for Flexible Objects

Pottayil Harisanker Menon Zachary Palmer Alexander Rozenshteyn Scott Smith

The Johns Hopkins University

{pharisa2, zachary.palmer, scott, arozens1}@jhu.edu

Abstract

Scripting languages are popular in part due to their extremely flexible objects. These languages support numerous object features, including dynamic extension, mixins, traits, and first-class messages. While some work has succeeded in typing these features individually, the solutions have limitations in some cases and no project has combined the results.

In this paper we define TinyBang, a small typed language containing only functions, labeled data, a data combinator, and pattern matching. We show how it can directly express all of the aforementioned flexible object features and still have sound typing. We use a subtype constraint type inference system with several novel extensions to ensure full type inference; our algorithm refines parametric polymorphism for both flexibility and efficiency. We also use TinyBang to solve an open problem in OO literature: objects can be extended after being messaged without loss of width or depth subtyping and without dedicated metatheory. A core subset of TinyBang is proven sound and a preliminary implementation has been constructed.

1. Introduction

Modern scripting languages such as Python and JavaScript have become popular in part due to the flexibility of their object semantics. In addition to supporting traditional OO operations such as inheritance and polymorphic dispatch, scripting programmers can add or remove members from existing objects, arbitrarily concatenate objects, represent messages as first-class data, and perform transformations on objects at any point during their lifecycle. This flexibility allows programs to be more concise and promotes a more separated design; in Ruby, for instance, refinements and other forms of “monkey patching” can be used to extend class definitions, allowing a user to add to the behavior of library classes without modifying the relevant library.

While a significant body of work has focused on typing these flexible object operations, the solutions proposed place restrictions on how the resulting objects can be used. The Bono-Fisher object calculus [8] defines a sound means by which objects may be functionally extended: functions within the calculus can define mixins and other higher-order object transformations. But objects in this calculus must be

“sealed” before they are messaged and sealed objects cannot be extended. Later works define models which allow objects to be extended at any point during their lifecycle, but these models either do not admit depth subtyping [26] or limit the precision of types which result from object extensions [6]. As we will show, these restrictions are too prohibitive to properly model the philosophy of script programming.

In this paper, we present TinyBang, a typed core language which contains very few features: *onioning* (an asymmetric data combinator for record-like data structures), *scapes* (partial functions which pattern-match their arguments) and labeled data. TinyBang has no explicit syntax for classes, objects, inheritance, object extension, overloading, or switch/case. We show, however, that objects and flexible operations on them, including those listed above, can be tersely expressed with only scapes and onions. These operations preserve width and depth subtyping and TinyBang’s type system infers correct types for them without any programmer annotation.

We do not expect programmers to write in TinyBang directly but instead in BigBang, a language including syntax for objects, classes, and so on that is not presented here; BigBang will desugar to TinyBang. While the fundamental ideas we present here could be used to give BigBang types and runtime semantics for objects directly, we focus on object encodings because they yield a few benefits. First: TinyBang’s object calculus operations can be defined as in-language functions, avoiding the need for object-specific metatheory. Second: our encoding is simple enough to be *translucent* in that it will be possible, though atypical, to “work under the hood” to manipulate objects on a lower level when necessary. Third: this translucency ensures an elegant encoding and attests to the expressiveness of the underlying language constructs.

In order to preserve translucency in the face of flexible operations like method override and first-class messages, we design our object encoding with a *variant-based* philosophy. In contrast with record-based encodings, in which objects are records and method invocation occurs by calling the appropriate member, variant-based encodings encode objects as message-processing functions and represent method invocation as a simple function call. First-class messages are then simply the argument to that function call, and method

names are the (optional) tags placed on such arguments. The variant-based philosophy not new; Actor models [2] are untyped variant-based objects, but the variant-based encoding is more challenging to type than record-based encodings. This paper shows how, using onions and scapes, the variant-based encoding is more translucent than a traditional record-based one.

Scapes: generalizing first-class cases In our variant-based encoding, method dispatch is essentially matching on a message value. But traditional match/case expressions are monolithic blocks: they lack an explicit composition operator. We must have extensible match/case expressions to allow our encoded objects to be extensible. To accomplish this, we use a (typed) notion of first-class cases, which has been explored previously in [7]. There, a construct \oplus is defined for extending a case with a single additional clause. Our scapes generalize first-class cases by supporting composition of arbitrary case expressions, not just the addition of one clause. Since the case extension of [7] is only adding case clauses to the end of an existing case, it cannot model inheritance or object concatenation: the former requires the subclass's methods to take priority over those of the superclass and the latter requires arbitrary composition of case expressions. We believe that scapes offer the first high-level algebraic case composition operator which can directly support object inheritance for a variant encoding of objects.

Scapes are in spirit similar to typed multimethods [10, 21, 22] but multimethods are a *nominal* dispatch on class names whereas scapes are a *structural* dispatch following ML/Haskell deep patterns; many of the scripting patterns we aim to model rely on structural dispatch.

Less related is the Pattern Calculus [15], which is a more fundamental algebra in that patterns themselves are first-class entities separate from the code to be executed in case of a match. However, variable bindings in this calculus are extremely complex, so we choose the case clause as our level of abstraction.

Dependent pattern types Another problem arises when using traditional match/case expressions in a typed, variant-based object encoding: all case branches must have the same output type. This essentially requires every method of an object to return the same type, which is clearly unacceptable. TinyBang solves this fundamental problem by including a weak form of dependent type, the *dependent pattern type*, to allow for heterogeneity in case branches. A dependent pattern type allows the type of a value to condition on the variant tag of the argument. We call these types *weakly* dependent because they depend only on this variant tag and not more detailed information from the value; the advantage of this weakness is that type inference is still decidable.

Our approach is a generalization of the conditional constraints originating in [4] and elaborated in [24]. Conditional constraints partly capture this dependency but only locally and so lose the dependency information in the presence of

side effects. Our dependent pattern types fully capture this dependency. Dependent pattern types are related to types that can be given to multimethods [10], but again are structural instead of being nominal.

Onions: type-indexed records supporting asymmetric concatenation and object extension A type-indexed record is one for which contents are projected using types rather than labels [27]. For example, consider the type-indexed record `{foo = 45; bar = 22; () ; 13}` which implicitly tags the untagged elements `()` and `13` with their types. Projecting `int` from this record would yield `13` (as the other integers are labeled). Similarly, one can project unlabeled functions from a type-indexed record. Our *onions* are a form of type-indexed record; this added flexibility is handy in many situations, as will be seen below.

Additionally, since (untagged) functions can be placed in our onions, we can re-use the onion record structure to hold our scape clauses and avoid needing two different extension operations as found in [7]. As we alluded above, *asymmetric* concatenation is the key to composing scapes and thus properly defining inheritance and overriding; onions thus support asymmetric concatenation as the default. Asymmetric record concatenation was initially proposed by Wand for modeling inheritance [29], but difficulties in obtaining principal types in unification-based type inference [30] caused a switch in research focus to symmetric notions of record concatenation. Unfortunately, these notions are not amenable to modeling method overriding. In standard record-based encodings of inheritance [9], there is no first-class record extension operation; inheritance requires the superclass to be known statically. This is implicitly a consequence of the difficulty of typing asymmetric record extension. Dynamically-typed OO languages have explored first-class object extenders to significant benefit; the use of typed asymmetric record concatenation can bring this flexibility to the typed world.

Subtype constraint types TinyBang uses a subtype constraint inference based type system [4, 14, 24]; in particular, it is most closely related to [24]. Compared to [24], our system does not need row types or conditional constraints to typecheck record concatenation, incorporates type-indexed records and first-class cases, and contains a more general notion of conditional type and a more general model of parametric polymorphism. Our approach to parametric polymorphism is based on flow analysis [28, 31] and improves on previous work for expressive but efficient contour sharing. It uses a restricted form of regular expressions to finitely abstract call strings, following program analyses [20].

While our subtype constraint system infers extremely precise and expressive types, these types are also by nature difficult to read and defy modularization. TinyBang's type system is therefore a whole-program analysis. Our broader research agenda includes acquiring the typical benefits of modularity by using other techniques. Lightweight, approximate type declarations, for instance, can be used to provide

```

e ::= x | lbl e | e & e | e & - π | e & . π | e & ! π  expressions
    | e e | e ope | def x = e in e | x = e in e
    | ℤ | (characters) | () | φ -> e                literals
op ::= + | - | == | <= | >=
x ::= (alphanumeric identifiers)
lbl ::= ‘(alphanumeric identifiers)
φ ::= x : φ | x | φ                                patterns
φ ::= tprim | lbl φ | (φ & ... & φ) | fun | any
π ::= tprim | lbl | fun                            projectors
tprim ::= int | char

```

Figure 2.1. TinyBang Syntax

well-defined type interfaces and to support dynamic loading. While we believe whole-program typechecking can be made efficient and practical, the means to do so are well beyond the scope of this paper.

Outline In the next section, we give an overview of how TinyBang can encode object features. Section 3 contains the operational semantics for MicroBang, a subset of TinyBang trimmed to the key features for readability. In Section 4 the type system for MicroBang⁰, a monomorphic restriction of the full MicroBang type system, is presented; we add polymorphism in Section 5 to yield the full MicroBang type system. (Type soundness of MicroBang is established in Appendix A.) Since MicroBang lacks several features of TinyBang, we then define a series of language extensions that close this gap: Section 6 adds fully wide and deep patterns and Section 7 describes how mutable state, onion filtering, and symmetric concatenation are added to MicroBang. We conclude in Section 8.

2. Overview

This section gives an overview of the TinyBang language and of how it supports flexible object operations and other scripting features. Section 2.1 describes the TinyBang language itself and shows how simple objects can be encoded. Section 2.2 describes how self-awareness is encoded and is more flexible in TinyBang than in previous work. Section 2.3 describes flexible operations that we can define over this object encoding, such as mixins and overloading. Section 2.4 compares this variant-based encoding with its record-based analogue.

2.1 Language Features for Flexible Objects

TinyBang’s syntax appears in Figure 2.1. The expression grammar contains the key features mentioned above: scapes (written $\phi \rightarrow e$), onions (written $e \& e$), and labeled data (written $\text{‘Foo } e$ for a label constructor ‘Foo). The remaining language features are orthogonal to our object encoding but are helpful simplifying the examples below.

In order to typecheck our encoding, we use a subtype constraint-based type system. The encoding we present is not intrinsically tied to such an approach, but subtype con-

straints provide numerous advantages. Compared to row typing systems, for instance, it is not necessary to explicitly annotate type transition sites such as upcasts [25]. Most importantly, subtype constraint systems are handily modified to accommodate the more unusual language features we describe below. We defer detailed discussion of our type system to Sections 4 through 7; this section’s use of types remains informal for presentation clarity.

Scapes as methods We begin by considering the oversimplified case of an object with a single method and no fields or self-awareness. In the variant encoding, such an object is represented by a function which matches on a single case. Note that, in TinyBang, *all* functions are written $\phi \rightarrow e$, with ϕ being a *pattern* to match against the function’s argument. Combining pattern match with function definition is also possible in ML and Haskell, but we go further: there is no need for any `match` syntax in TinyBang since `match` can be encoded as a pattern and its application. We call such pattern-matching functions *scapes*. For instance, consider the following object and its invocation:

```
1 def obj = (‘double x -> x + x) in obj (‘double 4)
```

The syntax `‘double 4` is a label constructor similar to an OCaml polymorphic variant; as in OCaml, the expression `‘double 4` has type `‘double int`. The scape `‘double x -> x + x` is a function which matches on any argument containing a `‘double` label and binds its contents to the variable x . Note that the expression `‘double 4` represents a *first-class message*; the object invocation is represented with its arguments as a variant.

Unlike a traditional `match` expression, an individual scape is a function and is only capable of matching one pattern. To express general match expressions with scapes, individual scapes are appended via the *onion* operation $\&$. (This conjoiner has many other uses which are discussed below.) Given two scape expressions e_1 and e_2 , the expression $(e_1 \& e_2)$ conjoins the patterns to make a scape with the conjoined pattern, and $(e_1 \& e_2) a$ will apply the scape which has a pattern matching a ; if both patterns match a , the rightmost scape (e_2) is given priority. We can thus write a dispatch on an object with two methods simply as:

```
1 def obj = (‘double x -> x + x)
2           & (‘isZero x -> x == 0) in obj ‘double 4
```

The above shows that traditional `match` expressions can be encoded using the $\&$ operator to join a number of scapes: one scape for each case. Our scape conjunction generalizes the first-class cases of [7] to support general appending of arbitrary cases; the aforementioned work only supports adding one clause to the end and so does not allow “override” of an existing clause or “mixing” of two arbitrary sets of clauses.

Dependent pattern types The above shows how to encode an object with multiple methods as an onion of scapes. But we must be careful not to type this encoding in the

way that match/case expressions are traditionally typed. The analogous OCaml match/case expression

```
1 let obj m = (match m with
2 | 'double x -> x + x
3 | 'isZero x -> x == 0) in ...
```

will not typecheck; OCaml match/case expressions must return the same type in all case branches. (The recent OCaml 4 GADT extension mitigates this difficulty but requires an explicit type declaration, type annotations, and only works under a closed world assumption). Instead, we give the scape a dependent pattern type `('double int → int) & ('isZero int → boolean)`.¹ If the scape is applied in the context where the type of message is known, the appropriate result type is inferred; for instance, invoking this method with `'isZero 0` always produces type `int` and not type `int ∪ boolean`. Because of this dependent typing, `match` expressions encoded in TinyBang may be heterogeneous; that is, each case branch may have a different type in a meaningful way. When we present the formal type system below, we show how these dependent pattern types extend the expressiveness of conditional constraint types in a dimension critical for typing objects.

Onions as records We now show how our data conjoiner, `&`, can act like a record constructor. For example, here is how we encode objects with multi-argument methods:

```
1 def obj = ('sum ('x x & 'y y) -> x + y)
2           & ('equal ('x x & 'y y) -> x == y)
3 in obj ('sum ('x 3 & 'y 2))
```

The `'sum` label on line 3 merely wraps another value, in this case an onion of two labels: our representation of a two-label record. This record-like onion is passed to the pattern `'x x & 'y y`; here, we use `&` to also denote pattern conjunction, which requires that the value must match both subpatterns to match the overall pattern. Observe from this example how there is no hard distinction in TinyBang between records and variants: there is only one class of label and a 1-ary record is the same as a 1-ary variant.

2.2 Self-Awareness and Resealable Objects

Up to this point objects have not been able to invoke their own methods, so the encoding is incomplete. To model self-reference we build on the work of [8], where an object exists in one of two states: as a prototype, which can be extended but not messaged, or as a “proper” object, which can be messaged but not extended. A prototype may be “sealed” to transform it into a proper object, at which point it may never again be extended.

Unlike the afocited work, our encoding permits sealed objects to be extended and then *resealed*. The flexibility of

¹This section uses simplified types which improve readability by avoiding constraint sets. The actual types used in TinyBang are described in Sections 4-7.

TinyBang allows the sharp phase distinction between prototypes and proper objects to be relaxed. All object extension below will be performed on sealed objects. Object sealing in TinyBang requires no special metatheory; it is defined directly as a function `seal`:

```
1 def fixpoint = f -> (g -> x -> g g x)
2   (h -> y -> f (h h) y) in
3 def seal = fixpoint (seal -> obj ->
4   obj & (msg -> obj ('self (seal obj) & msg))) in
5 def obj = ('double x -> x + x)
6           & ('quad x & 'self self ->
7             self ('double x) + self ('double x))
8 def sObj = seal obj in ...
```

The `seal` function accepts an object as an argument and returns it with a new message handler onioned onto its right. This message handler matches *every* argument and will therefore be used for every message the returned object receives. We call this message handler the *self binding scape* because it adds a `'self` component to every message sent to the object using a reference to the object as it stood at the time the seal occurred. The self binding scape also ensures that the value in `'self` is a sealed object, allowing methods to message it normally. As a result of sealing the object, a `'quad 4` message sent to `sObj` would produce the same effect as a `'self sObj` & `'quad 4` message sent to `obj`.

Extending previously sealed objects Previous work [6, 26] allows an object to be extended after it has been messaged. In [26], this power comes at the cost of depth subtyping, which is fundamental to modern object-oriented languages; for instance, it is necessary for the Java statement `Set<? extends Number> s = new HashSet<Integer>();` to typecheck. In [6], depth subtyping is admitted but the type system is nominal rather than structural; one result of this is that extended objects have only the type of the extension and not the type of the original object. We now show how TinyBang improves on previous work without making either of these concessions.

In the self binding scape above, the value of `self` is onioned onto the *left* of the message rather than the right; this choice is purposeful given the right-precedence nature of onioning. Because of this, any value of `'self` which exists in a message passed to a sealed object takes priority over the `'self` provided by the self binding scape. Consider the following continuation of the previous code:

```
1 def sixteen = sObj 'quad 4 in
2 def obj2 = sObj & ('double x -> x) in
3 def sObj2 = seal obj2 in
4 def four = sObj2 'quad 4 in ...
```

When this code is executed, the variable `sixteen` will hold the value 16. Even after the `'quad 4` message has been sent to `sObj`, we may extend it; in this case, `obj2` redefines how `'double` messages are handled. `sObj2` represents the sealed version of this new object. When the

'quad 4 message is sent to sObj2, it is passed to obj2 with a 'self component; that is, sObj2 ('quad 4) has the same effect as obj2 ('self sObj2 & 'quad 4). Because obj2 does not change how 'quad messages are handled, this has the same effect as sObj ('self sObj2 & 'quad 4). sObj is also a sealed object which adds a 'self component to the left; thus this has the same effect as obj ('self sObj & 'self sObj2 & 'quad 4). Because any pattern match will always match the rightmost 'self, the latter-sealed object is provided in the 'self added to the message passed to the 'quad-handling method and so any messages sent from that method will be dispatched to an object which includes the extensions in sObj2.

This code successfully typechecks object extension because, while we “tie the knot” on self using seal, we leave open the possibility of future *overriding* of self; it is merely a record element and we support record field override via asymmetric concatenation.

For examples in the remainder of the paper, we will assume that seal has been defined as above.

Onioning it all together Onions also provide a natural mechanism for including fields; we simply concatenate them to the scapes that represent the methods. Consider the following object which stores and increments a counter:

```
1 def obj = seal ('x 0 &
2   ('inc _ & 'self self ->
3     ('x x -> x = x + 1 in x) self))
4 in obj 'inc ()
```

Label construction implicitly creates a mutable cell², so the 'x label is used to store the counter's current value. The bottom line invokes the scape part of obj; the 'x label does not interfere because it is not a scape. (Note that syntax () here is the “empty onion”, the 0-ary conjunction of data items.)

In this body, self is passed to the inner scape. Because seal onions the target object onto the left of the self binding method, all of the labels from the unsealed object are still visible; thus, self has the same 'x label as obj and the cell within obj's 'x label is bound to the variable x. The code x = x + 1 in x is then executed, incrementing the label's contents.

It may seem unusual that obj is, at the top level, a heterogeneous “mash” of a record field (the 'x) and a function (the scape which handles 'inc). This is sensible because onions are *type-indexed* [27], meaning that they use the types of the values themselves to identify data. When matching on an onion, the rightmost value is projected; (7 & 3) + 1 is just 4. Scapes are a special case; instead of projecting the rightmost scape, all scapes are projected and application selects the rightmost scape which matches the argument. This type-

²The core MicroBang language we formalize is immutable, but in Section 7.1 we show how the MicroBang type system may be extended to include mutable state.

indexed view is useful because it leads to more concise code as seen in Section 2.3 below.

The above counter object code is quite concise considering that it defines a self-referential, mutable counter object using no syntactic sugar whatsoever in a core language with no explicit object syntax. But as we said before, we do not expect programmers to write directly in TinyBang under normal circumstances. Here we define just a few sugarings which we use in the examples throughout the remainder of this section, although the larger BigBang language would include sugarings for each of the features we are about to mention as well. When reading the sugarings below, reflect that it is still translucent in the sense we describe above: looking at desugared code is not prohibitive for programmers desiring more control.

$$\begin{aligned} o.x &\cong ('x\ x \rightarrow x) \circ \\ o.x = e_1 \text{ in } e_2 &\cong ('x\ x \rightarrow x = e_1 \text{ in } e_2) \circ \\ \text{if } e_1 \text{ then } e_2 &\cong (('True\ _ \rightarrow e_2) \& \\ &\quad \text{else } e_3 \quad ('False\ _ \rightarrow e_3))\ e_1 \\ e_1 \text{ and } e_2 &\cong (('True\ _ \rightarrow e_2) \& \\ &\quad ('False\ _ \rightarrow 'False\ ()))\ e_1 \end{aligned}$$

Using this sugar, the third line of the counter object above can be more concisely expressed as self.x = self.x + 1 in self.x.

2.3 Flexible Object Operations

The above shows both a typed, variant-based object encoding for TinyBang as well as two flexible object features: first-class messages are trivial to encode and objects may be extended after they are messaged without compromising expressiveness. We now focus on typed encodings of flexible object operations as discussed in Section 1. Traditionally, these abstractions are defined in a first-order sense; inheritance, for instance, can only be expressed if the type of the parent class is statically known. In contrast, TinyBang's encoding can type higher-order abstractions.

We show our encoding in terms of objects rather than classes for simplicity; applying these concepts to classes is straightforward, and working at the object layer also shows how TinyBang can express (functional) object extension. For clarity, we use the sugar defined above.

Default arguments Many scripting languages include a notion of a “default argument”: an argument which may be specified to a method but which will take on a default value if it is missing. Onions encode this behavior quite easily. For instance, consider:

```
1 def obj = seal
2   (('add ('x x & 'y y) -> x + y)
3     & ('sub ('x x & 'y y) -> x - y) ) in
4 def dflt = obj -> obj &
5   ('add a -> obj ('add ('x 1 & a))) in
6 def obj2 = dflt obj in
7 obj2 ('add ('y 3));
```

```
8 obj2 ('add ('x 7 & 'y 2))
```

Given an object which accepts a message `'add`, `df1t` will, in all `'add` messages sent without an `'x` component, include `'x 1`. Because the `'x 1` is onioned onto the left of `a`, it will have no effect if an `'x` is provided in the message. Thus, the invocations on the last lines will yield 4 and 9, respectively. Note also that, while most languages only permit default arguments in a first-order sense, TinyBang can encode default arguments as a higher-order transformation.

Overloading The pattern-matching semantics of scapes also provide a translucent mechanism whereby function (and thus method and operator) overloading can be defined. We might originally define negation on the integers as

```
1 def neg = x:int -> 0 - x in ...
```

Later code could then extend the definition of negation to include boolean values. Because operator overloading assigns new meaning to an existing symbol, we redefine `neg` to include all of the behavior of the old `neg` as well as new cases for `'True` and `'False`:

```
1 def neg = neg &! ('True _ -> 'False ())
2 &! ('False _ -> 'True ()) in ...
```

The `&!` operator is TinyBang's *symmetric* concatenation syntax, ensuring that new overloadings produce a type error if they override existing definitions; `&!` is not in our core MicroBang type system but is added as an extension in Section 7.3. Negation is now overloaded: `neg 4` evaluates to `-4`, and `neg 'True ()` evaluates to `'False ()` due to how scape application matches patterns. Overloading as defined above shares similarities with the $\lambda&$ -calculus [10] and typed multimethod-based language designs [21, 22]; these projects, however, dispatch on nominal (class) names and not on structural information as we do. Note that our type system assumes whole-program compilation. Many of the subtleties in implementations of overloading arise from attempts to modularize it [5, 21]; we get to live in a simpler universe.

Mixins The following example shows how a simple two-dimensional point object can be combined with a `mixin` providing extra methods:

```
1 def point = seal ('x 0 & 'y 0
2 & ('l1 _ & 'self self -> self.x + self.y)
3 & ('isZero _ & 'self self ->
4 self.x == 0 and self.y == 0) in
5 def mixin = (('near _ &
6 'self self -> (self 'l1 ()) < 4) in
7 def mixedPoint = seal (point & mixin) in
8 mixedPoint 'near ()
```

The `point` variable is our original point object. The `mixin` is merely a scape which calls the value passed as `self`. Because an object's methods are just scapes onioned together,

onioning the mixin into the point object is sufficient to produce the resulting mixed point; the `mixedPoint` variable contains an onion with `x` and `y` fields as well as all three scapes.

The above example is well-typed in TinyBang; parametric polymorphism is used to allow `point`, `mixin`, and `mixedPoint` to have different self-types. The `mixin` variable, the interesting part of the above code, has roughly the type `((('near unit & 'self α) \rightarrow boolean` where α is an object capable of receiving the `'l1` message and producing an `int`". `mixin` can be onioned with any object that satisfies these properties. If the object does not have these properties, a type error will result when the `'near` message is passed; for instance, `(seal mixin) ('near ())` is not typeable because `mixin`, the value of `self`, does not have a scape which can handle the `'l1` message.

TinyBang mixins are first-class values; the actual mixing need not occur until runtime. For instance, the following code selects a weighting metric to mix into a point based on some runtime condition `cond`.

```
1 def cond = (runtime boolean) in
2 def point = (as above) in
3 def w1 = ('weight _ & 'self self ->
4 self.x + self.y) in
5 def w2 = ('weight _ & 'self self ->
6 self.x - self.y) in
7 def mixedPoint = seal
8 (point & (if cond then w1 else w2)) in
9 mixedPoint 'weight ()
```

Traits can be encoded in a similar manner: each trait is an onion and objects are created by their concatenation. The encoding becomes slightly more complex when resolving conflicting methods; we omit the full discussion for brevity.

Inheritance, classes, and subclasses Typical object-oriented constructs can be defined in almost the same way. Object inheritance is accomplished in a fashion similar to mixins but with one exception: a variable `super` is bound to the original object and captured in the closure of the inheriting objects methods, allowing it to be reached for static dispatch. This allows overriding methods to invoke the methods they have overridden. The flexibility of the previously-defined `seal` function permits us to ensure that the inheriting object is used for future dispatches even in calls to overridden methods. Classes are defined quite simply as objects which generate other objects; subclasses are merely extensions of those object generating objects. We forgo any examples here for brevity.

2.4 A Record-Based Comparison

We have shown how an object model amenable to flexible operations can be defined using onions and scapes. In our encoding, we chose a variant-based approach. Since records and variants are natural duals, however, we now offer a high-level comparison with a record-based encoding.

TinyBang’s scapes and onions can also be used to build a pure record-based encoding of objects that would include support for first-class messages. Unlike previous encodings of first-class messages [23, 24], TinyBang provides a view which uses subtyping and dependent pattern types and does not require conditional constraints, row types, or a higher-kinded system. Such a record-based encoding is the true dual of the variant-based encoding presented above; the variant-encoded message ‘`msg arg`, for example, dualizes to `obj -> obj .msg arg` in the record-encoded system.

It is also possible to use first-class labels to form a record-based encoding of objects; first-class messages are then easily encoded with first-class labels. Fluid object types [13] are a recent system which uses this approach to infer types in scripting languages such as JavaScript. (Note that we do intend to investigate in the future whether the addition of first-class labels to TinyBang will give us additional added flexibility; they should not be theoretically challenging to add.)

We focus on a variant-based encoding because it leads to a more translucent and direct syntax. For instance, the variant first-class message ‘`msg arg` may be directly pattern matched, but the record-based function representation is harder to deconstruct. The flexible object operations we define above are also more opaque in the record-based view. In the above encoding, one adds a method `mthd` to an object `obj` simply by writing `obj & mthd`. In the record-based view, where methods are labeled functions which are projected when they are called, one must write the considerably more cumbersome

```
1 obj & 'mthd (( (_ -> mthd) &
2 ('mthd old -> old & mthd)) obj)
```

to account for the fact that the original object may or may not already have a ‘`mthd` component. The implementation of a `seal` method for the record-based view is very complex; we omit it here for brevity.

Much of the terseness in the TinyBang encoding arises due to the fact that it cleanly puts fields and methods into different syntactic sorts: labels and scapes, respectively. It is possible to get by with only one of these sorts much like it is possible to use only one of \wedge or \vee in Boolean logic: DeMorgan’s law allows one to be written in terms of the other, but the resulting expressions are opaque and unintuitive. In this sense, the encoding that we provide is inspired by the variant philosophy but is in fact a hybrid encoding in which methods are on the variant side of the record-variant duality and fields are on the record side. This hybrid view provides a very natural encoding of self-awareness (in the form of the `seal` function) without complicating field access.

3. MicroBang Operational Semantics

We now begin formal discussion of the operational semantics and typing of the language features presented in Section 2. Because the language is reasonably complex, we begin our discussion with MicroBang, a simplified TinyBang

$e ::= \overline{x=r}$	<i>expressions</i>
$r ::= v \mid x \mid x \text{ op } x$	<i>redices</i>
$v ::= x \mid \mathbb{Z} \mid \text{lbl } x \mid x \& x \mid () \mid \phi \rightarrow e$	<i>values</i>
$\phi ::= x : \varphi$	<i>patterns</i>
$\varphi ::= \text{int} \mid \text{lbl } x \mid \text{any}$	<i>primary patterns</i>
x	<i>variables</i>
$\text{op} ::= + \mid ==$	<i>operators</i>
$\text{lbl} ::= \text{'[a-zA-Z0-9_]}^+$	<i>label names</i>
$\pi ::= \text{int} \mid \text{lbl} \mid \text{fun}$	<i>projectors</i>

Figure 3.1. MicroBang Syntax

lacking complex patterns, mutable state, and the advanced onion operators $\&-$, $\&.$, and $\&!.$ This section defines a small-step operational semantics for MicroBang and subsequent sections address typing.

Notation For a given grammatic construct $*$, we let $[*_1, \dots, *_n]$ denote a list of $*$ ’s, often using the equivalent shorthand $*_{\square}^n$; the “ \square ” indicates which constructs are indexed, and the n and \square can be elided when they are obvious or not needed. Operator \parallel denotes list concatenation. For sets, we use similar indexing notation: $*_{\square}^{\downarrow}$ abbreviates $\{*_1, \dots, *_n\}$. We only use this notation on *finite* lists and sets in this paper.

3.1 MicroBang Grammar

We present the grammar for MicroBang in Figure 3.1. This grammar is written in A-normal form to simplify the small-step operational semantics presented below. The A-translation from a more traditional recursive grammar is standard and is not defined here for sake of brevity. As an example, the expression $4 + 5 == 9$ would translate to the A-normal expression $[x1 = 4, x2 = 5, x3 = x1+x2, x4 = 9, x5 = x3==x4]$: each program point is assigned its own variable and variable names are unique. For the remainder of the paper, we will assume that we work only over ANF programs e with unique variable names.

3.2 Operational Semantics

The operational semantics of MicroBang proceed by evaluating each redex in the expression from left to right. Let E be of the form $\overline{x=v}$; that is, E is the subset of e for which each redex is already a value. E constitutes both the environment and the evaluation context in our semantics. For such an $E = \overline{x=v}$, we define $E(x) = v$ to be a lookup function defined when $x = x_i$ for some $1 \leq i \leq n$. When v_i is a non-variable value, $v = v_i$; otherwise, $v = E(v_i)$ recursively.

Projection The projectors π in the grammar represent the different outermost value sorts. We use notation $v \in \pi$ to denote that the projector π matches the outermost structure of a given value v ; for example, ‘ $\mathbf{A} \ x \in \mathbf{A}$ ’. The special projector `fun` matches all scapes (and only scapes). No projector

matches onions directly. We then define a simple function $E\downarrow_{\pi}^*(x)$ which projects from a given value x all values which match the projector π ; there can be more than one result because if x is an onion, the result is the list of all values which match the projector.

Definition 3.1 (Value Projection).

$$E\downarrow_{\pi}^*(x) = \begin{cases} [] & \text{if } E(x) \text{ is non-onion, } E(x) \notin \pi \\ [E(x)] & \text{if } E(x) \text{ is non-onion, } E(x) \in \pi \\ E\downarrow_{\pi}^*(x_1) \parallel E\downarrow_{\pi}^*(x_2) & \text{if } E(x) = x_1 \& x_2 \end{cases}$$

In most cases, we only need the highest priority (rightmost) value from the onion. For those cases, we define the following function.

Definition 3.2 (Single Value Projection).

$$E\downarrow_{\pi}(x) = v_n \text{ where } E\downarrow_{\pi}^*(x) = \vec{v}$$

Note that the function $E\downarrow_{\pi}(x)$ is partial because it is undefined for cases in which the value x does not contain any values which match the projector.

Compatibility The next relation we use in our definition of MicroBang’s operational semantics is compatibility. Compatibility holds only when a given value matches a given pattern. If it does, then the match is with respect to a list of bindings E_0 which describes how the contents of the argument are assigned to the pattern’s variables. For convenience, we define this relation on both patterns (ϕ) and primary patterns (φ). We define $x \preceq_E \phi \setminus E_0$ by cases on the structure of the pattern as follows:

Definition 3.3 (Value Compatibility).

$$\begin{array}{ll} x_1 \preceq_E x_2 : \varphi \setminus E_0 \parallel [x_2 = x_1] & \text{if } x_1 \preceq_E \varphi \setminus E_0 \\ x_1 \preceq_E \text{int} \setminus [] & \text{if } E\downarrow_{\text{int}}(x_1) \in \mathbb{Z} \\ x_1 \preceq_E \text{lbl } x_2 \setminus [x_2 = x_3] & \text{if } E\downarrow_{\text{lbl}}(x_1) = \text{lbl } x_3 \\ x_1 \preceq_E \text{any} \setminus [] & \text{always} \end{array}$$

We write $x \not\preceq_E \phi$ to indicate that there exists no E_0 such that $x \preceq_E \phi \setminus E_0$.

We then use the above compatibility relation to define an application-driven form of compatibility. We write $x \preceq_E \overline{\phi \rightarrow e} \setminus E_0; e$ to denote that the argument x can be applied to one of the scapes in $\overline{\phi \rightarrow e}$; in this case, E_0 is the set of bindings from x to the pattern and e is the body of the scape to be applied. We define this relation below by induction on the length of the list of scapes. This definition will always use the rightmost applicable scape.

Definition 3.4 (Value Application Compatibility).

$$\begin{array}{ll} x \preceq_E [] \setminus E_0; e & \text{is false} \\ x \preceq_E \overline{\phi \rightarrow e} \parallel [\phi' \rightarrow e'] \setminus E_0; e' & \text{if } x \preceq_E \phi' \setminus E_0 \\ x \preceq_E \overline{\phi \rightarrow e} \parallel [\phi' \rightarrow e'] \setminus E_0; e'' & \text{if } x \not\preceq_E \phi' \wedge x \preceq_E \overline{\phi \rightarrow e} \setminus E_0; e'' \end{array}$$

Small-Step Semantics Using the relations above, we now define small-step computation for MicroBang. We take $\alpha(e)$ to be an α -conversion function which renames all bound variables in e to a fixed set of fresh names relative to the current context. Our single-step relation is then defined over closed e with unique variable names as follows; we give intuitions below.

Definition 3.5 (MicroBang Small-Step Semantics).

$$\begin{array}{l} E \parallel [x = x_1 + x_2] \parallel e \longrightarrow^1 E \parallel [x = n] \parallel e \\ \quad \text{when } E\downarrow_{\text{int}}(x_1) = n_1, E\downarrow_{\text{int}}(x_2) = n_2, n_1 + n_2 = n \\ E \parallel [x = x_1 == x_2] \parallel e \longrightarrow^1 E \parallel [x' = (), x = \text{‘True } x’}] \parallel e \\ \quad \text{when } E\downarrow_{\text{int}}(x_1) = n_1, E\downarrow_{\text{int}}(x_2) = n_2, n_1 = n_2, x' \text{ fresh} \\ E \parallel [x = x_1 == x_2] \parallel e \longrightarrow^1 E \parallel [x' = (), x = \text{‘False } x’}] \parallel e \\ \quad \text{when } E\downarrow_{\text{int}}(x_1) = n_1, E\downarrow_{\text{int}}(x_2) = n_2, n_1 \neq n_2, x' \text{ fresh} \\ E \parallel [x = x_1 x_2] \parallel e'' \longrightarrow^1 E \parallel \alpha(E_0 \parallel e' \parallel [x' = r]) \parallel [x = \alpha(x')] \parallel e'' \\ \quad \text{when } E\downarrow_{\text{fun}}^*(x_1) = \vec{v}, x_2 \preceq_E \vec{v} \setminus E_0; e' \parallel [x' = r] \end{array}$$

We then define small-step computation $e_0 \longrightarrow^* e_n$ to hold when $e_0 \longrightarrow^1 \dots \longrightarrow^1 e_n$ for some $n \geq 0$. Note that $e \longrightarrow^* E$ means that computation has resulted in a final value. We write $e \not\longrightarrow^1$ iff there is no e' such that $e \longrightarrow^1 e'$; observe $E \not\longrightarrow^1$ for any E .

Here are some intuitions for the above operational semantics rules. In the first case, we observe that the first unevaluated redex is an addition: $x_1 + x_2$. We first use projection to obtain the integer from each of x_1 and x_2 ; we then replace the redex with the sum of the results. If e.g. x_1 is defined as an integer (or as a variable which is inductively defined as an integer) in E , then projection is straightforward. If x_1 is an onion containing integers, then the rightmost integer is used; that is, $1 \& 2$ projects the integer 2 in keeping with the examples in Section 2. If x_1 contains no integers, then projection is undefined and no small step can occur, leaving the evaluation “stuck” (that is, $e \not\longrightarrow^1$). The small step evaluation rules for equality use the same approach as addition.

In the final case, the first unevaluated redex is an application $x_1 x_2$. We begin by projecting all scapes from x_1 ; we then use the compatibility relation to determine which scape matches the argument x_2 first. If no such scape exists, then the argument does not match any of the patterns and evaluation is stuck; otherwise, we find the pattern bindings E_0 and the body $e' \parallel [x' = r]$ of the matching scape. We describe the body of the scape in this manner because x' , as the last variable in the body, describes the result of the scape. We then replace the original application expression $[x = x_1 x_2]$ with $\alpha(E_0 \parallel e' \parallel [x' = r]) \parallel [x = \alpha(x')]$: the pattern bindings for the argument, the body of the matched scape, and a final definition to copy the result of the scape into x . The variables in these terms have been freshened by α -conversion wherever they originated from the body of the scape in order to ensure that names in the new expression are still unique; note that the freshening will freshen both occurrences of x' to the same fresh variable here.

We have defined a small-step operational semantics for MicroBang; we will now discuss how it may be typed.

4. A Monomorphic MicroBang Type System

In order to simplify the presentation of MicroBang’s type system, we first present a monomorphic type system we call MicroBang⁰. The MicroBang⁰ type system includes all features of MicroBang syntax and types except polymorphism. This allows us to show how scape application is typed with-

α		<i>type variables</i>
$\alpha ::= \overline{\alpha}$		<i>type variable sets</i>
$\Gamma ::= \overline{x : \alpha}$		<i>contexts</i>
$\tau ::= \mathbf{int} \mid \mathbf{lbl} \alpha \mid \alpha \& \alpha \mid () \mid \tau\phi \rightarrow \alpha \setminus C$		<i>types</i>
$\tau\phi ::= \alpha \sim \tau\phi$		<i>pattern types</i>
$\tau\varphi ::= \mathbf{int} \mid \mathbf{lbl} \alpha \mid \mathbf{any}$		<i>primary pattern types</i>
$c ::= \tau <: \alpha \mid \alpha <: \alpha \mid \alpha \alpha <: \alpha \mid \alpha \text{ op } \alpha <: \alpha$		<i>constraints</i>
$C ::= \overline{c}$		<i>constraint sets</i>

Figure 4.1. MicroBang⁰ Type Grammar

out the distraction of the complex polymorphism model. Polymorphic MicroBang is presented in Section 5.

Typechecking in MicroBang⁰ consists of two phases: initial derivation, in which the program is translated to a type variable and a set of initial constraints over that variable; and closure, during which the initial constraint set is deductively closed over a logical system and then checked for consistency. We begin our discussion with the type and constraint grammar shown in Figure 4.1.

The types τ in the type grammar represent concrete lower bounds in MicroBang⁰'s type system. They are, informally and respectively, primitives, labels of cells, the union of two other types, the empty union type, and functions. MicroBang⁰'s constraint grammar is restricted in that all subtype constraint upper bounds are type variables; this canonical form is nonetheless complete and facilitates reasoning. Note that we use $\alpha_1 \alpha_2 <: \alpha_3$ which is equivalent to the more standard notation $\alpha_1 <: \alpha_2 \rightarrow \alpha_3$.

4.1 Initial Derivation

Type checking in MicroBang⁰ is performed against a fixed, closed expression e . The first step of typechecking is to translate e into a type variable and a set of initial constraints. In order to define this operation concisely, we first give some supplementary definitions.

Initial type derivation is performed with respect to a context Γ . We still operate under the assumption from Section 3.1 that variable bindings in our programs are unique; thus, new bindings can be safely introduced to a context with set union. Throughout type derivation, we also require fresh type variables. We write $\check{\alpha}_i$ to denote the i th fresh variable relative to a given point in translation.

We define initial derivation as three functions: $\llbracket e \rrbracket_E^\Gamma$, $\llbracket x = r \rrbracket_R^\Gamma$, and $\llbracket \phi \rrbracket_P^\Gamma$. These functions define derivation over expressions, redex definitions, and patterns, respectively. Note that the letters E, R, and P are not variables; they distinguish the sort of initial derivation which is being performed.

We write $\llbracket e \rrbracket_E^\Gamma = \langle \alpha, C \rangle$ to indicate that an expression e in context Γ translates to the type α constrained by C . Expressions contain redex definitions; we write $\llbracket x = r \rrbracket_R^\Gamma = c$ to indicate that the redex definition $x = r$ translates to the constraint c . Because these definitions may contain scapes

$$\begin{aligned} \llbracket [x = r] \rrbracket_E^\Gamma &= \langle \check{\alpha}_1, \{ \llbracket x = r \rrbracket_R^{\Gamma \cup \Gamma'} \} \rangle \\ &\quad \text{where } \Gamma' = \{ x : \check{\alpha}_1 \} \\ \llbracket [x = r] \parallel e \rrbracket_E^\Gamma &= \langle \alpha', \{ \llbracket x = r \rrbracket_R^{\Gamma \cup \Gamma'} \} \cup C \rangle \\ &\quad \text{where } \Gamma' = \{ x : \check{\alpha}_1 \}, \llbracket e \rrbracket_E^{\Gamma \cup \Gamma'} = \langle \alpha', C \rangle \end{aligned}$$

$$\begin{aligned} \llbracket [x = x'] \rrbracket_R^\Gamma &= \Gamma(x') <: \Gamma(x) \\ \llbracket [x = \mathbf{lbl} x'] \rrbracket_R^\Gamma &= \mathbf{lbl} \Gamma(x') <: \Gamma(x) \\ \llbracket [x = x' \& x''] \rrbracket_R^\Gamma &= \Gamma(x') \& \Gamma(x'') <: \Gamma(x) \\ \llbracket [x = \phi \rightarrow e] \rrbracket_R^\Gamma &= \tau\phi \rightarrow \alpha \setminus C <: \Gamma(x) \\ &\quad \text{where } \begin{cases} \llbracket \phi \rrbracket_P^\Gamma = \langle \Gamma', \tau\phi \rangle, \\ \llbracket e \rrbracket_E^{\Gamma \cup \Gamma'} = \langle \alpha, C \rangle, \end{cases} \\ \llbracket [x = ()] \rrbracket_R^\Gamma &= () <: \Gamma(x) \\ \llbracket [x = \mathbf{Z}] \rrbracket_R^\Gamma &= \mathbf{int} <: \Gamma(x) \\ \llbracket [x = x' x''] \rrbracket_R^\Gamma &= \Gamma(x') \Gamma(x'') <: \Gamma(x) \\ \llbracket [x = x' \text{ op } x''] \rrbracket_R^\Gamma &= \Gamma(x') \text{ op } \Gamma(x'') <: \Gamma(x) \\ \llbracket [x_1 : \mathbf{int}] \rrbracket_P^\Gamma &= \langle \{ x_1 : \check{\alpha}_1 \}, \check{\alpha}_1 \sim \mathbf{int} \rangle \\ \llbracket [x_1 : \mathbf{lbl} x_2] \rrbracket_P^\Gamma &= \langle \{ x_1 : \check{\alpha}_1, x_2 : \check{\alpha}_2 \}, \check{\alpha}_1 \sim \mathbf{lbl} \check{\alpha}_2 \rangle \\ \llbracket [x_1 : \mathbf{any}] \rrbracket_P^\Gamma &= \langle \{ x_1 : \check{\alpha}_1 \}, \check{\alpha}_1 \sim \mathbf{any} \rangle \end{aligned}$$

Figure 4.2. MicroBang⁰ Initial Type Derivation

and each scape has a pattern, we write $\llbracket \phi \rrbracket_P^\Gamma = \langle \Gamma', \tau\phi \rangle$ to indicate that a pattern ϕ in context Γ translates to another context Γ' which contains the bindings for the pattern as well as a pattern type $\tau\phi$. Using this notation, we define the initial type derivation process in Figure 4.2.³

Much of this translation process is direct, substituting variables with their corresponding type variables and modeling value assignment as a lower bound. The translation of a scape creates a type which captures the constraints derived from its body. These constraints are introduced during closure only if the scape is used; otherwise, they have no effect on typechecking.

4.2 Constraint Closure

The second step in the typechecking process is to deductively close over the set of constraints produced by derivation. To define the closure relation, we must first specify other relations which mirror those used in the small-step semantics in Section 3. We begin by defining *concretization*, a simple relation which locates the lower bounds for a given type via transitivity chains in the constraint set.

Definition 4.1 (Concretization). $\tau \overset{C}{<}^* \alpha_n$ iff $\{ \tau <: \alpha_0, \alpha_0 <: \alpha_1, \dots, \alpha_{n-1} <: \alpha_n \} \subseteq C$ where $n \geq 0$.

Unlike the lookup function $E(x)$ from the operational semantics, concretization is a relation. In cases where the precise type is not statically known (such as a condition, the possible types of which are **True** () or **False** ()), we approximate it by using a union of the possible types. As a

³ Recall from the top of Section 3 that we use \square for indexing positions. For instance, we would write $x_{\square+1}^n$ to denote $[x_1 + 1, \dots, x_n + 1]$.

result, a given type variable may have many lower bounds and so there may be many concretizations for a single α_n .

Closure must also create boolean labels which contain the empty union (\emptyset) similar to how the operational semantics create a fresh x' when a boolean is created. To address this, we define a mapping $\text{ETV}(\cdot)$ to provide fresh variables:

Definition 4.2 (Empty Union Variables). For all α in the initial derivation, $\text{ETV}(\alpha) = \alpha'$ where α' is a fresh type variable not in the initial derivation.

Projection Next we define projection, a relation which extracts from a type variable all type lower bounds matching a given projector. We write $C \vdash \alpha \xrightarrow{\pi} \vec{\tau}$ to denote that $\vec{\tau}$ are the non-onion components of the lower-bounding types of α which match the projector π . Note that this relation is not a function; if α has multiple lower bounds, for instance, it will relate multiple type lists. In our definition, we use a relation $\tau \sqsubseteq \pi$ to determine if a non-onion type matches a given projector. It is true in exactly the following cases:

Definition 4.3 (Projector Match).

$$\begin{aligned} \text{int} &\sqsubseteq \text{int} \\ \text{lbl } \alpha &\sqsubseteq \text{lbl} \quad \text{for any } \alpha \\ \tau_\phi \rightarrow \alpha \setminus C &\sqsubseteq \text{fun} \quad \text{for any } \tau_\phi, \alpha, C \end{aligned}$$

Using this predicate, we define projection as follows:

Definition 4.4 (Projection).

$$\begin{aligned} C \vdash \alpha' \xrightarrow{\pi} [] &\quad \text{if } \tau \stackrel{C}{<} \alpha', \tau \not\sqsubseteq \pi, \tau \text{ is non-onion} \\ C \vdash \alpha' \xrightarrow{\pi} [\tau] &\quad \text{if } \tau \stackrel{C}{<} \alpha', \tau \sqsubseteq \pi \\ C \vdash \alpha' \xrightarrow{\pi} \vec{\tau}_1 \parallel \vec{\tau}_2 &\quad \text{if } \alpha_1 \& \alpha_2 \stackrel{C}{<} \alpha', \\ &\quad C \vdash \alpha_1 \xrightarrow{\pi} \vec{\tau}_1, C \vdash \alpha_2 \xrightarrow{\pi} \vec{\tau}_2 \end{aligned}$$

Projection from a non-onion type is quite simple: produce a singleton list containing the type if it matches the projector and produce an empty list if it does not. Projection from an onion involves concatenating the projection from each side; the result is a priority-ordered list. This relation is decidable, although calculating it efficiently is technically challenging in some rare cases. We discuss this in Appendix B.

We often want only the highest priority (rightmost) type from an onion which matches a given projector; this reflects the rightmost priority of onioning. For this purpose, we overload projection syntax with a single projection relation:

Definition 4.5 (Single Projection).

$$C \vdash \alpha \xrightarrow{\pi} \tau' \text{ holds iff } C \vdash \alpha \xrightarrow{\pi} \vec{\tau} \parallel [\tau'] \text{ holds for some } \vec{\tau}.$$

Note that this relation is partial; it does not hold when there are no types in any concrete form of α matching the projector π .

Compatibility We now define a type-level compatibility relation to model the compatibility relation of operational semantics. To do so, we define \dot{C} to range over ‘‘possible constraint sets’’; each \dot{C} is either a set of constraints C or the special symbol \ast . We define union over possible constraint sets as $\dot{C}_1 \cup \dot{C}_2 = C_1 \cup C_2$ when both $\dot{C}_1 = C_1$ and $\dot{C}_2 = C_2$; we define it as \ast when either \dot{C}_1 or \dot{C}_2 is \ast .

We write $\alpha \preceq_C \tau_\phi \setminus \dot{C}$ to describe how a type variable α may match the pattern τ_ϕ . When \dot{C} is C , this relation

indicates that those constraints allow the type information from α to flow into the variables in τ_ϕ (similar to how E_0 describes how data flows from an argument x into a pattern ϕ in the small-step semantics). When \dot{C} is \ast , this relation indicates that the argument does not match the pattern. The \ast notation is helpful here because α may have multiple lower bounds; it is thus possible for one of the forms of α to match the pattern while another does not.

We define type-level compatibility by casing on the pattern type. As in the operational semantics, we define compatibility on both patterns and primary patterns.

Definition 4.6 (Compatibility).

$$\begin{aligned} \alpha_1 \preceq_C \alpha_2 \sim \tau_\phi \setminus \{\alpha_1 <: \alpha_2\} \cup \dot{C} &\quad \text{if } \alpha_1 \preceq_C \tau_\phi \setminus \dot{C} \\ \alpha_1 \preceq_C \text{int} \setminus \emptyset &\quad \text{if } C \vdash \alpha_1 \xrightarrow{\text{int}} \text{int} \\ \alpha_1 \preceq_C \text{lbl } \alpha_2 \setminus \{\alpha_3 <: \alpha_2\} &\quad \text{if } C \vdash \alpha_1 \xrightarrow{\text{lbl}} \text{lbl } \alpha_3 \\ \alpha_1 \preceq_C \text{any} \setminus \emptyset &\quad \text{always} \\ \alpha_1 \preceq_C \text{int} \setminus \ast &\quad \text{if } C \vdash \alpha_1 \xrightarrow{\text{int}} [] \\ \alpha_1 \preceq_C \text{lbl } \alpha_2 \setminus \ast &\quad \text{if } C \vdash \alpha_1 \xrightarrow{\text{lbl}} [] \end{aligned}$$

Just as in the small-step semantics in Section 3, we now use this compatibility relation to define an application-driven form of compatibility. We write $\alpha \preceq_C (\tau_\phi' \rightarrow \alpha'_\square \setminus C'_\square) \setminus \alpha''; C''$ to indicate that α , the type in the argument position of an application, is compatible with the pattern of one of the scape types in the specified scape list. α'' and C'' represent the return type and constraint set of the matched pattern’s scape. We define this relation as follows:

Definition 4.7 (Application Compatibility).

$$\begin{aligned} \alpha \preceq_C [] \setminus \alpha''; C'' &\quad \text{is false} \\ \alpha \preceq_C \vec{\tau} \parallel [(\tau_\phi' \rightarrow \alpha' \setminus C')] \setminus \alpha''; C'' \cup C' &\quad \text{if } \alpha \preceq_C \tau_\phi' \setminus C'' \\ \alpha \preceq_C \vec{\tau} \parallel [(\tau_\phi' \rightarrow \alpha' \setminus C')] \setminus \alpha''; C'' &\quad \text{if } \alpha \preceq_C \tau_\phi' \setminus \ast, \\ &\quad \alpha \preceq_C \vec{\tau} \setminus \alpha''; C'' \end{aligned}$$

Again following the small-step semantics, this relation will always use the rightmost available scape. We write $\alpha \not\preceq_C \tau_\phi' \rightarrow \alpha'_\square \setminus C'_\square$ to denote that $\forall i \in \{1..n\}. \alpha \preceq_C \tau_\phi'_i \setminus \ast$; that is, that every pattern could fail to match α .⁴

Constraint Closure We are now equipped to define the constraint closure relation described above. We write $C_{\text{IN}} \xrightarrow{\text{CL}} C_{\text{OUT}}$ to denote that a single step of closure maps C_{IN} onto C_{OUT} . The definition of this relation appears in Figure 4.3, which uses the following notational sugar. In general, each rule assumes constraint closure to be performed over some set of constraints C_{IN} . We write a constraint set C_{NEW} as the conclusion of a rule to abbreviate $C_{\text{IN}} \xrightarrow{\text{CL}} C_{\text{IN}} \cup C_{\text{NEW}}$. We may also write a single constraint c as the conclusion of a rule to abbreviate $\{c\}$. If a constraint c appears in the premise of a rule, we take it to mean $c \in C_{\text{IN}}$.

⁴This definition of compatibility is sufficient for soundness but is overly conservative due to a form of union misalignment; it is possible for each pattern to fail to match α in a different way but for no single instance of α to fail to match all patterns. We solve this problem by modifying the relation to include a *fibration*, which describes concrete type structure, and then ensuring alignment between these fibrations. We defer presentation of this solution until Section 6 to simplify the initial type system we are describing here.

$$\begin{array}{c}
\text{APPLICATION} \\
\frac{\alpha_1 \alpha_2 <: \alpha_3 \quad C_{\text{IN}} \vdash \alpha_1 \xrightarrow{\text{fun}} \vec{\tau} \quad \alpha_2 \preceq_{C_{\text{IN}}} \vec{\tau} \setminus \alpha_4; C'}{C' \cup \{\alpha_4 <: \alpha_3\}} \\
\\
\text{INTEGER ADDITION} \\
\frac{\alpha_1 + \alpha_2 <: \alpha_3 \quad C_{\text{IN}} \vdash \alpha_1 \xrightarrow{\text{int}} \text{int} \quad C_{\text{IN}} \vdash \alpha_2 \xrightarrow{\text{int}} \text{int}}{\text{int} <: \alpha_3} \\
\\
\text{INTEGER EQUALITY} \\
\frac{\alpha_1 == \alpha_2 <: \alpha_3 \quad \alpha'_3 = \text{ETV}(\alpha_3) \quad C_{\text{IN}} \vdash \alpha_1 \xrightarrow{\text{int}} \text{int} \quad C_{\text{IN}} \vdash \alpha_2 \xrightarrow{\text{int}} \text{int}}{\{\text{True } \alpha'_3 <: \alpha_3, \text{False } \alpha'_3 <: \alpha_3, () <: \alpha'_3\}}
\end{array}$$

Figure 4.3. MicroBang⁰ Constraint Closure

Each constraint closure rule’s first premise requires the presence of a constraint for a particular grammar form. The Integer Addition rule, for instance, requires a constraint of the form $\alpha_1 + \alpha_2 <: \alpha_3$; Figure 4.2 shows us that each redex of the form $e_1 + e_2$ produces such a constraint. We then use projection to confirm that each type variable has `int` as a lower bound; if so, then we can conclude $\text{int} <: \alpha_3$, which indicates that the result also has a lower bound of `int`. (We consider type errors, such as when either side has a non-`int` lower bound, when we detect contradictions below.) The Integer Equality rule works in a similar fashion but produces two constraints; this is because the result is not known until runtime and a union of the types is a conservative approximation. In this rule, α' is the variable associated with α_3 which models the fresh x' in the operational semantics.

The application rule is also similar to the integer rules. Given a constraint that indicates that a function call occurred, we first use projection to determine which functions might have been available. We then use the compatibility relation to determine which of the scapes might have been suitable for the argument; for such a suitable scape, we conclude the constraints describing the body of that scape as well as an additional constraint $\alpha_3 <: \alpha_2$ which ensures that the result type of the scape’s computation flows into the output point for the call site.

In parallel to the operational semantics, we define a relation $C_0 \xrightarrow{\text{Cl}_y^*} C_n$ to hold when $C_0 \xrightarrow{\text{Cl}_y^1} \dots \xrightarrow{\text{Cl}_y^1} C_n$ for $n \geq 0$.

Contradictions After constraint closure is complete, detecting contradictions simply involves finding inconsistencies in the constraint set. We define such inconsistencies as follows:

Definition 4.8 (Inconsistency). A set of constraints C is *inconsistent* iff either

$$\begin{array}{l}
\exists \alpha_1 \alpha_2 <: \alpha_3 \in C \text{ such that } C \vdash \alpha_1 \xrightarrow{\text{fun}} \vec{\tau} \text{ and } \alpha_2 \not\preceq_C \vec{\tau}, \text{ or} \\
\exists \alpha_1 \text{ op } \alpha_2 <: \alpha_3 \in C \text{ such that } C \vdash \alpha_i \xrightarrow{\text{int}} \square \text{ for some } i \in \{1..2\}
\end{array}$$

A set of constraints is *consistent* if it is not inconsistent.

These conditions check for direct conflicts. In the latter condition, we detect non-integer arguments passed to an

integer binary operation. In the prior, we detect an onion of scapes α_0 for which some portion of its argument α_1 is incompatible with every scape in the onion. Note that this also detects the application of a non-scape: if α_0 contains no scapes, projection will produce an empty list which then trivially satisfies the incompatibility condition.

Typechecking is then simple. Given an expression e , we first calculate the set of constraints C such that $\llbracket e \rrbracket_E^0 = \langle \alpha, C \rangle$ (for some α). We then calculate the closure such that $C \xrightarrow{\text{Cl}_y^*} C'$. Finally, we determine if C' is inconsistent. If it is, then we reject e ; otherwise, e is declared type safe.

Because the polymorphic type system presented in the next section is a strict generalization of MicroBang⁰, we defer discussion of type soundness to the next section.

5. Polymorphism

We now extend MicroBang⁰ with parametric polymorphism to create a complete type system for MicroBang. Our polymorphism inference system is complex and so is placed in this separate section for readability. We are designing this system in the context of scripting languages, so we aim to realize the intuitions of types and subtyping held by programmers accustomed to dynamically-typed languages: a method call which can be understood using the philosophy of duck typing should simply work. For this reason, we base our system on expressive polymorphic set constraint type systems [3, 12, 14, 24, 31] adding several extensions for greater expressiveness.

Parametric polymorphism is important for expressiveness; we would like to avoid the arbitrary polymorphism cutoffs found in `let`-polymorphism and in local type inference. In order to obtain the maximal amount of polymorphism, every scape in MicroBang is inferred a polymorphic type; this approach is inspired by flow analyses [1, 28] and has previously been ported to a type constraint context [31]. In the ideal, each call site produces a fresh instantiation of the type variables. Unfortunately, this ideal cannot be implemented since a single program can have an unbounded number of function call sequences and so produce infinitely many instantiations. A standard solution to dealing with this case in the program analysis literature is to simply chop off call sequences at some fixed point. While such arbitrary cutoffs may work for program analyses, they work less well for type systems: they make the system hard for users to understand and potentially brittle to small refactorings.

We have developed an approach here starting from our previous work on this topic [16, 17, 31] to produce polymorphism on functions which is “nearly maximal” in that common programming patterns have expressive polymorphism, suffer no arbitrary cutoffs, and are not too inefficient. In this approach, a call to a non-recursive function is maximally polymorphic. Recursive call cycles are polymorphic only on the first traversal; type variables are reused when a recursive cycle is closed. Because scapes are higher-order in

MicroBang, it is not known at the start of closure which series of calls will be cyclic. Instead, the MicroBang type system optimistically instantiates fresh variables for each call sequence it encounters and then merges these variables when a recursive cycle is discovered.

This section makes adjustments to the MicroBang⁰ type system at each step – initial derivation, constraint closure, and contradiction analysis – to introduce polymorphism. We explain the mechanism we use to recognize recursive call cycles and to unify type variables when appropriate. We then present formal statements about the properties of our type system and describe our strategy for proving them correct.

5.1 Initial Derivation

To define initial type derivation for the MicroBang type system, we must first adjust our view of type variables. Type variables in MicroBang⁰, as in most type systems, are uninterpreted; they are terminal forms in the grammar. In the full MicroBang type system, however, we need to record how each variable is polyinstantiated. To do this, we use a *contour*, a structure which records a polyinstantiation context. The operations over contours are somewhat complex and quite difficult to motivate without a usage context; we therefore defer their presentation until Section 5.2. For now, we treat contours C as opaque, uninterpreted symbols.

The most direct way to associate a type variable with its contour information is to assign to it some structure. We begin by assuming that each redex and each pattern in the program is associated with some unique *identifier* ι . With this notation, our type variable grammar is as follows:

Definition 5.1 (Type Variable Grammar).

$$\begin{aligned} \dot{c} &::= c \mid \star \text{ possible contours} \\ \alpha &::= \langle \iota, \mathbb{Z}, \dot{c} \rangle \text{ type variables} \end{aligned}$$

Each type variable is now a triple between an identifier, an index, and a possible contour. The identifier names the point in the program for which the type variable was originally inferred. The index allows multiple type variables to be used when inferring a type for a given expression. Type variables also have a *possible contour* \dot{c} : either a contour C or the special symbol \star (which we describe here as “no contour”). The contour of a type variable describes how it was polyinstantiated; a type variable with no contour has not been polyinstantiated.

Initially, all type variables are inferred with no contour. In fact, the definition of initial type derivation given in Figure 4.2 is still largely suitable for the polymorphic system; we must make two small adjustments. First, rather than taking $\dot{\alpha}_i$ to represent an arbitrary fresh type variable, we give fresh type variables specific structure:

Definition 5.2 (Fresh Variables). In initial expression type derivation, $\dot{\alpha}_i = \langle \iota, i, \star \rangle$ where ι is the identifier of the redex for which we are currently deriving a type. The same is true in initial pattern derivation, where ι is the identifier of the pattern.

Second, we must capture the free type variables of functions during derivation so that they may be correctly polyinstantiated later; this parallels the process of type generalization in more traditional type systems. To accomplish this, we first modify the type grammar of the language so that scape types take the following form:

Definition 5.3 (Concrete Type Grammar).

$$\tau ::= \dots \mid \forall \alpha. \tau \phi \rightarrow \alpha \setminus C$$

The adjustments which appear in Definitions 5.1 and 5.3 are the only changes to the type grammar we must make for the polymorphic MicroBang type system. Otherwise, the type grammar we use is the same as in Figure 4.1.

Next, we define derivation over scape expressions, using $\text{FTV}(\cdot)$ as a function calculating free type variables:

Definition 5.4 (Initial Scape Type Derivation).

$$\begin{aligned} \llbracket x = \phi \rightarrow e \rrbracket_{\mathbb{R}}^{\Gamma} &= \forall \alpha. \tau \phi \rightarrow \alpha' \setminus C <: \Gamma(x) \\ \text{where } \begin{cases} \llbracket \phi \rrbracket_{\mathbb{P}}^{\Gamma} &= \langle \Gamma', \tau \phi \rangle, \\ \llbracket e \rrbracket_{\mathbb{E}}^{\Gamma \cup \Gamma'} &= \langle \alpha', C \rangle, \\ \alpha &= (\text{FTV}(\tau \phi) \cup \text{FTV}(C)) - \{ \alpha'' \mid x : \alpha'' \in \Gamma \} \end{cases} \end{aligned}$$

Initial type derivation in MicroBang is then all of the rules from Figure 4.2 amended with Definition 5.4 and interpreted using our new definition of $\dot{\alpha}_i$. The next step in defining our type system is constraint closure.

5.2 Constraint Closure

The MicroBang⁰ type system defined several relations to support its constraint closure process; those relations can be used in the MicroBang type system with little modification. We first observe that, because of our new scape type, we must modify the projection matching of Definition 4.4 to match against the new scape type syntax:

Definition 5.5 (Polymorphic Projector Match).

$$\forall \alpha. \tau \phi \rightarrow \alpha \setminus C \sqsubseteq \text{fun} \quad \text{for any } \alpha, \tau \phi, \alpha, C$$

We then can directly inherit from MicroBang⁰ the concretization (Definition 4.1), projection (Definition 4.4), and compatibility (Definition 4.6) relations.

The empty union variable mapping $\text{ETV}(\alpha)$ must be updated to observe our new type variable model; we use previously unused index position 0 for this purpose.

Definition 5.6 (Empty Onion Variables).

$$\text{ETV}(\langle \iota, n, \dot{c} \rangle) = \langle \iota, 0, \dot{c} \rangle.$$

The application-driven form of compatibility is nearly identical but must be updated for the new form of scape type described above. We define this updated relation as follows:

Definition 5.7 (Application Compatibility).

$$\begin{aligned} \alpha \preceq_C \llbracket \alpha''; C'' \rrbracket & \text{ is false} \\ \alpha \preceq_C \bar{\tau} \llbracket (\forall \alpha. \tau \phi' \rightarrow \alpha' \setminus C') \rrbracket \setminus \alpha'; C'' \cup C' & \text{ if } \alpha \preceq_C \tau \phi' \setminus C'' \\ \alpha \preceq_C \bar{\tau} \llbracket (\forall \alpha. \tau \phi' \rightarrow \alpha' \setminus C') \rrbracket \setminus \alpha''; C'' & \text{ if } \alpha \preceq_C \tau \phi' \setminus \star, \\ & \alpha \preceq_C \bar{\tau} \setminus \alpha''; C'' \end{aligned}$$

We write $\alpha \not\preceq_C \overrightarrow{\forall \alpha_{\square}. \tau \phi_{\square} \rightarrow \alpha'_{\square} \setminus C_{\square}}$ to denote that $\forall i \in \{1..n\}. \alpha \preceq_C \tau \phi_i \setminus \star$; that is, every pattern could fail to match the argument type α . Using this updated notation,

$$\begin{array}{c}
\text{APPLICATION} \\
\frac{\alpha_0 \alpha_1 <: \alpha_2 \quad C_{\text{IN}} \vdash \alpha_0 \xrightarrow{\text{fumb}} \vec{\tau} \quad \alpha_1 \preceq_{C_{\text{IN}}} \vec{\tau} \setminus \alpha_3; C' \quad c = \text{CNEW}(\alpha_2, C_{\text{IN}}) \quad C'' = \text{INST}(C' \cup \{\alpha_3 <: \alpha_2\}, c)}{C_{\text{IN}} \xrightarrow{\text{c}_1^1} \text{REPL}(C_{\text{IN}} \cup C'', c)}
\end{array}$$

Figure 5.1. MicroBang Polymorphic Application Rule

inconsistency as in Definition 4.8 is suitable for the full MicroBang type system.

We are also nearly ready to define the full MicroBang constraint closure. But while the initial derivation step only dealt with non-contoured type variables, constraint closure interacts with type variables containing any form of \dot{c} . Fortunately, constraint closure only requires three contour-related functions. We describe them at a high level for now and defer formal definitions until Section 5.3. These functions are:

- $\text{CNEW}(\alpha, C_{\text{IN}})$, a function which creates an appropriate contour for the call site described by variable α ;
- $\text{INST}(C, c)$, a function which polyinstantiates the free variables in a set of constraints C using contour c ; and
- $\text{REPL}(C, c)$, a function which uses information in c to unify type variables in recursive call cycles in C .

The Integer Addition and Integer Equality rules from Figure 4.3, when interpreted under these new definitions, are appropriate for use in the MicroBang type system; the Application rule needs to be modified; the revised rule appears in Figure 5.1. This rule begins much as it did in MicroBang^0 : it conditions on the presence of a constraint indicating an application. It then projects all of the scapes from α_0 and checks their compatibility with the argument α_1 . The rule then polyinstantiates the constraints (along with the output wiring constraint $\alpha_3 <: \alpha_2$) and unions them with the original constraint set C_{IN} . Finally, contour replacement is used to unify any cycles which were discovered by this process.

This constraint closure relation together with the initial derivation function described above provides the powerful call-site polymorphism that MicroBang requires. We formally define typechecking in Section 5.4. In order to complete the type system, however, we must first provide definitions for the three contour functions described above.

5.3 Contours

We now present the structure of contours and define the operations which closure requires of them. A contour is meant to track the polyinstantiation of a type variable. More specifically, it is meant to track the polyinstantiation contexts in which a given type variable is relevant. The type variable $\langle \iota, 1, \ast \rangle$, for instance, is the first type variable inferred at redex (or pattern) ι ; the type variable $\langle \iota, 1, c \rangle$ is a polyinstantiation of that type variable which describes the type of the redex or pattern ι only in the context represented by c .

We must therefore define contours such that they represent the contextual information relevant to our polymor-

phism model. A finite approximation of call strings is needed to avoid the unbounded computation of our operational semantics; to accomplish this, we use a restricted form of regular expressions over call strings to represent contours. We motivate our choice of restriction by example, using integers to represent identifiers and using lists of such identifiers to represent call strings (where each identifier represents a call site).

Contours by Example Using the above notation, consider the call string $[1, 2]$. This means that function call site 1 is called which then leads to a call at site 2. This call string can be represented as the simple regular expression $[1, 2]$ (we represent the concatenation of regular expressions as a list to fit with our eventual contour grammar).

When recursion happens, however, the call string will repeat a call site (e.g., $[1, 2, 3, 2]$). To guarantee termination, we restrict our regular expressions such that any identifier appears at most once. So, we will represent this call string using the contour expression $[1, \{2, 3\}]$, where a set represents the Kleene closure of its elements, i.e. “ $1(2 \mid 3)^*$ ” in traditional regular expression notation. Note that this expression also matches the call strings $[1]$, $[1, 3, 2, 3]$, *etc.* This information loss only applies to recursive call cycles; we can still extend the call string and the expression with non-recursive calls without loss. For instance, we represent the call string $[1, 2, 3, 2, 4]$ with the expression $[1, \{2, 3\}, 4]$.

In order to keep our contour expressions manageable and prevent unnecessary complexity in constraint closure, we also do not permit their nesting. For instance, the call string $[1, 2, 3, 2, 4, 1]$ would not extend the expression above to $[1, \{2, 3\}, 4, 2]$, nor would it treat the set as an element of a larger set as in $\{[1, \{2, 3\}, 4]\}$. Instead, we simply merge all of the elements between the two identical call sites, yielding $\{[1, 2, 3, 4]\}$. While this does represent a loss of information, we believe the cases in which this information loss is relevant to the programmer are obscure.

Although the expressions here are general enough to describe each polyinstantiation, constraint closure will occasionally require two contours to be *merged*, unifying the variables which exist in each contour with the corresponding variables in the other. In order to support this with no information loss, we define contours as sets of the restricted regular expressions presented here; such a contour matches a given call string if that string is matched by any of the regular expressions in its set.

Program analyses have previously use regular expressions to abstractly represent arbitrary call strings; see e.g. ΔCFA [20] and citations therein. Our work differs from existing program analyses because our goals do not match those of program optimizers. Compared to [20], for example, we create contours for each call site (rather than each function) as this finer granularity adds expressiveness and limits arbitrary cutoffs. On the other hand, we are not in-

$$\begin{aligned}
\mathcal{P} &::= \iota \mid \overleftarrow{\iota} && \text{contour parts} \\
\mathcal{S} &::= \overrightarrow{\mathcal{P}} && \text{contour strands} \\
\mathcal{C} &::= \overleftarrow{\mathcal{S}} && \text{contours}
\end{aligned}$$

Figure 5.2. MicroBang Contour Grammar

interested in the precise invariants on variables in the calling context that are derivable from a Δ CFA analysis.

Contour Grammar We now formalize the grammar of MicroBang contours using the above examples as a guide. Each term of the restricted regular expressions above was either a single, literal call site (in the form of an identifier) or a set of such call sites. We use the term *contour part* \mathcal{P} to refer to such a term. Each restricted regular expression was a list of these parts; we refer to these expressions as *contour strands* \mathcal{S} . As stated above, a contour \mathcal{C} is a set of these strands. We define this grammar in Figure 5.2.

Because contours represent regular expressions over call strings, it is helpful to formalize the meaning of contours in terms of the sets of call strings they match:

Definition 5.8 (Contour Meaning).

$$\begin{aligned}
(\text{Contours}) \quad \llbracket \mathcal{C} \rrbracket &= \bigcup \{ \llbracket \mathcal{S} \rrbracket \mid \mathcal{S} \in \mathcal{C} \} \\
(\text{Strands}) \quad \llbracket \overrightarrow{\mathcal{P}} \rrbracket &= \{ \overrightarrow{\mathcal{P}} \} \\
&\llbracket \overrightarrow{\mathcal{P}} \rrbracket = \{ \overrightarrow{\iota_1} \parallel \dots \parallel \overrightarrow{\iota_n} \mid \forall i \in \{1..n\}. \overrightarrow{\iota_i} \in \llbracket \mathcal{P}_i \rrbracket \} \\
(\text{Parts}) \quad \llbracket \overleftarrow{\iota} \rrbracket &= \{ \overleftarrow{\iota'} \mid \forall i \in \{1..n\}. \iota'_i \in \overleftarrow{\iota} \} \\
&\llbracket \iota \rrbracket = \{ \{\iota\} \}
\end{aligned}$$

Using the examples from above, we would say that $[1, 2, 3, 2, 4] \in \llbracket \mathcal{C} \rrbracket$ where \mathcal{C} is the contour containing just the strand $[1, \{2, 3\}, 4]$.

Contour Relations In order to define the three contour functions used in closure above, we require some auxiliary definitions. The first and simplest of these is contour extraction, a relation which produces the set of contours found in a type system grammatic construct such as a constraint set; we use $\langle \mathcal{C} \rangle$ to denote this. In most cases, this is just the natural catamorphism which unions results; for instance, $\langle \alpha_1 <: \alpha_2 \rangle = \langle \alpha_1 \rangle \cup \langle \alpha_2 \rangle$. The leaf cases are interesting for contours ($\langle \mathcal{C} \rangle = \{ \mathcal{C} \}$) and non-contours ($\langle \# \rangle = \emptyset$); for all other leaf cases, the extraction yields an empty set.

We next define contour subsumption: a contour \mathcal{C}_2 subsumes another contour \mathcal{C}_1 (written $\mathcal{C}_1 \leq \mathcal{C}_2$) if every call sequence described by \mathcal{C}_1 is also described by \mathcal{C}_2 . We also define contour overlap: a contour \mathcal{C}_1 overlaps a contour \mathcal{C}_2 (written $\mathcal{C}_1 \overline{\cap} \mathcal{C}_2$) if there is any call sequence represented by both of them. We define these relations below:

Definition 5.9 (Contour Relations).

$$\begin{aligned}
\mathcal{C}_1 \leq \mathcal{C}_2 &\text{ iff } \llbracket \mathcal{C}_1 \rrbracket \subseteq \llbracket \mathcal{C}_2 \rrbracket \\
\mathcal{C}_1 \overline{\cap} \mathcal{C}_2 &\text{ iff } \llbracket \mathcal{C}_1 \rrbracket \cap \llbracket \mathcal{C}_2 \rrbracket \neq \emptyset
\end{aligned}$$

Because contours are a restricted subset of regular expressions, subsumption and overlap can be computed efficiently.

Contour Creation The functions for creating new contours are described in Figure 5.3; we outline their behavior here.

$$\begin{aligned}
\text{MAKE}(\langle \iota, n, \mathcal{C} \rangle) &= \{ \mathcal{S} \parallel \llbracket \iota \rrbracket \mid \mathcal{S} \in \mathcal{C} \} \\
\text{SITES}(\iota) &= \{ \iota \} \\
\text{SITES}(\overleftarrow{\iota}) &= \overleftarrow{\iota} \\
\text{COLLAPSE}(\mathcal{C}) &= \{ \text{COLLAPSE}(\mathcal{S}) \mid \mathcal{S} \in \mathcal{C} \} \\
\text{COLLAPSE}(\mathcal{S}) &= \begin{cases} \text{COLLAPSE}(\mathcal{S}_1 \parallel \left(\bigcup \text{SITES}(\mathcal{P}_i) \right) \parallel \mathcal{S}_2) \\ \text{when } \begin{cases} \mathcal{S} = \mathcal{S}_1 \parallel \overrightarrow{\mathcal{P}} \parallel \mathcal{S}_2, n > 1, \\ \text{SITES}(\mathcal{P}_1) \cap \text{SITES}(\mathcal{P}_n) \neq \emptyset \end{cases} \\ \mathcal{S} \text{ otherwise} \end{cases} \\
\text{WIDEN}(\mathcal{C}, \mathcal{C}') &= \mathcal{C} \cup \left(\bigcup \{ \mathcal{C}' \mid \mathcal{C}' \in \langle \mathcal{C} \rangle \wedge \mathcal{C} \overline{\cap} \mathcal{C}' \} \right) \\
\text{CNEW}(\alpha, \mathcal{C}) &= \text{WIDEN}(\text{COLLAPSE}(\text{MAKE}(\alpha)), \mathcal{C})
\end{aligned}$$

Figure 5.3. MicroBang Contour Creation

The **MAKE** function gives a starting contour for a type variable representing a call site; the application rule in Figure 5.1 represents a call site by using the output type variable of the corresponding application constraint. This starting contour is the contour of the call site extended with the call site's identifier; in essence, the new contour represents pushing the call site onto a call stack.

The act of adding the call site to the starting contour may have introduced a cycle. To resolve this issue, the starting contour is then *collapsed*. **COLLAPSE** will find any cycles – that is, any sequence of contour parts which repeat a given ι , thus potentially representing call site recursion – and join them into a single Kleene closure set of ι . This enforces our invariant that the contour strands in the type system contain a given ι in at most one position.

By collapsing the contours we use in polyinstantiation, we guarantee termination; showing a worst case bound in the permutations of the call sites is trivial. Unfortunately, contour collapse does not guarantee a constraint set with disjoint contours. Closure over a constraint set in which contours overlap is quite involved and is exponentially complex in common cases. To remedy this problem, we *widen* the contour to subsume any contour it overlaps; this is as simple as unioning the strands of those contours with the strands of the new one. While the resulting contour still overlaps contours which appear in the constraint set, it can now be used to replace them. As a result, no overlapping contours exist after calculating the **REPL** of the constraint set.

Instantiation and Replacement The **INST** and **REPL** functions perform instantiation and replacement, respectively, for contours. We define here a three-argument version of contour instantiation, $\text{INST}(*, \alpha, \mathcal{C})$, where the set represents the bound variables which should not be instantiated; this is used when instantiation enters a scape type. The two-argument version of contour instantiation is defined in terms of this three-argument form: $\text{INST}(\mathcal{C}, \mathcal{C}') = \text{INST}(\mathcal{C}, \emptyset, \mathcal{C}')$.

The **INST** and **REPL** functions are quite similar in structure; they are largely the natural homomorphisms over the type system constructs. For instance, $\text{INST}(\tau <: \alpha, \alpha, \mathcal{C}) =$

$\text{INST}(\tau, \alpha, c) <: \text{INST}(\alpha, \alpha, c)$. The cases for which they are not naturally homomorphic appear below:

Definition 5.10 (Contour Function Cases).

$$\begin{aligned} \text{INST}(\forall \alpha_1. \tau \phi \rightarrow \alpha \setminus C, \alpha_2, c) &= \forall \alpha_1. \tau \phi \rightarrow \alpha \setminus \text{INST}(C, \alpha_1 \cup \alpha_2, c) \\ \text{INST}(\langle l, n, \dot{c} \rangle, \alpha, c) &= \begin{cases} \langle l, n, c \rangle & \text{if } \dot{c} = * \wedge \langle l, n, \dot{c} \rangle \notin \alpha \\ \langle l, n, \dot{c} \rangle & \text{otherwise} \end{cases} \\ \text{REPL}(\langle l, n, \dot{c} \rangle, c) &= \begin{cases} \langle l, n, c \rangle & \text{if } \dot{c} \leq c \\ \langle l, n, \dot{c} \rangle & \text{otherwise} \end{cases} \end{aligned}$$

The INST function is simply a capture-avoiding type variable substitution on the first argument. This allows the application rule to instantiate the set of constraints in a scape (which were all inferred with no contour) by giving them the contour provided by C_{NEW} .

The REPL function performs contour substitution using a given contour. The contour argument replaces contours appearing in the constraint set when they are subsumed by that contour argument; as a result, any two type variables which are both subsumed by this contour are now unified. This is the mechanism by which contours are merged to bound polyinstantiation.

The above definitions provide the three contour functions which were used in Section 5.2; this completes the definition of constraint closure.

5.4 Typechecking and Soundness

The relations above are sufficient to define a complete polymorphic type system for MicroBang. We formally define typechecking as follows:

Definition 5.11 (Typechecks). A closed expression e *typechecks* iff $\llbracket e \rrbracket_E^0 = \langle \alpha, C \rangle$, $C' = \text{INST}(C, \{\square\})$, and for any C'' , $C' \xrightarrow{c_1, *}_* C''$ implies C'' is consistent.

This definition of typechecking is similar to the description given in Section 4, but includes an additional step. Immediately after derivation, the resulting set of constraints is instantiated with the initial contour $\{\square\}$, which represents a global, top-level calling context. After this instantiation, we simply assert that no valid closure of C' contains inconsistencies.

Typechecking is both sound and decidable.

Theorem 5.1 (Type Soundness). For any closed e such that $e \xrightarrow{*} e'$, $e' \not\xrightarrow{1}$, and e' is not of the form E , we have that e does not typecheck.

A full proof of this theorem appears in Appendix A.

Theorem 5.2 (Decidability). Given a closed e , it is decidable whether e typechecks.

A full proof of this theorem appears in Appendix B.

This completes the MicroBang presentation. In the following two sections, we will extend MicroBang to include the features of TinyBang not present in MicroBang.

6. Deep and Wide Patterns

We begin extending the MicroBang type system toward TinyBang by adding deep and wide patterns. Deep patterns are those which can specify structure beyond a single label boundary; for instance, the TinyBang scape $\langle A \ B \ x \ \rightarrow \ x \rangle$ is a deep pattern. Wide patterns are those which can match their argument in multiple ways. The only form of wide pattern in TinyBang is the conjunctive pattern such as $\langle A \ x \ \& \ B \ y \rangle$. Conjunctive patterns only match an argument if all of their subpatterns match it. Here we extend MicroBang to MicroBang^ϕ, the grammar of which is the same as Figure 3.1 except for its pattern grammar:

Definition 6.1 (MicroBang^ϕ Pattern Grammar).

$$\begin{aligned} \phi &::= x : \varphi && \text{patterns} \\ \varphi &::= \text{int} \mid \text{lbl } x : \varphi \mid \varphi \ \& \ \varphi \mid \text{any } \text{primary patterns} \end{aligned}$$

This grammar supports both the wide and the deep patterns of TinyBang. Because pattern conjunction is associative, we will write $\varphi_1 \ \& \ \dots \ \& \ \varphi_n$ to denote $\varphi_1 \ \& \ (\dots (\varphi_{n-1} \ \& \ \varphi_n) \dots)$.

This section shows how the polymorphic type system for MicroBang can be extended to support both wide and deep patterns. We also provide a solution to the union alignment problem we mention in Section 4.2, as it is more pronounced when wide patterns are involved. For sake of brevity, we do not formally extend the operational semantics of MicroBang.

6.1 Derivation

In order to support the derivation step of type checking, we must extend derivation to cover our new pattern grammar. To do so, we must first extend the type grammar to provide types for the new patterns. The pattern types for MicroBang^ϕ appear below.

Definition 6.2 (MicroBang^ϕ Pattern Type Grammar).

$$\begin{aligned} \tau_\phi &::= \alpha \sim \tau_\varphi && \text{pattern types} \\ \tau_\varphi &::= \text{int} \mid \text{lbl } \alpha \sim \tau_\varphi \mid \tau_\varphi \ \& \ \tau_\varphi \mid \text{any } \text{primary pattern types} \end{aligned}$$

Overall, the type grammar for MicroBang^ϕ includes the changes from Definition 6.2 as well as the changes from the polymorphic model in Definitions 5.1 and 5.3. Otherwise, it has the same type grammar as MicroBang⁰ from Figure 4.1.

Using this type grammar, we can define initial type derivation. As with the type grammar, initial type derivation in MicroBang^ϕ is quite similar to the previous systems; it only changes for the new pattern grammar, the rules for which appear in Figure 6.1. The expression and redex rules of MicroBang^ϕ are the same as those shown in Figure 4.2 with the polymorphism adjustments from Definition 5.4.

6.2 Fibrations for Union Alignment

In order to support constraint closure, we must redefine compatibility in terms of these new patterns. However, conjunctive patterns present difficulty; the union alignment problem that we previously overlooked is considerably worse. For instance, consider $(\langle A \ x \ \& \ B \ y \ \rightarrow \ x + y \rangle \ z)$

$$\begin{aligned}
\llbracket x : \varphi \rrbracket_p^\Gamma &= \langle \{x : \dot{\alpha}_1\} \cup \Gamma', \dot{\alpha}_1 \sim \tau_\varphi \rangle \\
&\quad \text{where } \llbracket \varphi \rrbracket_p^\Gamma = \langle \Gamma', \tau_\varphi \rangle \\
\llbracket \text{int} \rrbracket_p^\Gamma &= \langle \emptyset, \text{int} \rangle \\
\llbracket \text{lbl } x : \varphi \rrbracket_p^\Gamma &= \langle \{x : \dot{\alpha}_1\} \cup \Gamma', \dot{\alpha}_1 \sim \tau_\varphi \rangle \\
&\quad \text{where } \llbracket \varphi \rrbracket_p^\Gamma = \langle \Gamma', \tau_\varphi \rangle \\
\llbracket \varphi_1 \& \varphi_2 \rrbracket_p^\Gamma &= \langle \Gamma_1 \cup \Gamma_2, \tau_{\varphi_1} \sim \tau_{\varphi_2} \rangle \\
&\quad \text{where } \llbracket \varphi_1 \rrbracket_p^\Gamma = \langle \Gamma_1, \tau_{\varphi_1} \rangle \text{ and } \llbracket \varphi_2 \rrbracket_p^\Gamma = \langle \Gamma_2, \tau_{\varphi_2} \rangle \\
\llbracket \text{any} \rrbracket_p^\Gamma &= \langle \emptyset, \text{any} \rangle
\end{aligned}$$

Figure 6.1. MicroBang^φ Initial Pattern Type Derivation

where argument z has type α and the constraint set contains $\{\text{'A int } <: \alpha, \text{'B int } <: \alpha\}$. It can be shown that both 'A and 'B can be projected from α individually, but we do not know that a single value will contain both of them.

The root of this problem is that the type variable α represents a union type and we are not ensuring that we have a consistent view of it across multiple projections. This so-called *union alignment problem* is well-known and invariably arises in union type systems; [11], for example, is a recent paper reviewing the problem and presenting a sound union elimination type rule. Unfortunately, sound union elimination rules in traditional union typing systems usually end up being weaker than desired: the rule must be syntactically constrained to eliminate only on union-typed expressions which are known to be evaluated only once (and thus to have made a *single* choice for which union branch was taken at runtime).

In a subtype constraint context, where (positive) union types take the form of multiple lower bounds on a type variable, the union eliminations can be viewed as occurring in canonical positions: each closure rule propagates per concrete lower bound, and thus can be viewed as performing an implicit union elimination by choosing a particular lower bound to propagate. For example, the α above is processed pointwise first as 'A int , and then as 'B int , fixing the choice each time. In order to achieve pointwise processing, any union decisions made for a given expression are marked and we constrain to make the *same* choices every time when considering a single function application; this achieves alignment. To record the particular union elimination decision(s) made, we introduce type *fibrations*, formally described below. We add fibrations to the projection relation, allowing us to require that multiple projections use the same fibration. We also add fibrations to the compatibility relation. The grammar of fibrations is as follows:

Definition 6.3 (Fibration Grammar).

$$f ::= * | \langle \tau, \vec{f} \rangle \text{ fibrations}$$

Each fibration f is used to describe decisions about a specific type variable at a specific point in the type. If the fibration is of the form $*$, it indicates that no decisions have been made about the type variable: it is unconstrained in context. If the fibration is of the form $\langle \tau, \vec{f} \rangle$, then a specific

$$\begin{aligned}
C \vdash \alpha_1 \xrightarrow{\text{int}} [\text{int}]; \langle \text{int}, [] \rangle &\quad \text{if } \text{int} \stackrel{C}{<:} \alpha_1 \\
C \vdash \alpha_1 \xrightarrow{\text{lbl}} [\text{lbl } \alpha_2]; \langle \text{lbl } \alpha_2, [f] \rangle &\quad \text{if } \text{lbl } \alpha_2 \stackrel{C}{<:} \alpha_1 \\
C \vdash \alpha_1 \xrightarrow{\text{fun}} [\tau]; \langle \tau, [f] \rangle &\quad \text{if } \tau \stackrel{C}{<:} \alpha_1, \\
&\quad \tau = \forall \alpha. \tau_\phi \rightarrow \alpha_2 \setminus C' \\
C \vdash \alpha_1 \xrightarrow{\pi} []; \langle \text{int}, [] \rangle &\quad \text{if } \pi \neq \text{int}, \text{int} \stackrel{C}{<:} \alpha_1 \\
C \vdash \alpha_1 \xrightarrow{\pi} []; \langle \text{lbl } \alpha_2, [f] \rangle &\quad \text{if } \pi \neq \text{lbl}, \text{lbl } \alpha_2 \stackrel{C}{<:} \alpha_1 \\
C \vdash \alpha_1 \xrightarrow{\pi} []; \langle \tau, [f] \rangle &\quad \text{if } \pi \neq \text{fun}, \tau \stackrel{C}{<:} \alpha_1, \\
&\quad \tau = \forall \alpha. \tau_\phi \rightarrow \alpha_2 \setminus C' \\
C \vdash \alpha_1 \xrightarrow{\pi} \vec{\tau}_1 \parallel \vec{\tau}_2; \langle \alpha'_1 \& \alpha'_2, [f_1, f_2] \rangle &\quad \text{if } \alpha'_1 \& \alpha'_2 \stackrel{C}{<:} \alpha_1, \\
&\quad C \vdash \alpha'_1 \xrightarrow{\pi} \vec{\tau}_1; f_1, C \vdash \alpha'_2 \xrightarrow{\pi} \vec{\tau}_2; f_2
\end{aligned}$$

Figure 6.2. MicroBang^φ Projection

lower bound has already been chosen for that type variable at that location in the type. The associated list of fibrations describe the decisions for type variables in τ ; it is a singleton list if τ is a label, for instance, and an empty list if τ is int .

Note that this deep structure is necessary: it is not sufficient to build a mapping between type variables and lower bounds. For example, consider the type α under $\{\text{int } <: \alpha, \text{'A } \alpha <: \alpha, \text{'B } \alpha <: \alpha\}$; each time the type recurses, the variable α may be given any of three lower bounds. When checking compatibility with the pattern $\text{'A 'B } _$, this type would use the fibration $\langle \text{'A } \alpha, [\langle \text{'B } \alpha, [*] \rangle] \rangle$. This fibration indicates that we chose $\text{'A } \alpha$ for the first occurrence α , we chose $\text{'B } \alpha$ for the second occurrence of α , and we did not make any further decisions. A simple mapping from α to a single lower bound is unable to express this case.

6.3 Constraint Closure

Using fibrations, we can now adjust the relations used in constraint closure to integrate deep patterns into the type system described in Section 5.

Projection We begin by redefining projection to include fibrations. We now write $C \vdash \alpha \xrightarrow{\pi} \vec{\tau}; f$ to indicate projection; this new relation simply includes an additional place for the fibration. For convenience, we elide the fibration f when it is unimportant. Given this definition, our new projection relation appears in Figure 6.2. This relation does not use the projection matching predicate from Definitions 4.3 and 5.5 because each case of projection must relate to a fibration based on the lower bounding type used.

The first two rules of this projection relation correspond to the first rule of MicroBang's projection relation from Definition 4.4; the next two rules correspond to the second rule of MicroBang's projection relation; and the last rule here corresponds to MicroBang projection's last rule. Note that the label fibrations above use the unconstrained fibration variable f ; in these cases, any fibration may relate. This is relevant for deep patterns; if a label $\text{'A } \alpha'$ is successfully projected from a type variable α , for example, we might later discover that decisions must be made for α' .

$\alpha_1 \preceq_C \alpha_2 \sim \tau_\varphi \setminus \{\alpha_1 <: \alpha_2\} \cup \dot{C}; f$	if $\alpha_1 \preceq_C \tau_\varphi \setminus \dot{C}; f$
$\alpha_1 \preceq_C \text{int} \setminus \emptyset; f$	if $C \vdash \alpha_1 \xrightarrow{\text{int}} \text{int}; f$
$\alpha_1 \preceq_C \text{lbl} \alpha_2 \sim \tau_\varphi \setminus \{\alpha_3 <: \alpha_2\} \cup \dot{C}; f'$	if $\alpha_3 \preceq_C \tau_\varphi \setminus \dot{C}; f,$ $f' = \langle \text{lbl} \alpha_3, [f] \rangle, C \vdash \alpha_1 \xrightarrow{\text{lbl}} \text{lbl} \alpha_3; f'$
$\alpha_1 \preceq_C \tau_{\varphi_1} \& \tau_{\varphi_2} \setminus \dot{C}_1 \cup \dot{C}_2; f$	if $\alpha_1 \preceq_C \tau_{\varphi_1} \setminus \dot{C}_1; f,$ $\alpha_1 \preceq_C \tau_{\varphi_2} \setminus \dot{C}_2; f$
$\alpha_1 \preceq_C \text{any} \setminus \emptyset; f$	always
$\alpha_1 \preceq_C \text{int} \setminus *; f$	if $C \vdash \alpha_1 \xrightarrow{\text{int}} *; f$
$\alpha_1 \preceq_C \text{lbl} \alpha_2 \setminus *; f$	if $C \vdash \alpha_1 \xrightarrow{\text{lbl}} *; f$

Figure 6.3. MicroBang^φ Compatibility

As with MicroBang, we often find it useful to discuss only the highest priority type which is projected from an onion. We also need to discuss the case in which no types can be projected from the onion with a given projector. We define notation for both cases below.

Definition 6.4 (MicroBang^φ Single Projection). $C \vdash \alpha \xrightarrow{\tau} \tau; f$ holds when $C \vdash \alpha \xrightarrow{\tau} \bar{\tau} \parallel [\tau]; f$ for some $\bar{\tau}$. Further, $C \vdash \alpha \xrightarrow{\tau} *; f$ holds when $C \vdash \alpha \xrightarrow{\tau} []; f$.

Compatibility We now define compatibility in MicroBang^φ. As with projection, we simply extend the previous compatibility relation, writing $\alpha \preceq_C \tau_\phi \setminus \dot{C}; f$. Compatibility is then defined in Figure 6.3.

The structure of this compatibility rule is slightly different from the one presented in Definition 4.6, but the intention is the same: each rule relates a type variable to a pattern either by producing a set of constraints (in the case that a type represented by the variable matches the pattern) or a $*$ (in the case that a type represented by the variable does not match the pattern). Because type variables may represent unions, it is possible for the same type variable and pattern to relate in multiple ways. The second rule in Figure 6.3, for instance, indicates that integers match the integer pattern under no constraints. The third rule indicates that labels match the pattern containing their label name by allowing the variable from the label type to flow to the variable in the label pattern. The last rule indicates that a label pattern does not match a variable if a label cannot be projected from that variable.

As with projection, we have redefined compatibility to use fibrations. In the case of compatibility with a conjunctive pattern (e.g., ‘**A** x & ‘**B** y -> ...), note that each compatibility relation against one of the subpatterns uses the same fibration as the other; this ensures union alignment for wide patterns. Likewise, the rule for label pattern matching enforces alignment for deep patterns; the fibration used when projecting the label type must immediately contain the fibration used when recursing into the label pattern’s subpattern.

A further advantage of adding fibrations to compatibility is that it allows us to force union alignment across multiple compatibility relations. We write the application-driven compatibility relation of MicroBang^φ as $\alpha \preceq_C \bar{\tau} \setminus \alpha'; C; f$ to indicate that α , fibrated as f , is compatible with some scape in $\bar{\tau}$. The α, α' , and C are the captured variables, return type,

and constraints of the matched scape. We define this relation as follows:

Definition 6.5 (Application Compatibility).

$\alpha \preceq_C [] \setminus \alpha''; C''; f$	is false
$\alpha \preceq_C \bar{\tau} \parallel [(\forall \alpha. \tau_\phi \rightarrow \alpha' \setminus C')] \setminus \alpha'; C'' \cup C'; f$	if $\alpha \preceq_C \tau_\phi \setminus C''; f$
$\alpha \preceq_C \bar{\tau} \parallel [(\forall \alpha. \tau_\phi \rightarrow \alpha' \setminus C')] \setminus \alpha''; C''; f$	if $\alpha \preceq_C \tau_\phi \setminus *; f,$ $\alpha \preceq_C \bar{\tau} \setminus \alpha''; C''; f$

Observe that the same fibration is used for all pattern compatibility checks; this forces application compatibility to use a consistent representation of the argument type α . For notational convenience, we omit the fibration when it is not important which fibration is used.

Finally, we take $\alpha \not\preceq_C \overrightarrow{\forall \alpha. \tau_\phi \rightarrow \alpha' \setminus C}$ to mean that $\exists f. \forall i \in \{1..n\}. \alpha \preceq_C \tau_{\phi_i} \setminus *; f$. Thus, a pattern match failure is only valid if it occurs for the same fibration at every pattern in the list of scapes.

Upgrade Complete! Using fibrations and the notational updates that appear above, the rest of the MicroBang^φ type system follows in the same fashion as the MicroBang type system. In fact, the definitions of closure, inconsistency, and typechecking are syntactically identical; we merely need to interpret Figure 5.1, Definition 4.8, and Definition 5.11 (respectively) in the context of the definitions presented in this section. We do not provide proofs of soundness and termination for this system, but there is little added complexity beyond fibration bookkeeping.

7. Bridging the Gap

While the previous sections have demonstrated how the most challenging features of TinyBang can be statically typed, there are still several language features which appear in TinyBang but not MicroBang^φ. While we do not have the space to formalize the typing of each of these features, this section shows several features of interest and discusses them in varying degrees of formality. In each case, we extend MicroBang (and not MicroBang^φ). There are no technical difficulties in extending MicroBang^φ with each of these features; we extend the simpler MicroBang type system for brevity.

7.1 State

The TinyBang language presented in Section 2 includes state: the construction of a label implies the creation of a cell, and variables bound to the contents of labels may be mutated. We now show how the operational semantics and type system for MicroBang can be amended to include these state semantics.

Syntax We begin by making changes to the MicroBang language grammar to accommodate the state semantics. Rather than treating an expression as a list of redex assignments, we treat an expression as a list of clauses, some of which are traditional redex assignments and some of which provide grammar for state operations. The resulting grammar changes appear below:

Definition 7.1 (Stateful MicroBang Grammar).

y	<i>cell variables</i>
$e ::= \vec{s}$	<i>expressions</i>
$s ::= x = r \mid y := x \mid y <- x \mid x <- y \mid y = y$	<i>clauses</i>
$v ::= x \mid \mathbb{Z} \mid \text{lbl } y \mid x \& x \mid () \mid \phi \rightarrow e$	<i>values</i>
$\varphi ::= \text{int} \mid \text{lbl } y \mid \text{any}$	<i>primary patterns</i>

The full grammar of this stateful MicroBang is the one appearing in Definition 7.1 along with the original rules for r , ϕ , x , op , lbl , and π from Figure 3.1. Although this grammar admits any form of clause at the end of an expression e , we only consider non-empty e in which the last clause is of the form $x = r$. The cell variables y in the new grammar are the only variables which actually appear in the TinyBang syntax; the x variables are only necessary for intermediate names induced by our ANF grammar. The stateful TinyBang syntax translates to this ANF grammar by using $y := x, \dots$ to create cells in place of `def $y = x$ in \dots` and by using $y <- x, \dots$ to update cells in place of `$y = x$ in \dots` . Further, all uses of a cell must be captured in an ANF variable; the TinyBang $y_1 + y_2$, for instance, is translated as $[x_1 <- y_1, x_2 <- y_2, x_3 = x_1 + x_2]$.

Semantics The new grammar clauses provide state semantics as described above. We also observe that although the value and pattern grammars are nearly identical to MicroBang, labels now contain cell variables instead of normal variables. As a result, label contents are always mutable.

Rather than interpreting E as a list of the form $\vec{x} = \vec{v}$ as we did before, we now interpret it as a list of clauses where each clause is of the form $x = v$, $y := x$, or $y = y$. Note that E is thus both the environment and the (cyclic) store. Lookup on a cell variable $E(y)$ is now defined inductively; it is $E(x)$ when $y := x \in E$ and $E(y')$ when $y = y' \in E$.

The projection function $E \downarrow_{\pi}^*(x)$ from Definition 3.1 can be used verbatim in this system, but we must redefine the label clause of compatibility in terms of the above syntax. This new definition is as follows:

Definition 7.2 (Stateful Value Compatibility).

$$x_1 \preceq_E \text{lbl } y_3 \setminus \{y_3 = y_2\} \quad \text{if} \quad E \downarrow_{\text{lbl}}(x_1) = \text{lbl } y_2$$

The definition of compatibility in stateful MicroBang is the one appearing in Definition 3.3 with the new clause from Definition 7.2.

Application Compatibility in Definition 3.4 can also be used verbatim in this system if interpreted under the new definition of compatibility. We can now define the new small step rules for stateful MicroBang:

Definition 7.3 (Stateful Small-Step Semantics).

$$\begin{aligned} E \parallel [x = x_1 == x_2] \parallel e &\longrightarrow^1 E \parallel [x' = () , y' := x', x = \text{True } y'] \parallel e \\ &\quad \text{when } E \downarrow_{\text{int}}(x_1) = n_1, E \downarrow_{\text{int}}(x_2) = n_2, n_1 = n_2, x', y' \text{ fresh} \\ E \parallel [x = x_1 == x_2] \parallel e &\longrightarrow^1 E \parallel [x' = () , y' := x', x = \text{False } y'] \parallel e \\ &\quad \text{when } E \downarrow_{\text{int}}(x_1) = n_1, E \downarrow_{\text{int}}(x_2) = n_2, n_1 \neq n_2, x', y' \text{ fresh} \\ E_1 \parallel [y := x_1] \parallel E_2 \parallel [y <- x_2] \parallel e &\longrightarrow^1 E_1 \parallel [y := x_2] \parallel E_2 \parallel e \\ E_1 \parallel [y := x_1] \parallel E_2 \parallel [x_2 <- y] \parallel e &\longrightarrow^1 \\ &\quad E_1 \parallel [y := x_1] \parallel E_2 \parallel [x_2 = x_1] \parallel e \end{aligned}$$

$$\begin{aligned} \llbracket s \rrbracket_E^\Gamma &= \langle \dot{\alpha}_1, \{ \llbracket s \rrbracket_R^{\Gamma \cup \Gamma'} \} \rangle \\ &\quad \text{where } \Gamma' = \{ \text{VAR}(s) : \dot{\alpha}_1 \} \\ \llbracket [s] \parallel e \rrbracket_E^\Gamma &= \langle \alpha', \{ \llbracket s \rrbracket_R^{\Gamma \cup \Gamma'} \} \cup C \rangle \\ &\quad \text{where } \Gamma' = \{ \text{VAR}(s) : \dot{\alpha}_1 \}, \llbracket e \rrbracket_E^{\Gamma \cup \Gamma'} = \langle \alpha', C \rangle \end{aligned}$$

$$\begin{aligned} \llbracket x = \text{lbl } y \rrbracket_R^\Gamma &= \text{lbl } \Gamma(y) <: \Gamma(x) \\ \llbracket y := x \rrbracket_R^\Gamma &= \text{cell}(\Gamma(x)) <: \Gamma(y) \\ \llbracket y <- x \rrbracket_R^\Gamma &= \Gamma(y) <: \text{cellS}(\Gamma(x)) \\ \llbracket x <- y \rrbracket_R^\Gamma &= \Gamma(y) <: \text{cellG}(\Gamma(x)) \\ \llbracket y = y' \rrbracket_R^\Gamma &= \text{cell}(\Gamma(y)) <: \text{cell}(\Gamma(y')) \end{aligned}$$

$$\llbracket x : \text{lbl } y \rrbracket_P^\Gamma = \langle \{ x : \dot{\alpha}_1, y : \dot{\alpha}_2 \}, \dot{\alpha}_1 \sim \text{lbl } \dot{\alpha}_2 \rangle$$

Figure 7.1. Stateful Initial Derivation Rules

The small step semantics for stateful MicroBang are then defined by the application and addition rules from Definition 3.5 and all of the rules in Definition 7.3.

Initial Derivation The semantics above create a model of state in which cells cannot directly contain other cells; labels must be used to nest them. To model this directly in the type system, we treat cells as a new form of constraint rather than a type constructor. The cell lower-bounding constraint carries a type variable representing the type of its contents; the upper-bounding constraints carry the variables which are used to read from or write to the cell. This allows us to incrementally build the type of the cell by adding constraints to the type carried by the cell lower bound, which in turn supports our model of flow-insensitive cell types. We specify the additions to the constraint grammar below:

Definition 7.4 (State Constraint Grammar).

$$\begin{aligned} c ::= \dots \mid \text{cell}(\alpha) <: \alpha \mid \text{cell}(\alpha) <: \text{cell}(\alpha) \\ \mid \alpha <: \text{cellG}(\alpha) \mid \alpha <: \text{cellS}(\alpha) \quad \text{constraints} \end{aligned}$$

The lower-bounding cell constraint is used to represent the creation of a cell; the two upper-bounding constraints represent, respectively, getting a value from and setting a value into a cell. We make use of these constraints by modifying the initial type derivation rules. The modified derivation rules (as well as the rule for assignment) appear in Figure 7.1. For these rules, we define $\text{VAR}(s)$ to be the variable on the left side of a given s , which is either an x or y ; we also allow bindings for both x and y to appear in Γ . Finally, we extend the function $\llbracket \cdot \rrbracket_R^\Gamma$ to operate on all clauses and we assume that all clauses are associated with a unique identifier ι (not just redices).

The initial type derivation for stateful MicroBang is defined by Figure 4.2 incorporating the changes from Definition 5.4 and Figure 7.1.

Constraint Closure Our next step in defining state is to define the operations in constraint closure. The projection, compatibility, and application compatibility definitions in Section 5 can be directly inherited for use in this type system.

$\frac{\alpha_1 == \alpha_2 <: \alpha_3 \quad \alpha'_3 = \text{ETV}(\alpha_3) \quad \alpha''_3 = \text{CTV}(\alpha_3)}{\text{C}_{\text{IN}} \vdash \alpha_1 \xrightarrow{\text{int}} \text{int} \quad \text{C}_{\text{IN}} \vdash \alpha_2 \xrightarrow{\text{int}} \text{int}}$
$\frac{\text{FORWARD PROPAGATION}}{\frac{\text{cell}(\alpha_1) <: \alpha_2 \quad \alpha_2 <: \text{cellG}(\alpha_3)}{\alpha_1 <: \alpha_3}}$
$\frac{\text{BACKWARD PROPAGATION}}{\frac{\text{cell}(\alpha_1) <: \alpha_2 \quad \alpha_2 <: \text{cellS}(\alpha_3)}{\text{cell}(\alpha_3) <: \alpha_2}}$
$\frac{\text{ALIAS PROPAGATION}}{\frac{\text{cell}(\alpha_1) <: \text{cell}(\alpha_2) \quad \text{cell}(\alpha_3) <: \alpha_4 \quad \alpha_3 \in \{\alpha_1, \alpha_2\}}{\{\text{cell}(\alpha_1) <: \alpha_4, \text{cell}(\alpha_2) <: \alpha_4\}}}$

Figure 7.2. New Stateful Closure Rules

To define the constraint closure rules themselves, we first must define $\text{CTV}(\cdot)$, a type variable function similar to $\text{ETV}(\cdot)$, to model the new fresh variable for cells used in the operational semantics. This function uses -1 , another as-of-yet unused index:

Definition 7.5 (Cell Variables). $\text{CTV}(\langle \iota, n, \dot{c} \rangle) = \langle \iota, -1, \dot{c} \rangle$.

Using the above, we define our new constraint closure rules in Figure 7.2. These rules assume that concretization is extended to cell lower bounds. The Integer Equality rule is redefined to ensure that it properly creates cells for labels. We also add two rules for cell type propagation. The remaining closure rules can be inherited from Figures 4.3 and 5.1.

The propagation rules in this figure propagate types through cells in a flow-insensitive manner: the static type of a cell is the union of all of the types which may be assigned to it. Forward propagation ensures that accessing a cell variable produces the type of its contents. Backward propagation ensures that any assignment to a cell causes it to be treated as if it may have had that type all along. Note that this model of state solves the polymorphic reference problem without a value restriction or equivalent. Inconsistency and typechecking definitions are directly inherited from Definitions 4.8 and 5.11, respectively.

7.2 Onion Filtering

The TinyBang language includes onion subtraction ($\&-$) and onion projection ($\&\cdot$), two features which functionally filter an onion structure. This filtering allows onions to be taken apart and recombined. We now introduce these operations as type constructors in the MicroBang language. We will use the type grammar from Figure 4.1 with changes from Definitions 5.1 and 5.3 as well as from Definition 7.6 below:

Definition 7.6 (Onion Filtering Types).

$$\tau ::= \dots \mid \alpha \&- \pi \mid \alpha \&\cdot \pi \text{ types}$$

Filtering is shallow, so it is described using a projector. The expressions $e \&\cdot \pi$ and $e \&- \pi$ produce values that

contain only those components which match (resp. do not match) the projector. Initial derivation over these expressions is as expected:

Definition 7.7 (Onion Filtering Initial Derivation).

$$\begin{aligned} \llbracket x = x' \&\cdot \pi \rrbracket_{\text{R}}^{\Gamma} &= \Gamma(x') \&\cdot \pi <: \Gamma(x) \\ \llbracket x = x' \&- \pi \rrbracket_{\text{R}}^{\Gamma} &= \Gamma(x') \&- \pi <: \Gamma(x) \end{aligned}$$

In the type system, we view $\&\cdot$ and $\&-$ as onion constructors similar to $\alpha_1 \&\alpha_2$; as a result, they should not be included where a non-onion type is expected to appear. To ensure this, we add the following rules to our definition of projection from Definition 4.4.

Definition 7.8 (Onion Filtering Projection).

$$\begin{aligned} \text{C} \vdash \alpha_1 \xrightarrow{\pi} \vec{\tau} \quad \text{if} \quad \alpha_2 \&- \pi' \text{C} <: \alpha_1, \pi \neq \pi', \text{C} \vdash \alpha_2 \xrightarrow{\pi} \vec{\tau} \\ \text{C} \vdash \alpha_1 \xrightarrow{\pi} \square \quad \text{if} \quad \alpha_2 \&- \pi' \text{C} <: \alpha_1, \pi = \pi' \\ \text{C} \vdash \alpha_1 \xrightarrow{\pi} \vec{\tau} \quad \text{if} \quad \alpha_2 \&\cdot \pi' \text{C} <: \alpha_1, \pi = \pi', \text{C} \vdash \alpha_2 \xrightarrow{\pi} \vec{\tau} \\ \text{C} \vdash \alpha_1 \xrightarrow{\pi} \square \quad \text{if} \quad \alpha_2 \&\cdot \pi' \text{C} <: \alpha_1, \pi \neq \pi' \end{aligned}$$

The first two rules deal with onion subtraction. The first rule requires that the projectors are not equal; this means that, in this case, the subtraction should not affect the result of projection. In such a case, projection continues into the type variable of the onion subtraction type. The second rule requires that the projector being used is equal to the projector being subtracted; in this case, we know that any type projected from the variable in the onion subtraction type would be subtracted here and so we simply return an empty list. The third and fourth rule are of the same form but with the projector conditions reversed.

In this way, we use projection to hide the structure introduced by the filtering types. Every other step in the system uses projection rather than destructing types directly, so we can interpret all of the remaining definitions in Section 5 in terms of this new projection relation in order to provide the onion filtering features.

7.3 Symmetric Onions

While we have shown that the asymmetric properties of onions concisely describe overriding behavior, it is sometimes useful to guarantee an override has not occurred. In Section 2.3, we highlight TinyBang's symmetric onioning operator: $\&!$. This operator creates an onion symmetrically, using the constraint-based nature of our type system to raise a type error if an override has occurred.

This symmetry requirement is represented by a new form of constraint $\alpha_1 \not\& \alpha_2$ – which indicates that two variables do not share a common component. This constraint is introduced whenever a symmetric onion is used; that is, $\llbracket x = x_1 \&! x_2 \rrbracket_{\text{R}}^{\Gamma} = \{\Gamma(x_1) \not\& \Gamma(x_2), \Gamma(x_1) \& \Gamma(x_2) <: \Gamma(x)\}$. We then complete type derivation and perform constraint closure as before. Implementing the symmetry guarantee is as simple as ensuring that no single projector succeeds in projecting a type on both sides of the symmetric onion. To accomplish this, we use the following definition of constraint set inconsistency:

Definition 7.9 (Symmetric Onion Inconsistency). C is inconsistent iff it is inconsistent by Definition 4.8 or the following is true:

$$\exists \alpha_1 \not\sim \alpha_2 \in C \text{ such that } \exists \pi. C \vdash \alpha_1 \xrightarrow{\pi} \tau_1 \wedge C \vdash \alpha_2 \xrightarrow{\pi} \tau_2$$

8. Conclusions

We have presented TinyBang, a core language in which a broad class of flexible object-oriented operations may be easily encoded and statically typed. Encodable operations include first-class messages, functional object extension, mixins, default arguments, and overloading. In particular, TinyBang’s object extension is more general than in previous works; an object can be extended even after it has been messaged without loss of depth subtyping or other critical properties.

No single feature of the TinyBang language is responsible for this flexibility. Instead, TinyBang’s novelty lies in the *combination* of improvements to the state of the art in type system design: a more general notion of first-class cases, simpler typing for asymmetric record concatenation, a more expressive dependent typing of case/match, concise syntax via type-indexed records, and a highly flexible method for inference of parametric polymorphism.

We defined the operational semantics, a type inference algorithm, and established type soundness for MicroBang, a subset of TinyBang with the most key features. To make a useful type theory, we needed to address how to maximally polyinstantiate functions yet preserve termination; for this, we defined a notion of polyinstantiation *contour* denoted by a form of regular expression. For wide and deep patterns, we had to introduce *fibrations* to solve a subtle question of union type alignment. To prove decidability of type inference, the issue of non-contractive recursive types had to be addressed.

The TinyBang type inference algorithm and a TinyBang interpreter have been implemented in Haskell, allowing us to confirm correctness of the type inference algorithm and operational semantics on code examples. All examples in this paper typecheck and run in the current implementation. The type inference implementation still needs significant tuning for improved efficiency.

The larger picture In practice, programmers need more than just the expressiveness of TinyBang; they also need syntactic sugar, an efficient compiler, and a rich set of debugging tools. We are presently working on a realistic language design and implementation for BigBang [19]; TinyBang is the core for BigBang. BigBang should be appealing because it will preserve most of the versatile spirit of dynamically-typed scripting languages while reaping the scalability and runtime performance made possible by static types.

Acknowledgments

We thank the ECOOP 2013 reviewers; our effort to improve the technical presentation was inspired by their thorough reviews. We also thank Nathaniel (Wes) Filardo for helpful comments.

References

- [1] O. Agenes. The cartesian product algorithm. In *Proceedings ECOOP’95*, volume 952 of *Lecture Notes in Computer Science*, 1995.
- [2] G. Agha. *Actors: a model of concurrent computation in distributed systems*. MIT Press, 1985.
- [3] A. Aiken and E. L. Wimmers. Type inclusion constraints and type inference. In *FPCA*, pages 31–41, 1993.
- [4] A. Aiken, E. L. Wimmers, and T. K. Lakshman. Soft typing with conditional types. In *POPL 21*, pages 163–173, 1994.
- [5] E. Allen, D. Chase, J. Hallet, V. Luchangco, J.-W. Maessen, S. Ryu, G. L. S. Jr., and S. Tobin-Hochstadt. The Fortress language specification, version 1.0. Technical report, Oracle Labs, 2008.
- [6] L. Bettini, V. Bono, and B. Venneri. Delegation by object composition. *Science of Computer Programming*, 76:992–1014, 2011.
- [7] M. Blume, U. A. Acar, and W. Chae. Extensible programming with first-class cases. In *ICFP ’06*, pages 239–250, 2006.
- [8] V. Bono and K. Fisher. An imperative, first-order calculus with object extension. In *ECOOP’98*, pages 462–497. Springer Verlag, 1998.
- [9] K. B. Bruce, L. Cardelli, and B. C. Pierce. Comparing object encodings. *Information and Computation*, 155(1-2):108–133, 1999.
- [10] G. Castagna, G. Ghelli, and G. Longo. A calculus for overloaded functions with subtyping. *Information and Computation*, 117(1):115–135, 1995.
- [11] J. Dunfield. Elaborating intersection and union types. In *Int’l Conf. Functional Programming*, 2012.
- [12] J. Eifrig, S. Smith, and V. Trifonov. Type inference for recursively constrained types and its application to OOP. In *MFPS*, Electronic Notes in Theoretical Computer Science. Elsevier, 1995.
- [13] A. Guha, J. G. Politz, and S. Krishnamurthi. Fluid object types. Technical Report CS-11-04, Brown University, 2011.
- [14] N. Heintze. Set-based analysis of ML programs. In *LFP*, pages 306–317. ACM, 1994.
- [15] B. Jay and D. Kesner. First-class patterns. *J. Funct. Program.*, 19(2):191–225, 2009.
- [16] A. Kulkarni, Y. D. Liu, and S. F. Smith. Task types for pervasive atomicity. In *OOPSLA*, pages 671–690. ACM, 2010.
- [17] Y. D. Liu and S. Smith. Pedigree types. In *International Workshop on Aliasing, Confinement and Ownership in object-oriented programming (IWACO)*, 2008.
- [18] D. B. MacQueen, G. Plotkin, and R. Sethi. An ideal model for recursive polymorphic types. *Information and Control*, 71:95–130, 1986.
- [19] P. H. Menon, Z. Palmer, A. Rozenshteyn, and S. Smith. Big Bang: Designing a statically-typed scripting language. In *International Workshop on Scripts to Programs (STOP)*, Beijing, China, 2012.
- [20] M. Might and O. Shivers. Environment analysis via Δ CFA. In *Proceedings of the 33rd Annual ACM Symposium on the Principles of Programming Languages (POPL 2006)*, pages 127–140, Charleston, South Carolina, January 2006.
- [21] T. D. Millstein and C. Chambers. Modular statically typed multimethods. In *ECOOP ’99*, pages 279–303. Springer-Verlag, 1999.
- [22] W. B. Mugridge, J. Hamer, and J. G. Hosking. Multi-methods in a statically-typed programming language. In *ECOOP ’91*, 1991.
- [23] S. Nishimura. Static typing for dynamic messages. In *POPL’98*, 1998.
- [24] F. Pottier. A versatile constraint-based type inference system. *Nordic J. of Computing*, 7(4):312–347, 2000.
- [25] D. Rémy and J. Vouillon. Objective ML: a simple object-oriented extension of ML. In *POPL ’97*, pages 40–53. ACM, 1997.
- [26] J. G. Riecke and C. A. Stone. Privacy via subsumption. *Inf. Comput.*, 172(1):2–28, Feb. 2002.

- [27] M. Shields and E. Meijer. Type-indexed rows. In *POPL*, pages 261–275, 2001.
- [28] O. Shivers. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie-Mellon University, 1991. TR CMU-CS-91-145.
- [29] M. Wand. Complete type inference for simple objects. In *Proc. 2nd IEEE Symposium on Logic in Computer Science*, 1987.
- [30] M. Wand. Corrigendum: Complete type inference for simple objects. In *LICS*. IEEE Computer Society, 1988.
- [31] T. Wang and S. F. Smith. Precise constraint-based type inference for Java. In *ECOOP'01*, pages 99–117, 2001. ISBN 3-540-42206-4.

A. MicroBang Type Soundness Proof

We present here a proof for Theorem 5.1, the type soundness theorem for MicroBang. Our proof strategy is to build a simulation relation between the small step semantics and the type system. We show that the initial derivation of a program simulates that program and that the simulation is preserved as the program executes. Finally, we use this simulation to show that every stuck program is simulated by an inconsistent constraint set.

Throughout this proof, the following notational sugar is used. This notation does not add meaning to the operational semantics or type system; it merely simplifies the presentation of some of the definitions below.

Definition A.1 (Proof Notation).

- We will write \mathbf{x} in lieu of \overline{x} to represent a set of variables.
- We will write \mathcal{C} in lieu of \overline{c} to represent a set of contours.

A.1 Variables

In order to define simulation, we must establish a convention for naming variables similar to the one given in Definition 5.1 for type variables. We are only concerned with programs for which small-step semantics are defined; we thus know that the left-hand side of all variable definitions in the program will be unique. We therefore define a bijection of the form $x \leftrightarrow \langle \iota, \mathbb{Z}, \dot{c} \rangle$; that is, this bijection maps each variable to a triple between an identifier, an integer, and a possible contour. This bijection is implicitly parametric in terms of the expression e that we are typechecking.

Definition A.2 (Variable Bijection). Let $\mathbf{x} = \{x'_i \mid e = \overrightarrow{x'_i = r'_i}, 1 \leq i \leq m\}$. For each redex clause $x = r$ identified by ι in the program we are typechecking, let $x \leftrightarrow \langle \iota, 1, \dot{c} \rangle$ where $\dot{c} = \{\}\}$ when $x \in \mathbf{x}$ and $\dot{c} = \bullet$ otherwise. In a pattern $x : \varphi$ identified by ι , we let $x \leftrightarrow \langle \iota, 1, \bullet \rangle$. In a pattern $x : \text{lbl } x'$ identified by ι , we let $x' \leftrightarrow \langle \iota, 2, \bullet \rangle$.

We leave the rest of this bijection unconstrained. We choose this bijection because it allows us to align fresh variables with the type variables that represent them. Throughout this proof, we will occasionally use a variable's triple in place of the variable itself; for instance, we may write $\{\langle \iota, 1, \bullet \rangle : \langle \iota, 1, \bullet \rangle\}$ to mean $\{x : \langle \iota, 1, \bullet \rangle\}$ such that $x \leftrightarrow \langle \iota, 1, \bullet \rangle$.

We will also need to be more precise in our definition of the α -conversion function discussed in Section 3.2. We begin by defining a function $\text{BV}(e)$ which produces the set of variables bound within a given expression e ; this function is only defined over expressions with unique variable names. We also define this function over patterns and scape values for convenience.

Definition A.3 (Bound Variables). We define $\text{BV}(\overrightarrow{x'_i = r'_i}) = \overrightarrow{x'_i}$. We let $\text{BV}(x : \text{lbl } x') = \{x, x'\}$ and we let $\text{BV}(x : \varphi) = \{x\}$ in all other cases of φ . We let $\text{BV}(\phi \rightarrow e) = \text{BV}(\phi) \cup \text{BV}(e)$.

This bound variables function produces a set which mirrors the purpose of the set of type variables captured by a

function type, but it is easier to compute given the structure of expressions in the operational semantics.

Next we define a function $\text{EINST}(*, \mathbf{x}, c)$ which freshens the free variables in the language construct $*$ except those appearing in \mathbf{x} . This function is the analogue of $\text{INST}(*, \alpha, c)$ in Definition 5.10 and, as such, is largely the natural homomorphism over the language constructs. The only case in which this function deviates is shown below:

Definition A.4 (Evaluation Instantiation).

$$\begin{aligned} \text{EINST}(\phi \rightarrow e, \mathbf{x}, c) &= \phi \rightarrow (\text{EINST}(e, \mathbf{x} \cup \text{BV}(\phi \rightarrow e), c)) \\ \text{EINST}(\langle \iota, n, \dot{c} \rangle, \mathbf{x}, c) &= \begin{cases} \langle \iota, n, c \rangle & \text{if } \dot{c} = \bullet \wedge \langle \iota, n, \dot{c} \rangle \notin \mathbf{x} \\ \langle \iota, n, \dot{c} \rangle & \text{otherwise} \end{cases} \end{aligned}$$

Given the above, we define an α -conversion function $\text{EFRESH}(x, e)$ which freshens variables in the given expression in the context of call site x . We define this function as:

Definition A.5 (Evaluation Freshening).

$$\begin{aligned} \text{EFRESH}(\langle \iota, n, \dot{c} \rangle, e) &= \text{EINST}(e, \emptyset, c') \\ \text{EFRESH}(\langle \iota, n, \dot{c} \rangle, x) &= \text{EINST}(x, \emptyset, c') \\ \text{where } c' &= \{S \parallel [\iota] \mid S \in c\} \end{aligned}$$

In this function, c' is analogous to the result of MAKE in Figure 5.3. We overload EFRESH to be defined over lists of scapes, which is calculated by applying EFRESH pointwise on each scape.

Using the freshening function as defined above, we can now clarify its use in the small step semantics in Definition 3.5. We do so in the following definition:

Definition A.6 (Clarified Application Semantics).

$$\begin{aligned} E \parallel [x = x_1 \ x_2] \parallel e'' \longrightarrow^1 E \parallel \text{EFRESH}(x, E_0 \parallel e' \parallel [x' = r]) \\ \parallel [x = \text{EFRESH}(x, x')] \parallel e'' \\ \text{when } E \downarrow_{\text{fun}}^*(x_1) = \overrightarrow{\phi \rightarrow e}, x_2 \preceq_E \overrightarrow{\phi \rightarrow e} \setminus E_0; e' \parallel [x' = r] \end{aligned}$$

Observe that Definition A.6 merely replaces the use of $\alpha(\cdot)$ with $\text{EFRESH}(x, \cdot)$. For this proof, we take the small-step semantics of MicroBang to be those in Definition 3.5 except the application case, which is taken from Definition A.6. In light of this precise definition of freshening, it is now possible to define the simulation relation for the soundness proof.

A.2 Simulation

We now define the simulation relation as a two place relation defined between a language construct and its corresponding type system construct. Values, for example, are simulated by types; redex assignments are simulated by constraints. We write $*_1 \preceq *_2$ to denote that language construct $*_1$ is simulated by type system construct $*_2$. We define this relation as follows:

Definition A.7 (Simulation). Simulation is defined by the rules in Figure A.1.

A.3 Contours

We begin constructing the type soundness proof by showing some properties of contours.

Lemma A.8. Given any contour $c \neq \emptyset$, $\llbracket c \rrbracket \neq \emptyset$.

variables:			
\ast	\simeq	\ast	
c	\simeq	c'	iff $c \leq c'$
$\langle \iota, n, \dot{c} \rangle$	\simeq	$\langle \iota, n, \dot{c}' \rangle$	iff $\dot{c} \simeq \dot{c}'$
\mathbf{x}	\simeq	α	iff $\forall x' \in \mathbf{x}. \exists \alpha \in \alpha. x' \simeq \alpha$
patterns:			
$x : \varphi$	\simeq	$\alpha \sim \tau_\varphi$	iff $x \simeq \alpha \wedge \varphi \simeq \tau_\varphi$
int	\simeq	int	
lbl x	\simeq	lbl α	iff $x \simeq \alpha$
any	\simeq	any	
values:			
x	\simeq	α	iff $x \simeq \alpha$
\mathbf{x}	\simeq	α	iff $\forall x \in \mathbf{x}. \exists \alpha \in \alpha. x \simeq \alpha$
n	\simeq	int	iff $n \in \mathbb{Z}$
lbl x	\simeq	lbl α	iff $x \simeq \alpha$
$x_1 \& x_2$	\simeq	$\alpha_1 \& \alpha_2$	iff $x_1 \simeq \alpha_1 \wedge x_2 \simeq \alpha_2$
$()$	\simeq	$()$	
$\phi \rightarrow e$	\simeq	$\forall \alpha. \tau_\phi \rightarrow \alpha \setminus C$	iff $\text{BV}(\phi \rightarrow e) \simeq \alpha \wedge \phi \simeq \tau_\phi \wedge e \simeq \langle \alpha, C \rangle$
redexes:			
$x = v$	\simeq	$\tau <: \alpha$	iff $x \simeq \alpha \wedge v \simeq \tau$
$x_3 = x_1 x_2$	\simeq	$\alpha_1 \alpha_2 <: \alpha_3$	iff $\forall i \in \{1..3\}. x_i \simeq \alpha_i$
$x_3 = x_1 \text{ op } x_2$	\simeq	$\alpha_1 \text{ op } \alpha_2 <: \alpha_3$	iff $\forall i \in \{1..3\}. x_i \simeq \alpha_i$
expressions:			
\square	\simeq	C	
$e \parallel [x = r]$	\simeq	$C \cup \{c\}$	iff $e \simeq C \wedge x = r \simeq c$
$e \parallel [x = r]$	\simeq	$\langle \alpha, C \rangle$	iff $e \parallel [x = r] \simeq C \wedge x \simeq \alpha$
environments:			
\square	\simeq	\emptyset	
$E \parallel [x = v]$	\simeq	$\Gamma \cup \{x : \tau\}$	iff $E \simeq \Gamma \wedge v \simeq \tau$

Figure A.1. Simulation Relation Definition

Proof. Because $c \neq \emptyset$, it is sufficient to show that $\llbracket S \rrbracket \neq \emptyset$ for any S in c . When S is the empty list, the result is $\{\square\}$ and so non-empty. Otherwise, it suffices to show that $\llbracket \mathcal{P} \rrbracket \neq \emptyset$ for any \mathcal{P} , which is immediate by inspection. \square

Lemma A.9. Suppose $c_1 \leq c_2$ and $c_1 \leq c_3$ and $c_1 \neq \emptyset$. Then $c_2 \not\leq c_3$.

Proof. By the definition of subsumption, we have $\llbracket c_1 \rrbracket \subseteq \llbracket c_2 \rrbracket$ and $\llbracket c_1 \rrbracket \subseteq \llbracket c_3 \rrbracket$. Thus, $\llbracket c_2 \rrbracket \cap \llbracket c_3 \rrbracket \supseteq \llbracket c_1 \rrbracket$. By Lemma A.8, we have that $\llbracket c_1 \rrbracket \neq \emptyset$ and so $c_2 \not\leq c_3$ by Definition 5.9. \square

Lemma A.10. Contour subsumption is transitive.

Proof. Trivial by definition of contour subsumption: \subseteq is transitive. \square

Lemma A.11. $\llbracket C_1 \cup C_2 \rrbracket = \llbracket C_1 \rrbracket \cup \llbracket C_2 \rrbracket$

Proof. Trivial by inspection of Definition 5.10. \square

Lemma A.12. If τ appears in C , then $\llbracket \tau \rrbracket \subseteq \llbracket C \rrbracket$. Similar statements are true about other constructs appearing in C .

Proof. Trivial by inspection of Definition 5.10. \square

Lemma A.13. $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$ iff $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket \subseteq \llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$; likewise, $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$ iff $\llbracket \llbracket \mathcal{P} \rrbracket \parallel S_1 \rrbracket \subseteq \llbracket \llbracket \mathcal{P} \rrbracket \parallel S_2 \rrbracket$.

Proof. For all of the below, let $S_1 = \overrightarrow{n} \mathcal{P}$ and let $S_2 = \overrightarrow{n} \mathcal{P}'$.

First, assume $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$. We intend to show that $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket \subseteq \llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$. By Definition 5.10, every element in $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$ is of the form $\bar{\iota}_1 \parallel \bar{\iota}_3$ where $\bar{\iota}_1 \in \llbracket S_1 \rrbracket$ and $\bar{\iota}_3 \in \llbracket \mathcal{P} \rrbracket$; we also know that every element of $\llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$ is of the form $\bar{\iota}_2 \parallel \bar{\iota}_3$ where $\bar{\iota}_2 \in \llbracket S_2 \rrbracket$. Because $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$, we have that any $\bar{\iota}_1 \parallel \bar{\iota}_3 = \bar{\iota}_2 \parallel \bar{\iota}_3$ for some $\bar{\iota}_2$; thus, $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket \subseteq \llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$.

We next assume $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket \subseteq \llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$ with the intent of proving that $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$. By Definition 5.10, every element in $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$ is of the form $\bar{\iota}_1 \parallel \bar{\iota}_3$ where $\bar{\iota}_1 \in \llbracket S_1 \rrbracket$ and $\bar{\iota}_3 \in \llbracket \mathcal{P} \rrbracket$; we also know that every element of $\llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$ is of the form $\bar{\iota}_2 \parallel \bar{\iota}_3$ where $\bar{\iota}_2 \in \llbracket S_2 \rrbracket$. For the first set, we can create a relation defining each element as equivalent if their values of $\bar{\iota}_1$ are equivalent; we can do likewise for the second set and $\bar{\iota}_2$. Each equivalence class is of the same size and contains a number of elements equal to the possible forms of $\bar{\iota}_3$, which is the cardinality of $\llbracket \mathcal{P} \rrbracket$. Since $\llbracket S_1 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket \subseteq \llbracket S_2 \rrbracket \parallel \llbracket \mathcal{P} \rrbracket$, we know that there are at least as many elements in the second set as there are in the first; since all equivalence classes are the same size, we know that there are at least as many equivalence classes in the second set as there are in the first. Because equivalence classes from these sets can be indexed by $\bar{\iota}_1$ and $\bar{\iota}_2$ respectively, we know that the set of all $\bar{\iota}_1$ is a subset of the set of all $\bar{\iota}_2$. By definition, $\bar{\iota}_1$ is an arbitrary member of $\llbracket S_1 \rrbracket$ and $\bar{\iota}_2$ is an arbitrary member of $\llbracket S_2 \rrbracket$; thus, $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$.

To prove that $\llbracket S_1 \rrbracket \subseteq \llbracket S_2 \rrbracket$ iff $\llbracket \llbracket \mathcal{P} \rrbracket \parallel S_1 \rrbracket \subseteq \llbracket \llbracket \mathcal{P} \rrbracket \parallel S_2 \rrbracket$, we use the same approach as above. \square

Lemma A.14. $\llbracket c_1 \rrbracket \subseteq \llbracket c_2 \rrbracket$ iff $\llbracket \{S \parallel \mathcal{P} \mid S \in c_1\} \rrbracket \subseteq \llbracket \{S \parallel \mathcal{P} \mid S \in c_2\} \rrbracket$; likewise, $\llbracket c_1 \rrbracket \subseteq \llbracket c_2 \rrbracket$ iff $\llbracket \{\llbracket \mathcal{P} \rrbracket \parallel S \mid S \in c_1\} \rrbracket \subseteq \llbracket \{\llbracket \mathcal{P} \rrbracket \parallel S \mid S \in c_2\} \rrbracket$.

Proof. In a fashion similar to Lemma A.14. In particular, the semantics of contours simply expands the number of forms of $\bar{\iota}_1$ and $\bar{\iota}_2$; the rest of the proof structure is identical. \square

A.4 Constraint Set Well-Formedness

The constraint closure process relies on the assumption that every contour in the constraint set is disjoint; that is, no two contours found in the constraint set overlap. This is necessary to ensure that closure is efficient; otherwise, two premises which share a term might apply at some intersection of their terms' contours. Furthermore, it is necessary to ensure that every constraint is meaningful; therefore, no constraint may contain the empty contour (\emptyset). We describe a contour set as *well-formed* when all of its contours are disjoint and none are the empty contour. A constraint set is well-formed when the contours which appear within it constitute a well-formed set. We formalize this as Definition A.15.

Definition A.15 (Well-Formed Contour Set). A contour set \mathcal{C} is well-formed when $\forall c_1 \in \mathcal{C}. c_1 \neq \emptyset \wedge \forall c_2 \in \mathcal{C}. c_1 \neq c_2 \implies c_1 \not\bowtie c_2$. A constraint set C is well-formed when $\langle C \rangle$ is well-formed.

We now show some important properties of well-formed constraint sets.

Lemma A.16. Let $C' = C_1 \cup C_2$ where C_1 and C_2 are well-formed. If $\langle C_1 \rangle \subseteq \langle C_2 \rangle$, then C' is also well-formed.

Proof. By Lemma A.11 and because $\langle C_1 \rangle \subseteq \langle C_2 \rangle$, we thus have $\langle C' \rangle = \langle C_1 \cup C_2 \rangle = \langle C_1 \rangle \cup \langle C_2 \rangle = \langle C_2 \rangle$. If C_2 is well-formed, we have that $\langle C_2 \rangle$ is well-formed; thus, $\langle C' \rangle$ is well-formed and so C' is well-formed. \square

Lemma A.17. Let C be a well-formed constraint set and let c_1 be a contour such that $c_1 \leq c_2$, $c_1 \leq c_3$, and $\{c_2, c_3\} \subseteq \langle C \rangle$. Then $c_2 = c_3$.

Proof. Because C is well-formed, we know that $c_1 \neq \emptyset$. By Lemma A.9, we have that $c_2 \not\bowtie c_3$. Again because C is well-formed, we have that $\forall c, c' \in C. c \neq c' \implies c \not\bowtie c'$. Since $c_2 \not\bowtie c_3$ we must therefore have $c_2 = c_3$. \square

Lemma A.18. Let \mathcal{C} be well-formed and let $\mathcal{C}' \subseteq \mathcal{C}$; then \mathcal{C}' is also well-formed. If C and C' are well-formed and $\langle C' \rangle \subseteq \langle C \rangle$, then $C \cup C'$ is well-formed. If C is well-formed and $C' \subseteq C$, then C' is also well-formed.

Proof. Because \mathcal{C} is well-formed, we have that $\emptyset \notin \mathcal{C}$ and that $\forall c_1, c_2 \in \mathcal{C}. c_1 \cap c_2 = \emptyset$. Because $\mathcal{C}' \subseteq \mathcal{C}$, we have that $\emptyset \notin \mathcal{C}'$. Likewise, we have that $\forall c_1, c_2 \in \mathcal{C}'. c_1 \cap c_2 = \emptyset$. Thus, \mathcal{C}' is well-formed.

If $\langle C' \rangle \subseteq \langle C \rangle$, the fact that $C \cup C'$ is well-formed reduces to the above.

If $C' \subseteq C$, then $\langle C' \rangle \subseteq \langle C \rangle$; this can be established from Definition 5.9 because $\langle \cdot \rangle$ is defined pointwise on constraint sets due to being a natural homomorphism. As a result, that C' is well-formed reduces to the above. \square

A.5 Simulation Properties

We next show useful properties of the simulation relation in Definition A.7.

Lemma A.19. Suppose $\overrightarrow{x_{\square} = r_{\square}} \preceq C$. Then $[x_1 = r_1, \dots, x_{i-1} = r_{i-1}, x_{i+1} = r_{i+1}, \dots, x_n = r_n] \preceq C$.

Proof. Immediate from Definition A.7. \square

Lemma A.20. Suppose $\overrightarrow{x'_{\square} = r'_{\square}} \preceq C'$ and $\overrightarrow{x''_{\square} = r''_{\square}} \preceq C''$. Let $\overrightarrow{x_{\square} = r_{\square}}$ be defined such that the latter list has been interjected into the former at some position k ; that is, for some k in $\{0..n\}$, we have $x_j = x'_j$ for all $j \leq k$, $x_j = x'_{j-m}$ for all $j > k + m$, and $x_j = x''_{j-k}$ otherwise (and likewise for r). Let $C = C' \cup C''$. Then $\overrightarrow{x_{\square} = r_{\square}} \preceq C$.

As a corollary, if $e \preceq C$ for some e and C , then $e \preceq C \cup C'$ for any C' .

Proof. By repeated application of Definition A.7, we have that $\prod \|[x'_1 = r'_1] \parallel \dots \parallel [x'_n = r'_n] \preceq C''' \cup \{c'_1\} \cup \dots \cup \{c'_n\}$ for some C''' such that $C''' \cup \{c'_1\} \cup \dots \cup \{c'_n\} = C'$. Likewise, $\prod \|[x''_1 = r''_1] \parallel \dots \parallel [x''_n = r''_n] \preceq C'''' \cup \{c''_1\} \cup \dots \cup \{c''_n\}$ for some C'''' such that $C'''' \cup \{c''_1\} \cup \dots \cup \{c''_n\} = C''$. We choose $C''' = C'$ and $C'''' = C''$ for convenience.

We aim to show that $\overrightarrow{x_{\square} = r_{\square}} \preceq C$. Because $C = C' \cup C''$, we have that $C = C' \cup C'' \cup c'_{\square} \cup c''_{\square}$. We can thus show that $\prod \|[x_1 = r_1] \parallel \dots \parallel [x_{n+m} = r_{n+m}] \preceq C \cup \{c_1\} \cup \dots \cup \{c_{n+m}\}$ by choosing each c as we chose x : we let $c_j = c'_j$ for all $j \leq k$, $c_j = c'_{j-m}$ for all $j > k + m$, and $c_j = c''_{j-k}$ otherwise. By Definition A.7, we are finished.

The corollary can be shown by the degenerate case in which $m = 0$. \square

Lemma A.21. Suppose $\mathbf{x} \preceq \alpha$ and $\mathbf{x}' \preceq \alpha'$. Then $\mathbf{x} \cup \mathbf{x}' \preceq \alpha \cup \alpha'$.

Proof. From $\mathbf{x} \preceq \alpha$ we have that for each x in \mathbf{x} there exists some α in α such that $x \preceq \alpha$; thus, there exists some α in $\alpha \cup \alpha'$ such that $x \preceq \alpha$. The same argument can be made regarding \mathbf{x}' . As a result, every x in $\mathbf{x} \cup \mathbf{x}'$ has some α in $\alpha \cup \alpha'$ such that $x \preceq \alpha$. By Definition A.7, we are finished. \square

Lemma A.22. If $e \preceq C$, then $x_2 = x_1 \in e \implies \exists \alpha_1, \alpha_2. \alpha_1 <: \alpha_2 \in C \wedge x_2 = x_1 \preceq \alpha_1 <: \alpha_2$. Corresponding statements hold true for the remaining expression and constraint forms.

Proof. Immediate from Definition A.7. \square

Lemma A.23. If $e \preceq C$ and C is well-formed, then $x_2 = x_1 \in e \wedge x_3 = x_2 \in e \implies \exists \alpha_1, \alpha_2, \alpha_3 \in C. \{\alpha_1 <: \alpha_2, \alpha_2 <: \alpha_3\} \in C \wedge \forall i \in \{1..3\}. x_i \preceq \alpha_i$. Similar statements hold for other sequences of constraints.

Proof. From Lemma A.22, we have that there exist some $\alpha_1, \alpha_2, \alpha_2'$, and α_3 such that $\{\alpha_1 <: \alpha_2, \alpha_2' <: \alpha_3\} \in C$, $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, $x_2 \preceq \alpha_2'$, and $x_3 \preceq \alpha_3$. It remains to show that $\alpha_2 = \alpha_2'$.

Let $x_2 \leftrightarrow \langle \iota_1, n_1, \dot{c}_1 \rangle$, let $\alpha_2 = \langle \iota_2, n_2, \dot{c}_2 \rangle$, and let $\alpha_2' = \langle \iota_2', n_2', \dot{c}_2' \rangle$. From $x_2 \preceq \alpha_2$ and $x_2 \preceq \alpha_2'$, we have that $\iota_1 = \iota_2 = \iota_2'$ and that $n_1 = n_2 = n_2'$. We have two cases: one in which $\dot{c}_1 = \spadesuit$ and one in which it does not. If $\dot{c}_1 = \spadesuit$, then $x_2 \preceq \alpha_2$ gives us by Definition A.7 that $\dot{c}_2 = \spadesuit$; similarly, $x_2 \preceq \alpha_2'$ gives us that $\dot{c}_2' = \spadesuit$. Thus, $\alpha_2 = \alpha_2'$.

Otherwise, $\dot{c}_1 = c_1$. By $x_2 \preceq \alpha_2$ and $x_2 \preceq \alpha_2'$, we have that $\dot{c}_2 = c_2$ and $\dot{c}_2' = c_2'$ such that $c_1 \leq c_2$ and $c_1 \leq c_2'$. Because C is well-formed, we know by Lemma A.17 that $c_2 = c_2'$; thus, $\alpha_2 = \alpha_2'$. \square

A.6 Initial Derivation Simulation

We now show that the initial type derivation of an expression simulates that expression. To do so, we begin by defining the notion of a *canonical* context.

Definition A.24 (Canonical Context). A context Γ is *canonical* iff for all mappings $\langle \iota, n, \dot{c} \rangle : \langle \iota', n', \dot{c}' \rangle$ in Γ , we have $\iota = \iota', n = n'$, and $\dot{c} = c \leq c' = \dot{c}'$.

We now show some lemmas over canonical contexts.

Lemma A.25. If Γ and Γ' are canonical, then $\Gamma \cup \Gamma'$ is canonical.

Proof. Trivial by definition of set union. \square

Lemma A.26. Given $\{x : \alpha\} \in \Gamma$ where Γ is canonical, $x \preceq \alpha$.

Proof. Let $x \leftrightarrow \langle \iota, n, \dot{c} \rangle$. By Definition A.24, $\dot{c} = c$ and $\alpha = \langle \iota, n, c' \rangle$ such that $c \leq c'$. By Definition A.7, $x \preceq \alpha$. \square

Our next objective is to show that the type and constraint set given by initial type derivation simulates the input expression. We begin by showing the simulation of initial pattern derivation.

Lemma A.27. For any ϕ and canonical Γ such that $\llbracket \phi \rrbracket_{\text{p}}^{\Gamma} = \langle \Gamma', \tau \phi \rangle$, we have that Γ' is canonical, every variable in Γ' has the non-contour \clubsuit , and that $\phi \preceq \tau \phi$.

Proof. By case analysis.

If $\phi = x : \text{int}$, then $\llbracket \phi \rrbracket_{\text{p}}^{\Gamma} = \langle \{x : \dot{\alpha}_1\}, \dot{\alpha}_1 \sim \text{int} \rangle$. By Definition 5.2, we have that $\dot{\alpha}_1 = \langle \iota, 1, \clubsuit \rangle$ where ι identifies this pattern. By Definition A.2, we have that $x \leftrightarrow \langle \iota, 1, \clubsuit \rangle$. Thus $\{x : \dot{\alpha}_1\}$ is canonical by Definition A.24. By Definition A.7, we have that $x \preceq \dot{\alpha}_1$; thus, we have $x : \text{int} \preceq \dot{\alpha}_1 \sim \text{int}$ and we are finished. Because the only variable appearing in the context is of the form $\dot{\alpha}_i$ and all such variables have the non-contour \clubsuit , the context has only variables with the non-contour.

Similar arguments hold for the $x : \text{lbl } x'$ and $x : \text{any}$ cases. \square

In order to support the lemmas over initial derivation, we will need to show some properties that align the bound variable set of a scape with the bound variable set of its type. We start by defining a function to retrieve the upper bound of a constraint; it is also overloaded to find the upper bounds of a constraint set.

Definition A.28 (Constraint Upper Bounds).

$$\text{UB}(c) = \begin{cases} \alpha & \text{when } c = \tau <: \alpha \\ \alpha & \text{when } c = \alpha' <: \alpha \\ \alpha & \text{when } c = \alpha' \alpha'' <: \alpha \\ \alpha & \text{when } c = \alpha' \text{ op } \alpha'' <: \alpha \end{cases}$$

$$\text{UB}(C) = \{\text{UB}(c) \mid c \in C\}$$

We then show the simple property that the upper bounds of scape type's constraint set simulate the bound variables of its scape.

Lemma A.29. For any e such that $\llbracket e \rrbracket_{\text{E}}^{\Gamma} = \langle \alpha', C \rangle$, we have $\text{BV}(e) \preceq \text{UB}(C)$.

Proof. Let $e = \overrightarrow{x \equiv r}$. For all $1 \leq i \leq n$, let $\alpha_i = \langle \iota_i, 1, \clubsuit \rangle$ where ι_i is the identity of $x_i = r_i$. Let $\Gamma' = \{x_i : \alpha_i \mid 1 \leq i \leq n\}$. Then by Figure 4.2, we have that $\alpha' = \alpha_n$ and that $C = \overrightarrow{c}$ where for all $1 \leq i \leq n$ we have $c_i = \llbracket x_i = r_i \rrbracket_{\text{R}}^{\Gamma \cup \Gamma'}$.

By Definition A.3, we have that $\text{BV}(e) = \{x_i \mid 1 \leq i \leq n\}$. By Definition A.28, we have that $\text{UB}(C) = \{\text{UB}(c_i) \mid 1 \leq i \leq n\}$. By Definition A.7, we have that $\{x_i \mid 1 \leq i \leq n\} \preceq \{\text{UB}(c_i) \mid 1 \leq i \leq n\}$ if $\forall x \in \{x_i \mid 1 \leq i \leq n\}. \exists \alpha \in \{\text{UB}(c_i) \mid 1 \leq i \leq n\}. x \preceq \alpha$. It is therefore sufficient to show that $\forall i \in \{1..n\}. x_i \preceq \text{UB}(c_i)$. More generally, it is sufficient to show that for all $\llbracket x' = r' \rrbracket_{\text{R}}^{\Gamma} = c'$ we have $x' \preceq \text{UB}(c')$. This is trivial by case analysis. \square

We next show some supporting lemmas.

Lemma A.30. If $\phi \preceq \tau \phi$ then $\text{BV}(\phi) \preceq \text{FTV}(\tau \phi)$.

Proof. By case analysis. If $\phi = x : \text{int}$ then by Definition A.7 we have $\tau \phi = \alpha \sim \text{int}$ such that $x \preceq \alpha$. We have that $\text{BV}(\phi) = \{x\}$ and that $\text{FTV}(\tau \phi) = \{\alpha\}$ and so this case is finished. A similar argument is made when $\phi = x : \text{any}$.

Otherwise, $\phi = x : \text{lbl } x'$ and by Definition A.7 we have $\tau \phi = \alpha \sim \text{lbl } \alpha'$ such that $x \preceq \alpha$ and $x' \preceq \alpha'$. We have that $\text{BV}(\phi) = \{x, x'\}$ and that $\text{FTV}(\tau \phi) = \{\alpha, \alpha'\}$ and so we are finished. \square

We now present the proof that initial derivation simulates as expected.

Lemma A.31. Let e be an expression $\overrightarrow{x \equiv r}$ with unique bindings closed over a canonical context Γ_1 such that no variable is bound in both e and Γ_1 and all variables in Γ_1 have the non-contour \clubsuit . If $\llbracket e \rrbracket_{\text{E}}^{\Gamma_1} = \langle \alpha, C \rangle$, then $e \preceq \langle \alpha, C \rangle$. Further, C is well-formed and every type variable in $\text{FTV}(C)$ is either in Γ_1 or simulates some variable in \overrightarrow{x} .

Proof. By structural induction on the size of e .

For all $1 \leq i \leq n$, let $\alpha_i = \langle \iota_i, 1, \clubsuit \rangle$ where ι_i is the identity of $x_i = r_i$. Let $\Gamma_2 = \{x_i : \alpha_i \mid 1 \leq i \leq n\}$. Observe that Γ_2 is canonical by Definition A.24. Let $\Gamma_3 = \Gamma_1 \cup \Gamma_2$. We observe that this context contains only variables with the non-contour. Also, by Lemma A.25 this context is canonical. For all $1 \leq i \leq n$, let $c_i = \llbracket x_i = r_i \rrbracket_{\text{R}}^{\Gamma_3}$; then $C = \overrightarrow{c}$ and $\alpha = \alpha_n$.

To show the property that $e \preceq \langle \alpha, C \rangle$, it is thus sufficient to show that $e \preceq C$. By induction on the length of e using Definition A.7, it is sufficient to show that $x_i = r_i \preceq c_i$ for all $1 \leq i \leq n$ to prove this property. To show the property that every type variable in $\text{FTV}(C)$ is either in Γ_1 or simulates some variable in \overrightarrow{x} , it is sufficient to show that every type variable in $\text{FTV}(c_i)$ is in Γ_3 for all $1 \leq i \leq n$. We defer showing the property that C is well-formed until the end of

the proof. We now show each of these properties for every c_i by case analysis on the redex r_i for each $1 \leq i \leq n$.

If r_i is x' , then $\llbracket x_i = r_i \rrbracket_{\mathbb{R}}^{\Gamma_3} = \Gamma_3(x') <: \Gamma_3(x)$. By Lemma A.26, $x \preceq \Gamma_3(x)$ and $x' \preceq \Gamma_3(x')$. It is trivial that the free type variables for this constraint are in Γ_3 ; they are retrieved directly from that context. Thus by Definition A.7, this case is finished. Similar arguments hold when r_i is $\text{lbl } x'$, $x' \& x''$, $x' x''$, or $x' \text{ op } x''$.

If r_i is $()$, then $\llbracket x_i = r_i \rrbracket_{\mathbb{R}}^{\Gamma_3} = () <: \Gamma_3(x)$. By Lemma A.26, we have $x \preceq \Gamma_3(x)$; by Definition A.7, we have $() \preceq ()$. Thus by Definition A.7, this case is finished. A similar argument holds when r_i is some $n \in \mathbb{Z}$.

If r_i is of the form $\phi \rightarrow e'$, then let $x'_\square = r'_\square = e'$. We have that $\llbracket x_i = r_i \rrbracket_{\mathbb{R}}^{\Gamma_3} = \forall \alpha. \tau \phi \rightarrow \alpha' \setminus C'$ such that (1) $\llbracket \phi \rrbracket_{\mathbb{P}}^{\Gamma_3} = \langle \Gamma_4, \tau \phi \rangle$ for some Γ_4 , (2) $\llbracket e' \rrbracket_{\mathbb{E}}^{\Gamma_3 \cup \Gamma_4} = \langle \alpha', C' \rangle$, and (3) $\alpha = (\text{FTV}(\tau \phi) \cup \text{FTV}(C')) - \alpha'$ where $\alpha' = \{\alpha'' \mid x : \alpha'' \in \Gamma_3\}$. By Lemma A.27, we have that Γ_4 is canonical and only contains variables with the non-contour.

Let $\Gamma_5 = \Gamma_3 \cup \Gamma_4$; observe that this context contains only variables with the non-contour and, by Lemma A.25, is canonical. Because (1) Γ_1 contains no variables bound in e , (2) Γ_2 contains only variables bound directly by e , (3) Γ_4 contains only variables bound by $\tau \phi$, and (4) variable bindings are unique in e , we know that Γ_5 does not contain any bindings that appear in e' . Thus by the induction hypothesis we have that $e' \preceq \langle \alpha', C' \rangle$ and C' is well-formed and every type variable in $\text{FTV}(C')$ is either in Γ_5 or simulates some variable x'_j where $1 \leq j \leq m$. Suppose we were to show that $\text{BV}(\phi \rightarrow e') \preceq \alpha$; then Definition A.7 gives us that $r_i \preceq \forall \alpha. \tau \phi \rightarrow \alpha' \setminus C'$, by Lemma A.26 we have $x \preceq \Gamma_3(x)$, and so by Definition A.7 we would have $x_i = r_i \preceq c_i$. So it remains to show that $\text{BV}(\phi \rightarrow e') \preceq \alpha$ and that every free type variable in c_i is in Γ_3 .

We have from above that free type variable in C' is either in Γ_5 or simulates some variable x'_j . This is to say that the variable is either in Γ_3 , in Γ_4 , or simulates some variable x'_j . The constraint c_i contains a type variable known to be in Γ_3 and the type of the scape. The type of the scape captures every free variable in $\tau \phi$ (those appearing in Γ_4) and every free variable in C' ; thus, the scape type has no free variables and so the only free variable in c_i is in Γ_3 . It remains to show that $\text{BV}(\phi \rightarrow e') \preceq \alpha$.

By Definition A.3 and the above, $\text{BV}(\phi \rightarrow e') \preceq \alpha$ is equivalent to $\text{BV}(\phi) \cup \text{BV}(e') \preceq (\text{FTV}(\tau \phi) \cup \text{FTV}(C')) - \alpha'$. By Lemma A.21, it is sufficient to show that $\text{BV}(\phi) \preceq \text{FTV}(\tau \phi) - \alpha'$ and that $\text{BV}(e') \preceq \text{FTV}(C') - \alpha'$. Because Γ_3 does not contain any variables in ϕ and because Γ_3 is canonical, it does not contain any variables in $\tau \phi$; therefore, the prior condition reduces to $\text{BV}(\phi) \preceq \text{FTV}(\tau \phi)$. Because $\phi \preceq \tau \phi$, Lemma A.30 gives us this property.

It thus remains to show that $\text{BV}(e') \preceq \text{FTV}(C') - \alpha'$. By Lemma A.29, it is sufficient to show that $\text{FTV}(C') - \alpha' \subseteq \text{UB}(C')$. From above, we know that every type variable in

$\text{FTV}(C')$ is either in α' or simulates a variable in $x'^{\overrightarrow{m}}$, which is the definition of $\text{BV}(e')$; thus, this case is finished.

We now show the deferred condition that C is well-formed. Every variable appearing in that constraint set is either from the above Γ_2 or Γ_4 (which have only variables with the non-contour), from Γ_1 (in which every variable has the non-contour), or is from some induction on this lemma. Since every variable in C has the non-contour, $\langle C \rangle = \emptyset$ and so C is well-formed. \square

A.7 Simulation Preservation

We now show that various relations used during constraint closure preserve constraint set simulation.

Lemma A.32. If $E \preceq C$, then $E(x) = v \wedge x \preceq \alpha \implies \tau \overset{C}{\prec} \alpha \wedge v \preceq \tau$.

Proof. For $E(x) = v$ to be defined, there must exist some $x = v' \in E$. We have two cases: either v' is of the form x' or it is not. If it is not, then $v = v'$ by the definition of $E(x)$. By $E \preceq C$ and Lemma A.22, we know that $\tau <: \alpha \in C$ such that $x \preceq \alpha$ and $v \preceq \tau$. By Definition 4.1, we have that $\tau \overset{C}{\prec} \alpha$.

In the case that v' is of the form x' , then there exists some $x = x' \in E$. By $E \preceq C$ and Lemma A.22, we have that there exists some $\alpha' <: \alpha \in C$ such that $x \preceq \alpha$ and $x' \preceq \alpha'$. By definition of $E(x)$, we know that $E(x) = E(x')$. By the induction hypothesis, we have that $\tau \overset{C}{\prec} \alpha' \wedge v \preceq \tau$; such induction is sound because we only consider closed e with unique variable names. Because $\tau \overset{C}{\prec} \alpha'$ and $\alpha' <: \alpha \in C$, we have by Definition 4.1 that $\tau \overset{C}{\prec} \alpha$. \square

Lemma A.33. If $E \preceq C$, $E \downarrow_{\pi}^*(x) = \overset{n}{v}$, and $x \preceq \alpha$, then $C \vdash \alpha \overset{n}{\tau} \wedge \forall i \in \{1..n\}. v_i \preceq \tau_i$.

Proof. First, we observe that the projection matching operation in Definition 4.3 properly simulates projector matching in the operational semantics; that is, $v \in \pi \wedge v \preceq \tau$ iff $\tau \sqsubseteq \pi$; this is shown by simple case analysis on π . We then show this lemma by induction on the size of the E preceding the definition of x (which is sound because e is closed and has unique variable names).

If $E(x)$ is a non-onion value v' and $v' \in \pi$, then by Lemma A.32 we have $\tau' \overset{C}{\prec} \alpha$ and $v' \preceq \tau'$ for some non-onion type τ' . We also have from the above that $\pi \sqsubseteq \alpha'$. By Definition 3.1, $\overset{n}{v} = [v']$; by Definition 4.4, $\overset{n}{\tau} = [\tau']$. Because $v_1 \preceq \tau_1$ and $n = 1$, this case is finished.

If $E(x)$ is a non-onion value v' and $v' \notin \pi$, then again by Lemma A.32 we have $\tau' \overset{C}{\prec} \alpha$ and $v' \preceq \tau'$ for some non-onion type τ' . We also have from the above that $\pi \not\sqsubseteq \alpha'$. By Definition 3.1, $\overset{n}{v} = []$; by Definition 4.4, $\overset{n}{\tau} = []$. Therefore this case is finished.

If $E(x)$ is an onion value $x_1 \& x_2$, then $\overset{n}{v} = \overset{n_1}{v'} \parallel \overset{n_2}{v''}$ where $n_1 + n_2 = n$, $E \downarrow_{\pi}^*(x_1) = \overset{n_1}{v'}$, and $E \downarrow_{\pi}^*(x_2) = \overset{n_2}{v''}$. By Lemma A.32 we have $\alpha_1 \& \alpha_2 \overset{C}{\prec} \alpha$ and $x_1 \& x_2 \preceq \alpha_1 \& \alpha_2$.

By Definition A.7, we have that $x_1 \preceq \alpha_1$ and $x_2 \preceq \alpha_2$. We therefore induct on this lemma for each side of the onion; the left induction gives $C \vdash \alpha_1 \xrightarrow{\tau_1} \tau_1'$ and $\forall i \in \{1..n_1\}. v_i' \preceq \tau_i'$ while the right induction gives $C \vdash \alpha_2 \xrightarrow{\tau_2} \tau_2'$ and $\forall i \in \{1..n_2\}. v_i'' \preceq \tau_i''$. By Definition 3.1, $\vec{v} = \vec{v}' \parallel \vec{v}''$; by Definition 4.4, $\vec{\tau} = \vec{\tau}' \parallel \vec{\tau}''$. We have from above that $\forall i \in \{1..n\}. v \preceq \tau$; we are therefore finished. \square

Lemma A.34. If $E \preceq C$, $E \downarrow_{\pi}(x) = v$, and $x \preceq \alpha$, then $C \vdash \alpha \xrightarrow{\tau} \tau \wedge v \preceq \tau$.

Proof. Immediate from Lemma A.33. \square

Lemma A.35. If $E \preceq C$, $x \preceq \alpha$, $\varphi \preceq \tau_{\varphi}$, and $x \preceq_E \varphi \setminus E_0$, then $\alpha \preceq_C \tau_{\varphi} \setminus C'$ for some C' such that $E_0 \preceq C'$.

Proof. By case analysis.

If $\varphi = \text{int}$, then because $\varphi \preceq \tau_{\varphi}$ and by Definition A.7 we have that $\tau_{\varphi} = \text{int}$. By Definition 3.3, the only case in which $\varphi = \text{int}$ results in compatibility is when $E \downarrow_{\text{int}}(x) = n$ and $n \in \mathbb{Z}$; in this case, $E_0 = []$. Because $x \preceq \alpha$, Lemma A.34 gives us that $C \vdash \alpha \xrightarrow{\text{int}} \tau$ such that $n \preceq \tau$. Because of this and $n \in \mathbb{Z}$, we have by Definition A.7 that $\tau = \text{int}$. By Definition 4.6, we have that $\alpha \preceq_C \text{int} \setminus \emptyset$. Because Definition A.7 gives us that $[] \preceq \emptyset$, this case is finished.

If $\varphi = \text{lbl } x'$, then because $\varphi \preceq \tau_{\varphi}$ and by Definition A.7 we have that $\tau_{\varphi} = \text{lbl } \alpha'$ such that $x' \preceq \alpha'$. By Definition 3.3, the only case in which $\varphi = \text{lbl } x'$ results in compatibility is when $E \downarrow_{\text{lbl}}(x) = \text{lbl } x''$; in this case, $E_0 = [x' = x'']$. Because $x \preceq \alpha$, Lemma A.34 gives us that $C \vdash \alpha \xrightarrow{\text{lbl}} \tau$ such that $\text{lbl } x'' \preceq \tau$. By Definition A.7, this gives us that $\tau = \text{lbl } \alpha''$ such that $x'' \preceq \alpha''$. By Definition 4.6, we have that $\alpha \preceq_C \text{lbl} \setminus \{\alpha'' <: \alpha'\}$. Because $x'' \preceq \alpha''$ and $x' \preceq \alpha'$, we have by Definition A.7 that $[x' = x''] \preceq \{\alpha'' <: \alpha'\}$; thus, this case is finished.

If $\varphi = \text{any}$, then because $\varphi \preceq \tau_{\varphi}$ and by Definition A.7 we have that $\tau_{\varphi} = \text{any}$. By Definition 3.3, we have that $E_0 = []$; by Definition 4.6, we have that $C' = \emptyset$. Because Definition A.7 gives us that $[] \preceq \emptyset$, we are finished. \square

Lemma A.36. If $E \preceq C$, $x \preceq \alpha$, $\phi \preceq \tau_{\phi}$, and $x \preceq_E \phi \setminus E_0$, then $\alpha \preceq_C \tau_{\phi} \setminus C''$ for some C'' such that $E_0 \preceq C''$.

Proof. By the grammar of patterns in Figure 3.1, we have that $\phi = x' : \varphi$ for some φ . Because $\phi \preceq \tau_{\phi}$, we have that $\tau_{\phi} = \alpha' \sim \tau_{\varphi}$ for some τ_{φ} such that $\varphi \preceq \tau_{\varphi}$ and $x' \preceq \alpha'$. By Definition 3.3, the only case in which compatibility holds over ϕ is when $x' \preceq_E \varphi \setminus E_0'$ for some E_0' ; in this case, $E_0 = E_0' \parallel [x' = x]$. By Lemma A.35, we have that $\alpha \preceq_C \tau_{\varphi} \setminus C'$ such that $E_0' \preceq C'$. By Definition 4.6, this gives us that $\alpha \preceq_C \alpha' \sim \tau_{\varphi} \setminus C' \cup \{\alpha <: \alpha'\}$.

It remains to show that $E_0' \parallel [x' = x] \preceq C' \cup \{\alpha <: \alpha'\}$. By Lemma A.20 it is sufficient to show that $E_0' \preceq C'$ (which we already know from above) and that $[x' = x] \preceq \{\alpha <: \alpha'\}$.

By Definition A.7, this is equivalent to showing that $x \preceq \alpha$ (which is a premise of this lemma) and that $x' \preceq \alpha'$ (which we know from the above). \square

Lemma A.37. If $E \preceq C$, $x \preceq \alpha$, $\phi_{\square} \xrightarrow{n} e_{\square} \preceq \vec{\tau}$, and $x \preceq_E \phi_{\square} \xrightarrow{n} e_{\square} \setminus E_0$; e' , then there exist some α' and C' such that $\alpha \preceq_C \vec{\tau} \setminus \alpha'$; C' and $E_0 \parallel e' \preceq \langle \alpha', C' \rangle$.

Proof. By induction on n .

Let $\vec{\tau} = \vec{\tau}_{\alpha_{\square}, \tau_{\phi_{\square}} \rightarrow \alpha_{\square}'' \setminus C_{\square}''}$. If $x \preceq_E \phi_{\square} \xrightarrow{n} e_{\square} \setminus E_0$, then by Lemma A.36 we have that $\alpha \preceq_C \tau_{\phi_{\square}} \setminus C'''$ such that $E_0 \preceq C'''$. We have from $\phi_{\square} \xrightarrow{n} e_{\square} \preceq \vec{\tau}$ and Definition A.7 that $e_{\square} \preceq \langle \alpha_{\square}'', C_{\square}'' \rangle$. By Definition 5.7 we have that $\alpha \preceq_C \vec{\tau} \setminus \alpha_{\square}''$; $C_{\square}'' \cup C'''$; thus, $\alpha' = \alpha_{\square}''$ and $C' = C_{\square}'' \cup C'''$. Because $e_{\square} \preceq \langle \alpha_{\square}'', C_{\square}'' \rangle$, we know that $e_{\square} \preceq C_{\square}''$; by Lemma A.20, we have that $E_0 \parallel e_{\square} \preceq C_{\square}'' \cup C'''$. Because E_0 was prepended (rather than appended) to e_{\square} , the last variable in the list is still the last variable in e_{\square} ; thus, we have $E_0 \parallel e_{\square} \preceq \langle \alpha_{\square}'', C_{\square}'' \cup C''' \rangle$ and this case is finished.

Otherwise, we simply induct on this lemma. We know that induction is possible; if $n = 0$ then the statement that $x \preceq_E \phi_{\square} \xrightarrow{n} e_{\square} \setminus E_0$; e' is a contradiction. \square

Variable instantiation preserves the simulation relation. To demonstrate this property, we first prove a supplementary lemma.

Lemma A.38. Suppose that $e \preceq C$, that $\mathbf{x} \preceq \alpha$, and that $c_1 \preceq c_2$. Then for any x' appearing in e , there is a corresponding α' appearing in C such that $E_{\text{INST}}(x', \mathbf{x}, c_1) \preceq \text{INST}(\alpha', \alpha, c_2)$.

Proof. For any x' in e , Lemma A.22 gives us that there is a corresponding α' in C such that $x' \preceq \alpha'$. By Definition A.2, we have that $x' \leftrightarrow \langle \iota, n, \dot{c}_1 \rangle$ and $\alpha' = \langle \iota, n, \dot{c}_2 \rangle$ such that $\dot{c}_1 \preceq \dot{c}_2$.

Suppose that $x' \in \mathbf{x}$ and $\dot{c}_1 = \ast$. In this case, $E_{\text{INST}}(x', \mathbf{x}, c_1) = x'$. Because $x' \preceq \alpha'$, we know that \dot{c}_2 is \ast . Because of this and $\mathbf{x} \preceq \alpha$, we also know that $\alpha' \in \alpha$. Thus, $\text{INST}(\alpha', \alpha, c_2) = \alpha'$. Because $x' \preceq \alpha'$, this case is finished.

Otherwise, suppose that $\dot{c}_1 \neq \ast$. In that case, $E_{\text{INST}}(x', \mathbf{x}, c_1) = x'$. Because $x' \preceq \alpha'$, we also know that $\dot{c}_2 \neq \ast$; thus, $\text{INST}(\alpha', \alpha, c_2) = \alpha'$. Thus this case is finished.

Finally, suppose that $\dot{c}_1 = \ast$ but $x' \notin \mathbf{x}$. In this case, $E_{\text{INST}}(x', \mathbf{x}, c_1) = \langle \iota, n, c_1 \rangle$. Because $x' \preceq \alpha'$, we know that \dot{c}_2 is \ast . By Definition A.7, the fact that $\mathbf{x} \preceq \alpha$ means that every $\alpha'' \in \alpha$ has some $x'' \in \mathbf{x}$ such that $x'' \preceq \alpha''$. For variables without a contour (such as α'), this implies that $x'' \leftrightarrow \langle \iota', n', \ast \rangle$ such that $\alpha'' = \langle \iota', n', \ast \rangle$. Thus, because $\dot{c}_1 = \dot{c}_2 = \ast$, $x' \in \mathbf{x}$ implies $\alpha' \in \alpha$. Using this, we conclude that $\text{INST}(\alpha', \alpha, c_2) = \langle \iota, n, c_2 \rangle$. Because $c_1 \preceq c_2$ and by Definition A.7, we are finished. \square

We now show that simulation is preserved through instantiation.

Lemma A.39. Suppose that $e \preceq C$, that $\mathbf{x} \preceq \alpha$, and that $c_1 \preceq c_2$. Then $\text{EINST}(e, \mathbf{x}, c_1) \preceq \text{INST}(C, \alpha, c_2)$.

Proof. By induction on Definition A.7.

For variables, simulation is demonstrated by Lemma A.38. For all lower bound types other than functions and all upper bound types, patterns, constraints and contexts, simulation is parametric in variables and so is also demonstrated by Lemma A.38.

For function lower bound types, we must show that for any $v = \phi \rightarrow e'$ appearing in e , there is a corresponding $\tau = \forall \alpha'. \tau \phi \rightarrow \alpha \setminus C'$ appearing in C such that $\text{EINST}(v, \mathbf{x}, c_1) \preceq \text{INST}(\tau, \alpha, c_2)$. For any such v in e , Lemma A.22 gives us that such a τ exists and that $v \preceq \tau$. By Definition A.7, this means that $\phi \preceq \tau \phi$, $e' \preceq C'$, and that $\text{BV}(v) \preceq \alpha$. Since each of these terms are smaller than v or τ accordingly, we induct on the definition of this lemma and so function lower bound types are also simulated.

By Definition A.7, we know that for each $x = v \in e$, there is a corresponding $c \in C$ such that $x = v \preceq c$. Suppose that $\text{EINST}(x = v, \mathbf{x}, c_1) = x' = v'$, $\text{INST}(c, \alpha, c_2) = c'$, $\text{EINST}(e, \mathbf{x}, c_1) = e'$, and $\text{INST}(C, \alpha, c_2) = C'$. Because we have that simulation of each constraint is preserved by instantiation, we have that $x' = v' \preceq c'$. Since this is true for each $x' = v'$ and c' , we have that each $x' = v'$ in e' is simulated by some c' in C' ; thus $\text{EINST}(E, \mathbf{x}, c_1) \preceq \text{INST}(C, \alpha, c_2)$. \square

Lemma A.40. If $e \preceq C$ and $c_1 \preceq c_2$, then $\text{EFRESH}(e, c_1) \preceq \text{INST}(C, c_2)$.

Proof. Immediate from Lemma A.39. \square

Contour replacement also preserves simulation. We use a strategy similar to that above.

Lemma A.41. If $x \preceq \alpha$ then $x \preceq \text{REPL}(\alpha, C)$.

Proof. Let $\text{REPL}(\alpha, C) = \alpha'$. Then, let $x \leftrightarrow \langle \iota_1, n_1, \dot{c}_1 \rangle$ and $\alpha = \langle \iota_2, n_2, \dot{c}_2 \rangle$ and $\alpha' = \langle \iota_3, n_3, \dot{c}_3 \rangle$. Because $x \preceq \alpha$, we know that $\iota_1 = \iota_2$ and $n_1 = n_2$; by Definition 5.10, contour replacement is naturally homomorphic down to (but not including) contours and so we know that $\iota_2 = \iota_3$ and $n_2 = n_3$. It remains to show that $\dot{c}_1 \preceq \dot{c}_3$.

$x \preceq \alpha$ also gives us that $\dot{c}_1 = \ast$ iff $\dot{c}_2 = \ast$. If $\dot{c}_1 = \ast$, then $\dot{c}_2 = \ast$ and, by Definition 5.10, $\dot{c}_3 = \ast$; this case is finished.

If $\dot{c}_1 \neq \ast$ then $\dot{c}_1 = c_1$ and $\dot{c}_2 = c_2$. By Definition A.7, it suffices to show that $c_1 \leq c_3$ where $c_3 = \dot{c}_3$. If $c_2 \not\leq c$, then $\dot{c}_3 = c_2$ and so $\dot{c}_1 \leq \dot{c}_3$ and this case is finished.

Otherwise, $c_2 \leq c$. Then $\dot{c}_3 = c$ and it suffices to show that $c_1 \leq c$. By Lemma A.10, it suffices to show that $c_1 \leq c_2$ and that $c_2 \leq c$; the latter we have from this case. We prove $c_1 \leq c_2$ by Definition A.7 because $x \preceq \alpha$ and so $c_1 \preceq c_2$. \square

Lemma A.42. If $E \preceq C$ then $E \preceq \text{REPL}(C, C)$.

Proof. Because contour replacement is a natural homomorphism in all cases not proven by Lemma A.41, this argument follows in parallel to that of Lemma A.39. \square

A.8 Contour Properties of Functions and Relations

We now prove contour-related properties about functions and relations which are necessary to show simulation of constraint closure.

Lemma A.43. $C \leq \text{COLLAPSE}(C)$

Proof. This is to show that $\llbracket C \rrbracket \subseteq \llbracket \{\text{COLLAPSE}(S) \mid S \in C\} \rrbracket$. By Definition 5.10, this is to show that $\bigcup \{\llbracket S \rrbracket \mid S \in C\} \subseteq \bigcup \{\llbracket S \rrbracket \mid S \in \text{COLLAPSE}(C)\}$, which is equivalent to $\bigcup \{\llbracket S \rrbracket \mid S \in C\} \subseteq \bigcup \{\llbracket S \rrbracket \mid S \in \{\text{COLLAPSE}(S) \mid S \in C\}\}$, which reduces to $\bigcup \{\llbracket S \rrbracket \mid S \in C\} \subseteq \bigcup \{\llbracket \text{COLLAPSE}(S) \rrbracket \mid S \in C\}$. It is therefore sufficient to prove that $\llbracket S \rrbracket \subseteq \llbracket \text{COLLAPSE}(S) \rrbracket$ for all S .

We proceed by strong induction on the number of cycles in S ; by ‘‘cycle’’, we mean a pair of distinct parts \mathcal{P}_1 and \mathcal{P}_2 in S such that $\text{SITES}(\mathcal{P}_1) \cap \text{SITES}(\mathcal{P}_2) \neq \emptyset$. The base case is trivial, since in the absence of a cycle we have $\text{COLLAPSE}(S) = S$. In the inductive case, let $S = \xrightarrow{n} \overline{\mathcal{P}}$. Let cycles in S be represented by pairs of indices (i, j) where $i < j$ and $\text{SITES}(\mathcal{P}_i) \cap \text{SITES}(\mathcal{P}_j) \neq \emptyset$; let $\overline{(i_\square, j_\square)}$ be the set of all cycles in S . Finally, let $S' = \overline{\mathcal{P}'} = [\mathcal{P}_1, \dots, \mathcal{P}_{i_k-1}, (\text{SITES}(\mathcal{P}_{i_k}) \cup \dots \cup \text{SITES}(\mathcal{P}_{j_k})), \mathcal{P}_{j_k+1}, \dots, \mathcal{P}_n]$ for some $1 \leq k \leq m$. We can construct the set of cycles in S' by forming a mapping M from each index in S to indices in S' such that all $1 \leq i'' < i_k$ map to themselves, all $i_k \leq i'' \leq j_k$ map to i_k , and all $j_k < i'' \leq n$ map to themselves minus $j_k - i_k$. For each (i, j) in the cycles of S , either $M[i] = M[j]$ or $(M[i], M[j])$ is a cycle in S' ; in particular, we know that $M[i_k] = M[j_k]$, which indicates that cycle k has been eliminated. From this, we have that the set of cycles in S' is strictly smaller than the set of cycles in S ; thus, it is sufficient by the induction hypothesis to show that $\llbracket S \rrbracket \subseteq \llbracket S' \rrbracket$. By repeated application of Lemma A.13, it remains to show that $\llbracket [\mathcal{P}_{i_k}, \dots, \mathcal{P}_{j_k}] \rrbracket \subseteq \llbracket [\text{SITES}(\mathcal{P}_{i_k}) \cup \dots \cup \text{SITES}(\mathcal{P}_{j_k})] \rrbracket$.

It is sufficient to show that $\llbracket \overline{\mathcal{P}} \rrbracket \subseteq \llbracket [\bigcup \text{SITES}(\mathcal{P}_\square)] \rrbracket$. This is immediate; the latter contains every call sequence composed of any elements in $\overline{\mathcal{P}}$ and the prior contains no call sequences containing elements not in $\overline{\mathcal{P}}$. \square

Lemma A.44. $\forall C. C \leq \text{WIDEN}(C, C)$

Proof. For any C , we have that $\text{WIDEN}(C, C) = c' \supseteq c$ by inspection. By definition, $c \leq c'$ iff $\llbracket c \rrbracket \subseteq \llbracket c' \rrbracket$. This reduces to $(\bigcup \{\llbracket S \rrbracket \mid S \in C\}) \subseteq (\bigcup \{\llbracket S \rrbracket \mid S \in c'\})$. Because $c \subseteq c'$, every element in the comprehension on the left also appears in the comprehension on the right; thus, $c \leq c'$. \square

Lemma A.45. If $\text{WIDEN}(C, C) = c'$, then $\forall c'' \in (C). c'' \not\leq c' \implies c'' \leq c'$.

Proof. By definition, $c' = c \cup (\bigcup \{c'' \mid c'' \in \langle C \rangle \wedge c \not\prec c''\})$. We therefore have for any c'' in $\langle C \rangle$ that $c'' \subseteq c'$. By Definition 5.10, we thus have that $\llbracket c'' \rrbracket \subseteq \llbracket c' \rrbracket$ and so $c'' \leq c'$. \square

Lemma A.46. If $C \vdash \alpha \xrightarrow{\tau} \bar{\tau}$ for some well-formed C , then $\langle \bar{\tau} \rangle \subseteq \langle \alpha \rangle \cup \langle C \rangle$.

Proof. By inspection. Each rule produces a type which is either (1) devoid of contours, (2) constructed from τ , (3) constructed from constraints found in C , or (4) derived from another projection. \square

Lemma A.47. If $\alpha \preceq_C \tau \phi \setminus C'$, then $\langle C' \rangle \subseteq \langle C \rangle \cup \langle \tau \phi \rangle \cup \langle \alpha \rangle$.

Proof. By inspection. The only cases in which constraints are added to the constraint set C' in the compatibility relation are the cases for $\alpha'_1 \sim \tau \phi$ and $\text{lbl} \alpha'_2$. In the prior case, the added constraint is comprised of α and α'_1 ; in the latter, it is comprised of α'_2 and some α'_3 which resulted from projection. By Lemma A.46, we know that $\langle \alpha'_3 \rangle \subseteq \langle \alpha \rangle \cup \langle C \rangle$. \square

Lemma A.48. If $\text{INST}(C, \alpha, c) = C'$, then $\langle C' \rangle \subseteq \langle C \rangle \cup \{c\}$.

Proof. Let $\alpha' = \text{INST}(\alpha, \alpha, c)$. We first prove that $\langle \alpha' \rangle \subseteq \langle \alpha \rangle \cup \{c\}$ by case analysis. If the contour of α is \clubsuit and $\alpha \notin \alpha$, then the contour of α' is c and so $\langle \alpha' \rangle = \{c\}$. If the contour of α is \clubsuit but $\alpha \in \alpha$, then the contour of α' is \clubsuit and so $\langle \alpha' \rangle = \emptyset$. Finally, if the contour of α is c' , then the contour of α' is also c' and so $\langle \alpha' \rangle = c' = \langle \alpha \rangle$. Thus, $\langle \alpha' \rangle \subseteq \langle \alpha \rangle \cup \{c\}$.

Because contour instantiation is largely naturally homomorphic, the remainder of this argument proceeds in parallel to Lemma A.39. \square

Lemma A.49. If $\text{REPL}(C, c) = C'$, then $\langle C' \rangle \subseteq \{c\} \cup \{c' \mid c' \in \langle C \rangle \wedge c' \not\prec c\}$.

Proof. Let $\alpha' = \text{REPL}(\alpha, c)$. We first prove that $\langle \alpha' \rangle \subseteq \{c\} \cup \{c' \mid c' \in \langle \alpha \rangle \wedge c' \not\prec c\}$. This is done by simple case analysis. If the contour of α is \clubsuit then contour replacement is a natural homomorphism and so $\alpha' = \alpha$; thus, $\langle \alpha' \rangle = \emptyset \subseteq \{c\} \cup \emptyset$. Otherwise, let the contour of α be c' and let the contour of α be c'' . If $c' \leq c$, then $c'' = c$; thus we have $\langle \alpha' \rangle = \{c\} \subseteq \{c\} \cup \emptyset$. If $c' \not\leq c$, then $c'' = c'$; thus we have $\langle \alpha' \rangle = \{c'\} \subseteq \{c\} \cup \{c'\}$.

Because contour replacement is largely naturally homomorphic, the remainder of this argument proceeds in parallel to Lemma A.39. \square

A.9 Closure Simulation

We now show our Preservation Lemma. Informally, this lemma says that, for any expression e simulated by a constraint set C , every legal small step corresponds to some legal type closure step such that simulation still holds over the new expression and closure set. This is formalized as follows:

Lemma A.50 (Preservation). If $e \preceq C$, C is well-formed, and $e \xrightarrow{1} e'$, then there exists some C' such that $C \xrightarrow{\text{cl}} C'$, $e' \preceq C'$, and C' is well-formed.

We now give a proof of Lemma A.50.

Proof. By case analysis on the operational semantics small step rule. Our strategy for each case is first to show that at least one type closure rule applies for each small step. We then show that the resulting expression is simulated by the resulting constraint set.

Integer Addition: If $e = E \llbracket [x = x_1 + x_2] \rrbracket e_{\text{REST}}$, then $e' = E \llbracket [x = n] \rrbracket e_{\text{REST}}$ for some n . We also have that $E \downarrow_{\text{int}}(x_1) = n_1$, that $E \downarrow_{\text{int}}(x_2) = n_2$, and that $n = n_1 + n_2$. Because addition is well-defined, we further have that both n_1 and n_2 are elements of \mathbb{Z} .

We now show that the Integer Addition closure rule applies in this case. Because $x = x_1 + x_2$ appears in e and because $e \preceq C$, Lemma A.22 gives us that there exists some $\alpha_1 + \alpha_2 <: \alpha$ in C such that $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, and $x \preceq \alpha$. Because $E \downarrow_{\text{int}}(x_1) = n_1$ and $x_1 \preceq \alpha_1$, Lemma A.33 gives us that $C \vdash \alpha_1 \xrightarrow{\text{int}} \tau_1$ and $n_1 \preceq \tau_1$; the same statement can be made regarding x_2, α_2, n_2 , and τ_2 .

Because we have shown that the Integer Addition closure rule applies, we know that $C \xrightarrow{\text{cl}} C'$ where $C' = C \cup \{\text{int} <: \alpha\}$. It remains to show that $e' \preceq C'$ and that C' is well-formed. We first observe that by Lemma A.19 we have $E \llbracket e_{\text{REST}} \rrbracket \preceq C$. By Lemma A.20 and $[x = n] \preceq \{\text{int} <: \alpha\}$ (which is immediate from the above and Definition A.7), we have that $e' \preceq C'$. To show that C' is well-formed, we observe that the only contour appearing in $\{\text{int} <: \alpha\}$ is that in α , which is a type variable appearing in C ; thus $\langle \{\text{int} <: \alpha\} \rangle \subseteq \langle C \rangle$. By Lemma A.18, we therefore have that C' is well-formed.

True Integer Equality: If $e = E \llbracket [x = x_1 == x_2] \rrbracket e_{\text{REST}}$, then $e' = E \llbracket [x' = ()], x = \text{'True } x' \rrbracket e$ for a fresh x' chosen such that $x' \leftrightarrow \langle \iota, 0, \dot{c} \rangle$ where $x \leftrightarrow \langle \iota, n, \dot{c} \rangle$. We also have that $E \downarrow_{\text{int}}(x_1) = n$ and that $E \downarrow_{\text{int}}(x_2) = n$ for some $n \in \mathbb{Z}$.

We now show that the Integer Equality closure rule applies in this case. Because $x = x_1 == x_2$ appears in e and because $e \preceq C$, Lemma A.22 gives us that there exists some $\alpha_1 == \alpha_2 <: \alpha$ in C such that $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, and $x \preceq \alpha$. Because $E \downarrow_{\text{int}}(x_1) = n_1$ and $x_1 \preceq \alpha_1$, Lemma A.33 gives us that $C \vdash \alpha_1 \xrightarrow{\text{int}} \tau_1$ and $n_1 \preceq \tau_1$; the same statement can be made regarding x_2, α_2, n_2 , and τ_2 .

Because we have shown that the Integer Equality closure rule applies, we know that $C \xrightarrow{\text{cl}} C'$ where $C' = C \cup \{\text{'True } \alpha' <: \alpha, \text{'False } \alpha' <: \alpha, () <: \alpha'\}$ where $\alpha' = \text{ETV}(\alpha)$. Because $x \preceq \alpha$, we know that $\alpha = \langle \iota, n, \dot{c} \rangle$ where $\dot{c}' = \clubsuit$ if $\dot{c} = \clubsuit$ and $\dot{c}' \leq \dot{c}$ otherwise. By Definition 5.6, we have that $\alpha' = \langle \iota, 0, \dot{c}' \rangle$; thus, $x' \preceq \alpha'$. By Definition A.7 and the above, we have that $[x' = ()], x = \text{'True } x' \preceq \{\text{'True } \alpha' <: \alpha, () <: \alpha'\}$; by Lemma A.20, we have that $[x' = ()], x = \text{'True } x' \preceq$

{‘True $\alpha' <: \alpha$, ‘False $\alpha' <: \alpha$, $() <: \alpha'$ }. By Lemma A.20, we therefore have that $e' \preceq C'$. It remains to show that C' is well-formed; this is demonstrated by Lemma A.18 as above.

A similar argument is made when $e' = E \parallel [x' = ()], x = \text{‘False } x'] \parallel e$.

Application: If $e = E \parallel [x = x_1 \ x_2] \parallel e_{\text{REST}}$, then $e' = E \parallel \text{EFRESH}(x, E_0 \parallel e''' \parallel [x' = r]) \parallel [x = \text{EFRESH}(x, x')] \parallel e_{\text{REST}}$ where $E \downarrow_{\text{fun}}^*(x_1) \xrightarrow{n} \phi \rightarrow e''$ and $x_2 \preceq_E \phi \rightarrow e'' \setminus E_0; e''' \parallel [x' = r]$.

We now show that the Application closure rule applies in this case. Because $x = x_1 \ x_2$ appears in e and because $e \preceq C$, Lemma A.22 gives us that there exists some $\alpha_1 \ \alpha_2 <: \alpha$ in C such that $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, and $x \preceq \alpha$. Because $E \downarrow_{\text{fun}}^*(x_1) \xrightarrow{n} \phi \rightarrow e''$ and $x_1 \preceq \alpha_1$ and by Lemma A.33, we have that there exists some $\bar{\tau}$ such that $\phi \rightarrow e'' \preceq \bar{\tau}$ and $C \vdash \alpha_1 \xrightarrow{\text{fun}} \bar{\tau}$. Because $x_2 \preceq_E \phi \rightarrow e'' \setminus E_0; e''' \parallel [x' = r]$, the previous, and by Lemma A.37, we have that there exist some α_{LAST} and C_{NEW} such that $\alpha_2 \preceq_C \bar{\tau} \setminus \alpha_{\text{LAST}}; C_{\text{NEW}}$ such that $E_0 \parallel e''' \preceq \langle \alpha_{\text{LAST}}, C_{\text{NEW}} \rangle$. Because C_{NEW} and INST are total functions, we have shown that the Application closure rule applies.

Let $e_{\text{NEW}} = E_0 \parallel e'''$ and let $x_{\text{LAST}} = r_{\text{LAST}}$ be the last element in e_{NEW} . Then $e_{\text{NEW}} \preceq C_{\text{NEW}}$ and $x_{\text{LAST}} \preceq \alpha_{\text{LAST}}$. Because $x \preceq \alpha$, we have by Definition A.7 that $[x = x_{\text{LAST}}] \preceq \langle \alpha_{\text{LAST}} <: \alpha \rangle$.

Let $c = C_{\text{NEW}}(\alpha_2, C)$, let $e'_{\text{NEW}} = \text{EFRESH}(e_{\text{NEW}}, c)$, and let $C'_{\text{NEW}} = \text{INST}(C_{\text{NEW}}, c)$. By Lemma A.40, we have $e'_{\text{NEW}} \preceq C'_{\text{NEW}}$. Let $\text{EFRESH}(x_{\text{LAST}}, c) = x'_{\text{LAST}}$. Let $C'' = \text{INST}(\langle \alpha_{\text{LAST}} <: \alpha \rangle, c)$. By Definition 5.10, we have that $C'' = \{\text{INST}(\alpha_{\text{LAST}}, \emptyset, c) <: \text{INST}(\alpha, \emptyset, c)\}$. The contour of α is not $*$ as it appears at top level in C ; thus, $C'' = \{\alpha'_{\text{LAST}} <: \alpha\}$ where $\alpha'_{\text{LAST}} = \text{INST}(\alpha_{\text{LAST}}, \emptyset, c)$. By Lemma A.40, we thus have that $[x = x'_{\text{LAST}}] \preceq C''$; thus by Lemma A.20 we have that $e'_{\text{NEW}} \parallel [x = x'_{\text{LAST}}] \preceq C'_{\text{NEW}} \cup C''$. It should be observed that the prior represents the expression which replaces the application at the call site in the operational semantics; the latter represents the set of constraints after instantiation in the Application Closure rule. We let $C_{\text{NEW}} \cup C'' = C'''$.

By Lemma A.19, we have that $E \parallel e_{\text{REST}} \preceq C$; thus by Lemma A.20 we have that $E \parallel \text{EFRESH}(x, E_0 \parallel e''' \parallel [x' = r]) \parallel [x = \text{EFRESH}(x, x')] \parallel e_{\text{REST}} \preceq C \cup C'''$. Finally, we let $C' = \text{REPL}(C \cup C''', c)$ and by Lemma A.42, we have $E \parallel \text{EFRESH}(x, E_0 \parallel e''' \parallel [x' = r]) \parallel [x = \text{EFRESH}(x, x')] \parallel e_{\text{REST}} \preceq C'$.

It remains to show that $\text{REPL}(C \cup C''', c)$ is well-formed. We know from premises that C is well-formed. Our strategy is to use the extractions of the contour sets involved in the application rule to show that each step in this process is also well-formed. Because α_1 appears in C , we have from

Lemma A.46 that $(\bar{\tau}) \subseteq (C)$. From this, because α_2 appears in C and by Lemma A.47, we have that $(C_{\text{NEW}}) \subseteq (C)$.

By Lemma A.48, we have that $(C'_{\text{NEW}}) \subseteq (C) \cup \{c\}$. We also have that $(C'') \subseteq (C) \cup \{c\}$. This gives us $(C'_{\text{NEW}} \cup C'') \subseteq (C) \cup \{c\}$ by Lemma A.11, which is equivalent to $(C''') \subseteq (C) \cup \{c\}$. Again by Lemma A.11, we have $(C \cup C''') \subseteq (C) \cup \{c\}$, which is equivalent to $(C') \subseteq (C) \cup \{c\}$. By Lemma A.49, we have that $(C') \subseteq \{c\} \cup \{c' \mid c' \in (C \cup C''') \wedge c' \not\leq c\}$. Because $(C \cup C''') \subseteq (C) \cup \{c\}$, this can be weakened to $(C') \subseteq \{c\} \cup \{c' \mid c' \in (C) \cup \{c\} \wedge c' \not\leq c\}$; since $c \leq c$, this simplifies to $(C') \subseteq \{c\} \cup \{c' \mid c' \in (C) \wedge c' \not\leq c\}$. For notational convenience, let $\mathcal{C} = \{c' \mid c' \in (C) \wedge c' \not\leq c\}$. It remains to show that (C') is well-formed; by Lemma A.18 it is sufficient to show that $\{c\} \cup \mathcal{C}$ is well-formed.

Since \mathcal{C} is a subset of (C) and C is well-formed, it is sufficient to show that $\forall c' \in \mathcal{C}. c' \not\leq c$. By Lemma A.45, we have that for every c' in (C) , we have either $c' \not\leq c$ or $c' \leq c$. By definition, \mathcal{C} contains only those c' such that $c' \not\leq c$; thus, we know that $\forall c' \in \mathcal{C}. c' \not\leq c$. Therefore, $\mathcal{C} \cup \{c\}$ is well-formed and so C' is well-formed. \square

To complete our simulation of closure, we show stuck expressions are modeled by simulation as inconsistencies in the constraint set.

Lemma A.51. If $e \preceq C$, there exists no e' such that $e \rightarrow^1 e'$, and e is not of the form E , then C is inconsistent.

Proof. We begin by observing that every e not of the form E must be factorable into some $E \parallel [x = r] \parallel e''$ where r is not of the form v . We proceed by case analysis on r and the conditions of Definition 3.5 which lead to a stuck e . In each case, we show that the corresponding constraint set is inconsistent.

If r is of the form $x_1 \text{ op } x_2$, we observe that e is only stuck when $E \downarrow_{\text{int}}(x_i)$ is undefined for $i = 1$ or $i = 2$. By Definition 3.2, this implies that $E \downarrow_{\text{int}}^*(x_i) = \square$. Because $[x = x_1 \text{ op } x_2]$ is in e , Lemma A.22 gives us that $\{\alpha_1 \text{ op } \alpha_2 <: \alpha\}$ for $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, and $x \preceq \alpha$. Because $x_i \preceq \alpha_i$, Lemma A.33 gives us that $C \vdash \alpha_i \xrightarrow{\text{int}} \square$. The presence of the constraint $\alpha_1 + \alpha_2 <: \alpha$ in C such that $C \vdash \alpha_i \xrightarrow{\text{int}} \square$ is inconsistent by Definition 4.8.

Otherwise, r is of the form $x_1 \ x_2$. Let $E \downarrow_{\text{fun}}^*(x_1) \xrightarrow{n} \phi_{\square} \rightarrow e'_{\square}$. In this case, e is only stuck if there exists no E_0 and e'' such that $x_2 \preceq_E \phi_{\square} \rightarrow e'_{\square} \setminus E_0; e''$. Inspection of Definition 3.4 reveals that this is only the case if for all $1 \leq j \leq n$ there is no E_0 such that $x_2 \preceq_E \phi_j \setminus E_0$. Inspection of Definition 3.3 reveals that, for each such ϕ_j , this is only the case under two conditions: either (1) $\phi_j = x_3 : \text{int}$ and $E \downarrow_{\text{int}}(x_2)$ is undefined, or (2) $\phi_j = x_3 : \text{lbl } x_4$ and $E \downarrow_{\text{lbl}}(x_2)$ is undefined. Thus, one of these two conditions must be true for each j in $\{1..n\}$.

Because $[x = x_1 \ x_2]$ is in e , Lemma A.22 gives us that $\alpha_1 \ \alpha_2 <: \alpha$ exists in C such that $x_1 \preceq \alpha_1$, $x_2 \preceq \alpha_2$, and $x \preceq \alpha$. Because $x_1 \preceq \alpha_1$, Lemma A.33 gives us that $C \vdash \alpha_1 \xrightarrow{\text{fun}} \forall \alpha_{\square} . \tau \phi_{\square} \rightarrow \alpha'_{\square} \setminus C_{\square}$ and, among other things, that $\phi_j \preceq \tau \phi_j$ for each j in $\{1..n\}$.

For each j in $\{1..n\}$, we now show that $\alpha_2 \preceq_C \tau \phi_j \setminus \star$. For such a j , one of the two conditions above holds. If $\phi_j = x_3 : \text{int}$, then $\phi_j \preceq \tau \phi_j$ gives us that, for some α_3 such that $x_3 \preceq \alpha_3$, $\tau \phi_j = \alpha_3 \sim \text{int}$. We also have that $E \downarrow_{\text{int}}(x_2)$ is undefined, which by Definition 3.2 gives us that $E \downarrow_{\text{int}}^*(x_2) = \square$. Because $x_2 \preceq \alpha_2$, Lemma A.33 gives us that $C \vdash \alpha_2 \xrightarrow{\text{int}} \square$; as a result, $\alpha_2 \preceq_C \tau \phi_j \setminus \star$.

Otherwise, $\phi_j = x_3 : \text{lbl } x_4$; in this case, $\phi_j \preceq \tau \phi_j$ gives us that, for some α_3 and α_4 such that $x_3 \preceq \alpha_3$ and $x_4 \preceq \alpha_4$, $\tau \phi_j = \alpha_3 \sim \text{lbl } \alpha_4$. We also have that $E \downarrow_{\text{lbl}}(x_2)$ is undefined, which by Definition 3.2 gives us that $E \downarrow_{\text{lbl}}^*(x_2) = \square$. Because $x_2 \preceq \alpha_2$, Lemma A.33 gives us that $C \vdash \alpha_2 \xrightarrow{\text{lbl}} \square$; as a result, $\alpha_2 \preceq_C \tau \phi_j \setminus \star$.

We have thus shown that for all j in $\{1..n\}$, $\alpha_2 \preceq_C \tau \phi_j \setminus \star$. Equivalently by notational sugar, we have shown that $\alpha_2 \not\prec_C \vec{\tau}$ given $C \vdash \alpha_1 \xrightarrow{\text{fun}} \vec{\tau}$ for a constraint $\alpha_1 \ \alpha_2 <: \alpha$ in C ; by Definition 4.8, C is inconsistent. \square

A.10 Proof of Soundness

We now prove Theorem 5.1. Our strategy for doing so is to show that simulation holds after the initial derivation and after each small step evaluation. Once small step evaluation is complete, we use this simulation to show that evaluation can only be stuck if the constraint set is inconsistent.

Proof. Given a closed e with unique variable bindings, we have by Lemma A.31 that, if $\llbracket e \rrbracket_E^{\emptyset} = \langle \alpha, C \rangle$, then C is well-formed and $e \preceq \langle \alpha, C \rangle$; we also have that every type variable in C has the non-contour \star . As a result, we have that $\langle C \rangle = \emptyset$. By Definition A.7, we have that $e \preceq C$. Let $e' = \text{EFRESH}(e, \{\square\})$ and let $C' = \text{INST}(C, \{\square\})$; by Lemma A.40, we have that $e' \preceq C'$. Observe that, by Definition A.5, e' is α -equivalent to e .

We next induct on the length of constraint closure. Let $e_0 = e'$ and let $C_0 = C'$. Further, let $e_0 \xrightarrow{1} \dots \xrightarrow{1} e_n$ for some $n \geq 0$. We start with the base case that $e_0 \preceq C_0$ and that C_0 is well-formed. We use Lemma A.50 to prove our inductive step: if $e_i \preceq C_i$ and C_i is well-formed, then there exists some $C_i \xrightarrow{\text{cl}_i} C_{i+1}$ such that $e_{i+1} \preceq C_{i+1}$ and C_{i+1} is well-formed. We therefore have by induction that $e_n \preceq C_n$ and that C_n is well-formed.

We have by premise that $e_n \not\rightarrow^1$ and that e_n is not of the form E . By Lemma A.51, we have that C_n is inconsistent. By Definition 5.11, e does *not* typecheck. \square

B. MicroBang Termination

We show here that the MicroBang typechecking definition presented in Section 5 is decidable. We first present a proof of decidability via a counting argument. We then informally

discuss how typechecking can be made computationally feasible.

B.1 Proof of Termination

We begin by showing each relation used in defining the type system is decidable, and each function is computable. We then show that the number of constraints which may appear in any closure is bounded in terms of the size of the original program and that the number of closure steps necessary to decide typechecking is finite.

Lemma B.1. Given a closed e , $\llbracket e \rrbracket_E^{\emptyset}$ is a computable function.

Proof. By direct induction on the definition in Figure 4.2. \square

Lemma B.2. Relation $\tau \prec_{<:\star} \alpha$ is decidable.

Proof. By enumeration of the possible chains $\tau <: \alpha_0, \dots, \alpha_{n-1} <: \alpha$. Each cyclic chain (a chain where the same type variable occurs more than once) has the same lower bounds as the acyclic chain obtained by pruning the cycle, and so it suffices to consider acyclic chains only. Because there are a fixed, finite number of constraints in C , there is a fixed number of acyclic chains which means the fixed set of all possible lower bounds can be computed and the relation holds iff τ is in that set. \square

The decidability of the projection relation and computability of a function which finds all $\vec{\tau}$ that can be projected are complex. This is due to the (rare in practice) case where onions can directly recurse without an intervening label. Since onioning is a form of type intersection, such recursions are non-contractive and are usually syntactically ruled out as they have no well-defined semantics [18], but we support non-contractive recursions. The following series of Definitions and Lemmas give us this decidability and computability result.

We begin by defining a relation similar to the projection shown in Definition 4.4 and then align this new definition with the old version to show the decision procedure. The definition below uses the notation $\epsilon_{\pi}(\alpha)$ to indicate that a type variable α is *nullable*: it may produce an empty list for the given projector π .

Definition B.3 (Nullable). $\epsilon_{\pi}(\alpha)$ if and only if either (1) there exists a non-onion τ such that $\tau \prec_{<:\star} \alpha$ and $\tau \not\sqsubseteq \pi$ or (2) $\alpha_1 \ \alpha_2 \prec_{<:\star} \alpha$ and inductively both $\epsilon_{\pi}(\alpha_1)$ and $\epsilon_{\pi}(\alpha_2)$ hold.

Lemma B.4. Relation $\epsilon_{\pi}(\alpha)$ is decidable.

Proof. By Lemma B.2, we have that concretization is decidable. There are finitely many $\&$ constraints in C and it suffices never to visit the same $\&$ constraint twice by analogy with the cycle pruning in Lemma B.2. \square

Definition B.5 (Modified Projection).

$$\begin{array}{ll}
C \vdash_0^\epsilon \alpha' \xrightarrow{\tau} [\tau] & \text{if } \tau \stackrel{C}{\prec} \alpha', \tau \sqsubseteq \pi \\
C \vdash_0^\epsilon \alpha' \xrightarrow{\tau_1} \parallel \tau_2 & \text{if } \alpha_1 \& \alpha_2 \stackrel{C}{\prec} \alpha', \\
& C \vdash_0^\epsilon \alpha_1 \xrightarrow{\tau_1}, C \vdash_0^\epsilon \alpha_2 \xrightarrow{\tau_2} \\
C \vdash_0^\epsilon \alpha' \xrightarrow{\tau_2} & \text{if } \alpha_1 \& \alpha_2 \stackrel{C}{\prec} \alpha', \epsilon_\pi(\alpha_1), \\
& C \vdash_0^\epsilon \alpha_2 \xrightarrow{\tau_2} \\
C \vdash_0^\epsilon \alpha' \xrightarrow{\tau_1} & \text{if } \alpha_1 \& \alpha_2 \stackrel{C}{\prec} \alpha', \epsilon_\pi(\alpha_2), \\
& C \vdash_0^\epsilon \alpha_1 \xrightarrow{\tau_1} \\
C \vdash^\epsilon \alpha \xrightarrow{\tau} & \text{if } C \vdash_0^\epsilon \alpha \xrightarrow{\tau} \text{ or } \epsilon_\pi(\alpha) \text{ and } \tau = \square
\end{array}$$

Lemma B.6. Relation $C \vdash \alpha \xrightarrow{\tau}$ is decidable.

Proof. We first show that the modified projection relation $C \vdash^\epsilon \alpha \xrightarrow{\tau}$ at the bottom of Definition B.5 is equivalent to the original projection relation $C \vdash \alpha \xrightarrow{\tau}$ in Definition 4.4. The modified definition $C \vdash_0^\epsilon \alpha \xrightarrow{\tau}$ at the top of Definition B.5 by inspection inlines the empty list clause (the first clause) of the original definition into the $\&$ clause, and that constructively produces a non-empty list in each case. Thus, the only case where $C \vdash \alpha \xrightarrow{\tau} / C \vdash_0^\epsilon \alpha \xrightarrow{\tau}$ differ is where the \square clause is the top-level result, and $C \vdash^\epsilon \alpha \xrightarrow{\tau}$ explicitly aligns that top-level behavior.

The modified relation $C \vdash^\epsilon \alpha \xrightarrow{\tau}$ is now shown to be decidable, and by the above equivalence, the original relation is decidable. Observe that the list τ is divided into two non-empty sublists in any use of the $\&$ clause, so it is possible to exhaustively test each of the n possible factorings of a length n list, and $\epsilon_\pi(\alpha)$ was shown decidable in Lemma B.4. At the leaf case, we rely on concretization and projection matching; the latter is trivially decidable and the prior was shown to be decidable via Lemma B.2. \square

Now we show that the set of all possible τ that project from a given type variable is computable. We show this by showing how the equivalent set of lists with duplicate types in each list removed, is computable. `DEDUP` is a function which removes all duplicate types in a list of types, preserving the rightmost occurrence in the list only. Formally:

Definition B.7 (Deduplication).

$$\begin{array}{ll}
\text{DEDUP}(\vec{\tau}) & = \text{DEDUP}(\emptyset, \vec{\tau}) \\
\text{DEDUP}(\vec{\tau}, \square) & = \square \\
\text{DEDUP}(\vec{\tau}, \vec{\tau} \parallel [\tau']) & = \begin{cases} \text{DEDUP}(\vec{\tau}, \vec{\tau}) & \text{when } \tau' \in \vec{\tau} \\ \text{DEDUP}(\vec{\tau} \cup \{\tau'\}, \vec{\tau} \parallel [\tau']) & \text{otherwise} \end{cases}
\end{array}$$

Lemma B.8. There is a computable function which takes a C , π , and α as argument and returns a set $\vec{\tau}$ such that for all $\vec{\tau}'$ we have $\text{DEDUP}(\vec{\tau}') \in \vec{\tau}$ if and only if $C \vdash \alpha \xrightarrow{\tau'}$.

Proof. The core of the algorithm is to perform non-deterministic disjunctive computation on each lower bound encountered (subject to some restrictions for termination). If we encounter a non-onion type under no other restrictions, computation is simple: we produce the singleton list containing that type.

If we encounter an onion type, we first explore the right side and then explore the left, concatenating the resulting type list sets pointwise. As we explore the left side, we maintain knowledge of the types which appeared in the right side's list and do not select proof subtrees which include them; this prevents generation of duplicates and is also key to termination: every non-onion expansion makes some form of progress.

For the case onions recurse into themselves without an intervening label, special care is needed. For instance, consider the constraint $\alpha_1 \& \alpha_2 \prec \alpha_2$. In order to prevent divergence in such a case, we track each type variable and *the number of times* we have expanded it to an onion lower bound in this proof tree branch. After this number of passes reaches a certain threshold, we replace all expansions of onion lower bounds on that variable with \square , pruning off that branch. We assert that the number of times a given onion is expanded can be conservatively be set to the number of distinct (onion and non-onion) lower-bounding types in C .

We first observe that this choice ensures decidability and then defend that it is sufficient for correctness. Because of this decision, there is now a maximum size to the proof tree; each variable with an onion lower bound is only expanded a fixed number of times and there are finitely many such variables. Because non-onion cases are always leaves and because each variable has finitely many lower bounds, the number of proof trees is finite and can be exhaustively checked.

We next assert that it is sufficient to build any proof tree described by the recursive onions to be large enough to hold any legal permutation of the lower-bounding types in C . This is because additional onion structure exploration will never yield more legal variable positions, there are only finitely many types that can be placed in any position. \square

Lemma B.9. The function which given some C , α , and π returns the set of all τ in C such that $C \vdash \alpha \xrightarrow{\tau}$ is a computable function.

Proof. By Lemma B.8, there is a computable function which produces the set of deduplicated lists such that the original lists were the result of projection. Because deduplication never removes the rightmost element of a list (trivial by inspection of Definition B.7), the set of rightmost types from the lists is equal to the set of rightmost types from the lists produced by projection. \square

Lemma B.10. Relation $\alpha \preceq_C \tau \setminus \dot{C}$ is decidable.

Proof. Inspection of Definition 4.6 reveals that compatibility is syntax-directed with leaf cases being either axiomatic or relying on single projection. By Lemma B.9, we have that single projection is decidable. \square

Lemma B.11. Relation $\alpha \preceq_C \vec{\tau} \setminus \alpha'; C'$ is decidable.

Proof. By induction on the length of $\bar{\tau}$. Inspection of Definition 4.7 reveals that the truth of this term is either axiomatic (when the length of $\bar{\tau}$ is zero), reliant upon compatibility (shown to be decidable in Lemma B.10), or defined in terms of a smaller list. \square

Lemma B.12. Relations $c_1 \leq c_2$ and $c_1 \not\leq c_2$ are decidable.

Proof. The grammar and meaning of contours is a subset of regular expressions and subsumption and overlap on regular expressions is well-known to be computable. \square

Lemma B.13. The contour extraction ($(\cdot)_\#$), instantiation (INST), and replacement (REPL) are decidable.

Proof. These functions are trivially computable by induction on their respective first arguments. Each step not defined as a natural homomorphism is trivially computable. For instantiation, the scape case calculates a finite union and the variable case performs a presence test on a finite set. For replacement, the variable case performs subsumption testing (decidable by Lemma B.12). The leaf cases of extraction are constant. \square

Lemma B.14. The functions used in contour creation – C_{NEW} , C_{COLLAPSE} , and C_{WIDEN} – are decidable.

Proof. C_{NEW} is trivially decidable. $C_{\text{COLLAPSE}}(c)$ is trivially decidable if each $C_{\text{COLLAPSE}}(s)$ is decidable for each $s \in c$. $C_{\text{COLLAPSE}}(s)$ is computable given some s by induction on the number of cycles in s . $C_{\text{WIDEN}}(c, C)$ is decidable given c and C because contour extraction and overlap are known to be decidable from Lemmas B.12 and B.13. \square

Lemma B.15. Given C and C' , $C \xrightarrow{\text{CL}} C'$ is decidable.

Proof. We proceed by showing that there is a computable function from C to C'' such that $C \xrightarrow{\text{CL}} C''_i$ for all $i \in \{1..m\}$ and that there exists no other C''' such that $C \xrightarrow{\text{CL}} C'''$; that is, the set of possible next steps is computable and finite. Inspection of the closure rules reveals that the first premise of each rule tests for the presence of a constraint in C of a certain form; because C is finite, we can enumerate all such constraints. For each such constraint, the remainder of the premises in each rule are computable by one of the above lemmas. For application, we use Lemma B.8 to produce the set of possible $\bar{\tau}$ scapes that may apply; this Lemma produces a list without duplicates, and it is clear that the presence of duplicates is irrelevant to the result of application compatibility, Definition 5.7. Because each rule has finitely many preconditions, this process is computable. Thus we can compute the set of constraints C'' ; because each rule only adds finitely many constraints, each C''_i is finite. This relation is then determined simply by $C' \in C''$. \square

Lemma B.16. It is decidable whether a given C is inconsistent.

Proof. By inspection of Definition 4.8, each condition of inconsistency first checks for a constraint of a given form; this check can be achieved by enumeration because C is finite. For each such constraint, it then performs a series of checks known to be decidable from the above lemmas. \square

We now prove Theorem 5.2.

Proof. The proof proceeds by showing that each closure rule can only introduce constraints drawn from an initial fixed finite set of possible constraints.

Given a closed e , computing $\llbracket e \rrbracket_E^\emptyset = \langle \alpha, C \rangle$ is decidable from Lemma B.1.

Clearly a fixed set of identifiers ι exist in e (and thus in C) and no type system rule introduces new identifiers. Additionally, new contours are only created by C_{NEW} . This function guarantees that (1) all contour strands contain a given identifier only once and that (2) no contour strands contain an empty set as a part. As a result, there is a fixed, finite number of unique contour strands that could ever occur in any closure sequence of C and thus there is a fixed bound on the number of distinct contours possible in any closure. Because the initial expression e is fixed, there is also a fixed number of variable indices that could be used in a given closure sequence. These facts together give us that, given a C , we can define a fixed, finite set of type variables which is a superset of any variable that could arise during any closure sequence.

Since the possible contours in any closure sequence are fixed in advance, there a fixed finite set can be defined which bounds all of the τ that could occur in any constraint over any closure sequence: by inspection of the closure rules, no new non-scape types τ are created in new constraints which are not just substitutions on types already occurring in C . The number of possible patterns τ_ϕ is also fixed because new constraints added in closure contain no new patterns. Finally, function types contain a constraint set and all substitutions on such a set are fixed and finite since the number of possible substitutions is bounded by the above and so the number of constraints that could be added to closure via application have a fixed bound.

Since given initial e the possible type variables and types have a fixed bounded size, there is also a fixed, finite set bounding the constraints which may appear in any closure; let C^* be that set. Therefore, any C' such that $C \xrightarrow{\text{CL}} C'$ must be in the power set of C^* ; because C^* is finite, its power set is also finite. Let this power set be denoted 2^{C^*} and let n be the size of this set.

The set C' such that for all $i \in \{1..m\}$ we have $C \xrightarrow{\text{CL}} C'_i$ is also decidable. We know this set to be a subset of 2^{C^*} and that any acyclic chain of closures $C \xrightarrow{\text{CL}} C'_{j_1} \xrightarrow{\text{CL}} \dots \xrightarrow{\text{CL}} C'_i$ (for $i \in \{1..m\}$ and j_k all in $\{1..n\}$) has at most n steps (since there are only n such sets in 2^{C^*}). Because the relation holds over the first and last sets and is not based on the intervening sets, any cyclic chain of closures need

not be considered as it is equivalent to some acyclic chain. Because these acyclic chains are all composed of elements of 2^{C^*} , the number of such acyclic chains is at most the sum of the permutations of the subsets of 2^{C^*} , which is finite. By enumeration of these permutations, we can exhaustively determine which elements of 2^{C^*} are in C' and which are not; therefore, C' is decidable.

Finally, for each such C'_i , Lemma B.16 gives us that it is decidable whether or not that set is inconsistent. Since it is decidable whether any of these sets is inconsistent, it is decidable whether or not e typechecks. \square

B.2 Efficient Typechecking

The proof of Theorem 5.2 in the previous section used a counting argument for sake of simplicity. While this shows that the MicroBang type system is decidable, it does not show that it is computationally feasible. Because our objective is to produce a usable, highly flexible scripting language, we will now informally discuss how the complexity of typechecking is tractable.

Constraint Closure Confluence The first reduction in complexity starts with the way typechecking is phrased. A program e is typesafe only if, given its initial constraint C , every C' for which $C \xrightarrow{C^*} C'$ is consistent; computing every such constraint set would require considerable effort. This could be avoided if constraint closure were confluent; then, the closure work would be reduced to a computable, deterministic function. The constraint closure relation is not trivially confluent; the contours which it produces will vary based upon the order in which the constraint closure rules are used. This is, in particular, due to the unioning of contours in the widening step of contour creation. Up to contour meaning ($\llbracket \cdot \rrbracket$), however, these contours are equivalent. It is therefore sufficient to close over the initial constraint set to any of its fixed points and simply determine if that constraint set is inconsistent.

Projection Lemma B.8 shows that the set of orderings of scapes which will project from a type variable is computable, but it uses an excessively complex algorithm. Algorithms for projecting the single highest priority element from a type variable are quite simple, but – due to onion types which recurse directly and not under a label – the task of projecting multiple types in priority order is difficult in some corner cases. This is because immediately recursive onions are a form of *non-contractive* type because $\&$ shares enough similarity with intersection type, and intersection is not a contractive type operator [18]. It is well known that working with non-contractive recursive types is difficult [18] and, for that reason, it is standard to work only with contractive types. While eliminating non-contractive types from TinyBang would be as simple as a syntactic check prohibiting immediately recursive onion types, we believe that non-

contractive types will make it possible to statically type certain higher-order object programming patterns.

There are only two places that our priority-based projection is used. The first is in the definition of single projection which, as mentioned above, can be implemented with simple, efficient algorithms. The second is in the Application closure rule, where it is used to determine the order in which scapes are applied. Observe that the use of priority-based projection here is inefficient; the projection tries to model each and every concrete type tree and individually determine the priority list which results from it, and there are many more such type trees than there are priority lists in every case. Furthermore, the set of priority-ordered lists contain significant redundancies, since the ordering only matters in cases where the types matched by scape patterns overlap.

In order to correct both sources of inefficiency, we plan to specialize an operation to address the behavior of the application rule over an onion of scapes. We rely on the fact that all types expressible in our constraint-based system can be expressed in a regular tree algebra and we can use operations over regular trees to calculate how much of the argument type is matched at any given point during the process. By defining the operation to handle all lower bounds at a given call site simultaneously, we effectively group the application cases into equivalence classes and process each equivalence class only once. This approach also allows us to consider each recursive onion type only once; because it considers all lower bounds of the recursive type variable simultaneously, we know that no further information can be gained by exploring it again.

Contours The only remaining concern regarding the complexity of the TinyBang type system is the generation of type contours. As with any polymorphic system, the polyinstantiation of variables gains expressiveness at the cost of complexity and, as with any interesting polymorphic type system (such as those in the ML family of languages), typechecking TinyBang is exponentially complex in pathological scenarios. We have attempted to ensure, however, that TinyBang typechecking will be polynomial in practice.

The manner in which contours are created, for instance, is intended to prevent exponential complexity. Since contours on type variables represent (often infinite) sets of calling contexts, it is not impossible to imagine that a closure rule may fire for some intersection of two variables' calling contexts. In that case, constraints would need to be constructed over the intersection of these calling contexts and propagated appropriately. This deep reasoning about type propagation is only helpful in very specific cases of overlapping recursive call cycles where the relevant input types are statically known; such reasoning also generates exponentially many constraints in the size of the program. Because this precision is costly and not very useful, we choose to completely unify the variables involved in any call cycle; this is accomplished at the end of the contour creation process by

the `WIDEN` function, which ensures that any new contour will be unified with a contour it overlaps (even partially). This ensures that each contour in constraint closure represents a disjoint calling context, significantly reducing the complexity of closure.

In general, we expect that a polynomial-in-practice implementation of the `TinyBang` typechecker is feasible, but we concede that it will take more theoretical effort than other, more traditional typecheckers due to subtle issues like those mentioned above.