

A foundation for trait-based metaprogramming

John Reppy Aaron Turon

University of Chicago

{jhr, adrassi}@cs.uchicago.edu

Abstract

Schärli *et al.* introduced traits as reusable units of behavior independent of the inheritance hierarchy. Despite their relative simplicity, traits offer a surprisingly rich calculus. Trait calculi typically include operations for resolving conflicts when composing two traits. In the existing work on traits, these operations (method exclusion and aliasing) are *shallow*, *i.e.*, they have no effect on the body of the other methods in the trait. In this paper, we present a new trait system, based on the Fisher-Reppy trait calculus, that adds *deep* operations (method hiding and renaming) to support conflict resolution. The proposed operations are deep in the sense that they preserve any existing connections between the affected method and the other methods of the trait. Our system uses Riecke-Stone dictionaries to support these features. In addition, we define a more fine-grained mechanism for tracking trait types than in previous systems. The resulting calculus is more flexible and expressive, and can serve as the foundation for *trait-based metaprogramming*, an idiom we introduce. A companion technical report proves type soundness for our system; we state the key results in this paper.

1. Introduction

A *trait* is a simple collection of methods that represent the partial implementation of a class. Methods defined in a trait are *provided* methods; any method referenced in a trait but not provided by it is a *required* method. Traits are similar to abstract classes, but with two important differences: traits cannot introduce state, and they do not lie within the inheritance hierarchy. The primary mechanism of class reuse is inheritance, which is an asymmetric operation. With traits, reuse takes place via *composition*, a symmetric concatenation of two traits. In addition to composition, trait calculi include fine-grained operations for manipulating traits as method suites. Ultimately, traits are *inlined* into classes during class formation. Unfulfilled trait requirements must be provided by the class at the time of inlining.

Traits were introduced by Schärli *et al.* in the setting of SMALLTALK [SDNB03]. A companion paper explored the use of traits to refactor the SMALLTALK collection classes, with encouraging results: a 10% reduction in method count for a 12% reduction in overall code size [BSD03]. Since SMALLTALK is untyped, the original work on traits did not include a type system. Fisher and Reppy gave a calculus for statically-typed traits [FR04, FR03]. Other work subsequently developed typed trait calculi for JAVA [LS04, SD05]. Quitslund performed a simple analysis of the Swing Java library, which suggests that code reuse can also be improved by adding traits to Java [Qui04]. A fair amount of activity has followed, with trait implementations underway or completed for C#, JAVA, PERL, and SCALA.¹

¹See <http://www.iam.unibe.ch/~scg/Research/Traits/> for more information.

We present a calculus, based on the Fisher-Reppy polymorphic trait calculus [FR03], with support for trait privacy, hiding and deep renaming of trait methods, and a more granular trait typing. Our calculus is more expressive (it provides new forms of conflict-resolution) and more flexible (it allows after-the-fact renaming) than the previous work. Traits provide a useful mechanism for sharing code between otherwise unrelated classes. By adding deep renaming, our trait calculus supports sharing code between *methods*. For example, the JAVA notion of synchronized methods can be implemented as a trait in our system and can be applied to multiple methods in the same class to produce synchronized versions. We term this new use of traits *trait-based metaprogramming*.

We review the standard trait operations in Section 2 and describe our additions in Section 3; the presentation is organized around a series of examples, culminating with the introduction of our metaprogramming idiom. The formal description of our proposal begins with Section 4, which outlines our notation and the syntax of the calculus. Section 5 describes the static semantics. The system types traits at the time of trait formation, rather than trait inlining. It improves on previous work by tracking trait requirements at a per-method, rather than per-trait basis, which allows spurious requirements to be dropped as a trait is manipulated. Properly handling evaluation in the presence of privacy can be somewhat subtle. We utilize Riecke-Stone dictionaries [RS02] to handle privacy; this technique also provides support for renaming. Our approach is detailed in Section 6. Finally, we outline the proof of type soundness in Section 7. After the formal description, we take stock of the calculus in a broader context. Section 8 details other work in traits and relates our presentation of traits to other constructs, *e.g.* mixins and aspects. The calculus raises some interesting questions for language design. We examine these questions and conclude in Section 9.

A technical report version of this paper includes the complete formal model and a detailed proof of type soundness [RT06].

2. Traits

Traits originate from the observation that classes serve two often conflicting purposes [SDNB03]. From one perspective, classes are meant to generate objects, which means they must be complete and monolithic. At the same time, classes act as units of reuse via inheritance, and from this perspective they should be small, fine-grained, and possibly incomplete. Inheritance must straddle these two roles, which often forces the designer of a class hierarchy to choose between interface cleanliness and implementation cleanliness.

Consider the case of two classes in different subtrees of the inheritance hierarchy which both implement some common protocol. To avoid code duplication, the protocol implementation should be shared between the two classes. In a single inheritance framework, the common code can be lifted to a shared superclass, but doing so pollutes the interface of the superclass, affecting all of its subclasses. With multiple inheritance, the protocol implementation can reside in a new superclass that is inherited along with the existing

superclasses, but multiple inheritance complicates the implementation of subclasses (e.g., with respect to instance variables).

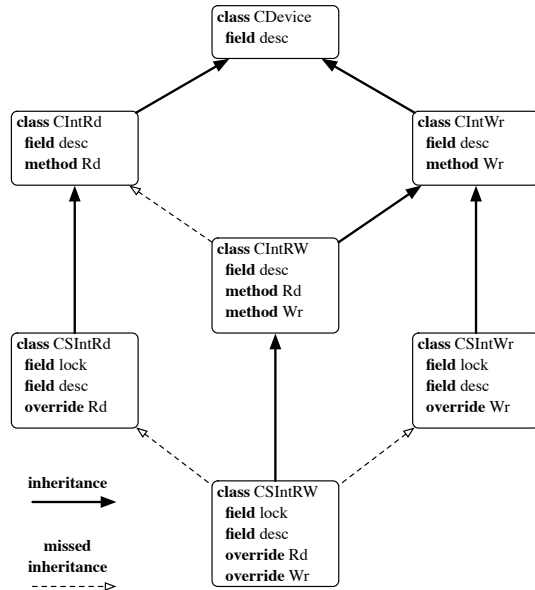


Figure 1. Synchronized readers and writers in a single inheritance framework

Figure 1 illustrates these issues with a concrete class hierarchy in a single inheritance framework. The root of the hierarchy is the `CDevice`² class, which implements I/O on a file descriptor. It has two subclasses for, respectively, reading and writing integers on the device. Defining a class that supports both reading and writing (`CIntRW`) requires reimplementing one of the methods (denoted by the dashed arrow in the figure). While we could lift the `Rd` and `Wr` methods to the `CDevice` class, doing so would pollute the interface of the original `CIntRd` and `CIntWr` subclasses as well as new subclasses such as boolean readers and writers. The class hierarchy is further extended with support for synchronized reading and writing by adding a lock. Single inheritance again forces a reimplementations of methods.

In contrast to inheritance, which specifies the relationship between a family of classes, traits allow the implementation of a single class to be factored into multiple, structured parts. Classes are retained as hierarchically organized object generators, but traits are introduced as flatly organized, partial class implementations. Traits thus assist in separating the roles distinguished above. Figure 2 shows how traits can allow code reuse without having to define methods too high in the hierarchy; in this example, four traits are used to generate six classes. Note that the traits in the example have field requirements, not just method requirements.³ The **override** annotation on provided methods signifies that the method invokes the super-method with the same name.

Operations on traits

Traits are partial class implementations, but they are restricted to providing only a simple set of methods. In particular, traits cannot introduce state. The methods of a trait are only loosely coupled: they can be freely removed and replaced by other implementations.

²For concrete names, we prefix classes with C and traits with T.

³Field requirements were introduced in [FR04].

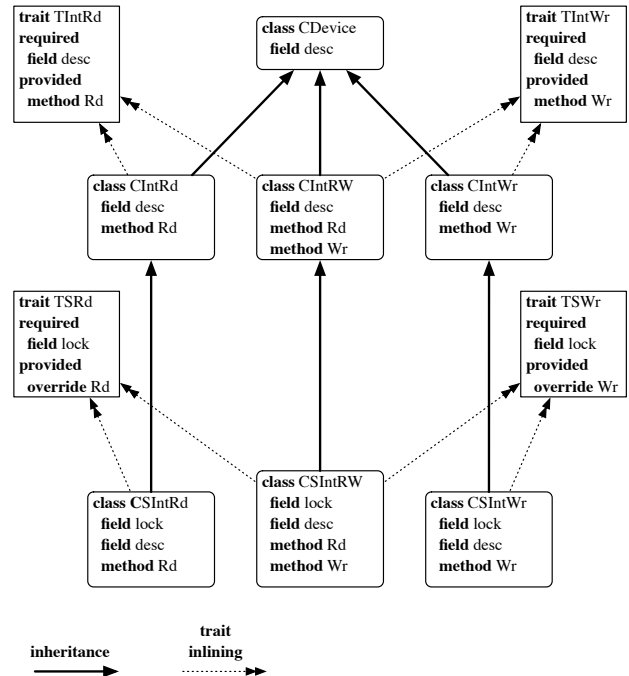


Figure 2. Using traits to implement synchronized readers and writers

These properties make traits more nimble and lighter-weight than either multiple inheritance [Str94] or mixins [BC90, FKF98].

Two traits can be combined via *trait composition*, written $T_1 + T_2$. The resulting trait is a flat merger of the operands. Composing two traits may fulfill method requirements for one or both of the traits. In a statically-typed setting, composition requires its operands to be disjoint; method conflicts must be resolved explicitly, using other trait operations. Under this definition, trait composition is commutative, obviating the need for a linear ordering found with single inheritance mixins.

Trait composition is a symmetric alternative to reuse via inheritance. By defining additional operations on traits, one can capture more complex idioms of reuse. The typical suite of trait operations include *alias* and *exclude*, which are useful for conflict resolution. As an illustration, assume that we have two traits, `TPoint` and `TColored`, which both provide a `toString` method, and that we wish to compose them into a single trait, `TCPoint`. In the simplest case, we can resolve such a conflict by choosing one method or the other. We exclude the unwanted method implementation, and compose the resulting trait:

```
TCPoint = TPoint + (TColored exclude toString)
```

`TCPoint` will provide a single `toString` method, using the implementation from `TPoint`. In addition, any invocations of `toString` from within other methods provided by `TColored` will now invoke `TPoint`'s implementation, making the `exclude`-compose combination similar to method override.

Sometimes it is useful to retain both implementations of a conflicting method, perhaps providing a new implementation combining the two. The *alias* operation creates a new name for an existing method:

```

TCPoint = {
  provides toString() : string {
    self.strP() + ": " + self.strC();
  }
} + ((TPoint alias toString as strP)
  exclude toString)
+ ((TColored alias toString as strC)
  exclude toString)

```

It is important to realize that, while `strP` and `strC` are available in this version of `TCPoint`, the original `toString` methods have effectively been overridden. Invocations of `toString` from the other methods provided by `TPoint` and `TColored` will use the new, combined `toString` implementation provided by `TCPoint`. The combination of aliasing and excluding thus yields a *shallow* renaming: existing references to an aliased method continue to refer to the original method name.

3. Traits with hiding and deep renaming

We extend the trait system of the previous section with two new trait operations, *hide* and *rename*, which act as the deep variants of *exclude* and *alias*. These two operations provide new forms of conflict resolution when composing traits. Each is also useful in isolation: hiding without composition yields trait privacy, while renaming yields an idiom we term *trait-based metaprogramming*, the most exciting aspect of our work.

The *hide* operation permanently binds a provided method to a trait, while hiding the method’s name. A new method with the same name can be introduced as a new provided or required method of the trait, but existing references to the method from other provided methods are statically bound to its implementation at the time of hiding. Returning to the `TCPoint` example, we can write

```
TCPoint = TPoint + (TColored hide toString)
```

Here, `TCPoint` will provide the `toString` implementation from `TPoint`. Unlike the exclusion example, any references from within `TColored` to the `toString` method will continue to refer to `TColored`’s implementation. In effect, `TColored`’s `toString` implementation is provided by `TCPoint`, but in a nameless and inaccessible form. Where combining exclusion and composition leads to overriding, the combination of hiding and composition yields *shadowing*.

Method hiding can also be used to hide trait implementation details — *i.e.*, as a form of trait privacy. A trait implements some collection of behavior, and it may need to make use of new methods that are specific to its implementation but are not appropriate to provide publicly. Although such “helper methods” could be made private after they are inlined into a class, we believe that traits should not only factor out, but also encapsulate, units of behavior. The importance of trait privacy will depend to some extent on the strength of the surrounding language features. In a language with a powerful module system, for example, it is likely that trait privacy could be achieved through signature ascription instead [FR99]. In any case, trait privacy is a free by-product of introducing method hiding for conflict resolution.

Loosening the connection between method names and method implementations suggests the possibility of a deep renaming operation, as opposed to shallow renaming with *alias-exclude*. As with the other operations, pairing renaming with composition provides a new form of conflict resolution:

```
TCPoint = (TPoint rename toString to strP)
+ (TColored rename toString to strC)
```

This version of `TCPoint` does not provide a `toString` method at all. Existing references to `toString` from within `TPoint` and `TColored` now refer to `strP` and `strC`, respectively; this is the sense in which the renaming is “deep.” `TCPoint` can now be

extended with another implementation of `toString`, even one with a different type.

Of course, renaming can be used alone in order to align a “misnamed” provided method with an existing class hierarchy or signature constraint. Required methods may be renamed for the same reason. Furthermore, super-method names may be renamed, so that for example all the invocations of `super.foo` within a trait may be renamed to invoke `super.bar`. The rationale for this last form of renaming will become more clear in the following discussion.

Trait-based metaprogramming

Schärli *et al.* summarized their system with the following equation [SDNB03]:

$$\text{Class} = \text{Superclass} + \text{State} + \text{Traits} + \text{Glue}$$

A trait represents a flat fragment of a class’s behavior, and class behavior is closely tied to naming. But a trait can also be seen as a collection of named provided methods *parameterized* over a collection of named required methods (and fields). Thus, traits capture a relationship between two families of named methods. Often this relationship carries as much information as the names of the methods themselves.

Reconsider the example of traits usage presented in Figure 2. The trait `TIntRd` requires a field `desc` and provides a method `Rd`. In a broad sense, it does not matter what we call these entities (`Rd`, `read`, `Read`, *etc.*), so long as (1) the names convey “descriptor” and “read,” respectively, and (2) the relationship between the entities remains intact. After-the-fact renaming allows the names to be changed as needed.

But what can we say about the traits `TSRd` and `TSWr` from the same example? Each requires the field `lock` and provides a single method that wraps a super-method invocation with synchronization code. For these traits, the relationship between the provided methods and the trait’s requirements carries essentially *all* of the relevant information; the provided method names `Rd` and `Wr` are incidental. We can capture the relationship in a single trait, `TSync`, which is parameterized by type:

```

trait TSync<ty1, ty2> = {
  provides Op(x : ty1) : ty2 {
    self.lock.Acquire();
    super.Op(x) before
    self.lock.Release();
  }
  requires field lock : LockObj
}

```

The generic method `Op` and its associated super-send can be renamed and used to override a method with a synchronized wrapper. Returning to the readers and writers example, we could rename `Op` to `Rd` and `super.Op` to `super.Rd`, and likewise for `Wr`. The key insight is that `TSync` can be instantiated several times, once for each method we want to synchronize. Figure 3 shows the readers/writers example using `TSync`; we now have six classes derived from three traits. It is worth noting that, with the ability to rename field requirements, different instantiations of `TSync` could use different locks. Our model does not support field renaming, but this is for simplicity only; the feature would be a fairly straightforward addition.

Traits were intended to capture specific, named behavior at the class level. Examples like `TSync` shift the focus to behavior at the method family level. With the latter perspective it becomes sensible to inline a trait several times into a class; we label this idiom *trait-based metaprogramming*. The power of the technique is that we can freely mix flat requirements for the class (*e.g.*, `lock`) with provided methods that may be instantiated several times. Traits

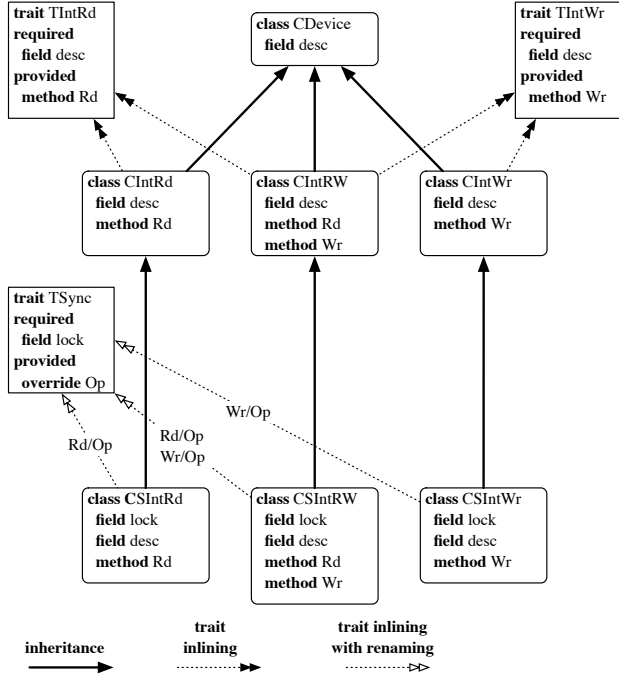


Figure 3. Using traits with renaming to implement synchronized readers and writers

provide the structure for specifying method relationships, and renaming provides the flexibility to apply traits in multiple contexts within a single class.

The metaprogramming technique we have outlined is somewhat *ad hoc*: examples like `TSync` would be best served by an explicit notion of method name abstraction. A good starting point would be a “lambda calculus of traits,” with method renaming playing the role of substitution. It is easy to imagine further extensions. A sophisticated trait-metalanguage could codify certain patterns of trait use, perhaps providing a mechanism like `JAVA`’s **synchronized** keyword for appropriately instantiating traits like `TSync`. Borrowing a line from aspect-oriented programming, a language might also allow *join points* to be specified for trait application. Such possibilities are exciting and deserve to be explored, but before we can define abstraction mechanisms for higher-level calculi, we need a well-defined notion of substitution.

Starting with the next section, we give a detailed semantics of a calculus with hiding and renaming for traits, which can serve as the foundation for trait-based metaprogramming.

4. A trait calculus with hiding and renaming

To model our proposed features, we have developed a statically-typed trait calculus loosely based on Fisher and Reppy’s [FR03]. The calculus is meant to give a rigorous semantics of the new trait operations, and does not represent a surface language design. We model the trait language in considerable detail, but restrict the rest of the language to essential features. Specifically, our class definition form does not support definition of methods directly. Instead, we leave this responsibility to traits. This choice is purely to minimize the complexity of the formal system; method definition in classes could be easily added.

The syntax for our calculus is given in Figure 4. To keep the syntax lightweight, we separate names into disjoint “universes,”

$P ::= D; P \mid e$ $D ::= t = (\vec{\alpha})T$ $\quad \mid c = C$ $\quad \mid x = e$ $T ::= t(\vec{\tau})$ $\quad \langle \mu; \rho \rangle$ $\quad T_1 + T_2$ $\quad T \text{ exclude } m$ $\quad T \text{ hide } m$ $\quad T \text{ alias } m \text{ as } m'$ $\quad T \text{ rename } r \text{ to } r'$ $\mu ::= m(x : \tau_1) : \tau_2 \{e\}$ $\rho ::= \langle \langle r : \tau_r \mid r \in \mathcal{R} \rangle \rangle$ $\theta ::= \langle \langle m : \tau_m @ \rho_m \mid m \in \mathcal{M} \rangle \rangle$ $\quad \mid \Lambda(\vec{\alpha}).\theta$ $C ::= c$ $\quad \mid \text{nil}$ $\quad \mid I \text{ in } T \text{ extends } C$ $I ::= \lambda(x : \tau).(\text{super } e_1) \oplus e_2$ $\chi ::= \tau \rightarrow \{ \{ l : \tau_l \mid l \in \mathcal{L} \} \}$ $e ::= x$ $\quad \lambda(x : \tau).e$ $\quad e_1 e_2$ $\quad \text{new } C e$ $\quad \text{self}$ $\quad \text{super}.m$ $\quad e.m$ $\quad e.f$ $\quad e_1.f := e_2$ $\quad \{f = e_f \mid f \in \mathcal{F}\}$ $\quad e_1 \oplus e_2$ $\quad ()$ $\tau ::= \alpha$ $\quad \langle l : \tau_l \mid l \in \mathcal{L} \rangle$ $\quad \tau_1 \rightarrow \tau_2$ $\quad \{f : \tau_f \mid f \in \mathcal{F}\}$ $\quad \text{unit}$	<p>program</p> <p>trait declaration</p> <p>class declaration</p> <p>expression declaration</p> <p>polymorphic trait name</p> <p>trait formation</p> <p>trait composition</p> <p>method exclusion</p> <p>method hiding</p> <p>method aliasing</p> <p>method renaming</p> <p>method</p> <p>inlining assumptions</p> <p>trait type</p> <p>polymorphic trait type</p> <p>class name</p> <p>empty class</p> <p>subclass formation</p> <p>constructor</p> <p>class type</p> <p>variable</p> <p>function abstraction</p> <p>function application</p> <p>object instantiation</p> <p>host object</p> <p>super-method dispatch</p> <p>method dispatch</p> <p>field selection</p> <p>field update</p> <p>field record</p> <p>field concatenation</p> <p>unit value</p> <p>type variable</p> <p>object type</p> <p>function type</p> <p>field record type</p> <p>unit type</p>
---	---

Figure 4. Trait calculus syntax

	Items	\in	Sets	\subset	Universe
Field names	f	\in	\mathcal{F}	\subset	\mathcal{F}_U
Method names	m	\in	\mathcal{M}	\subset	\mathcal{M}_U
Labels	l	\in	\mathcal{L}	\subset	$\mathcal{L}_U = \mathcal{F}_U \cup \mathcal{M}_U$
Super methods	$\text{super}.m$	\in	\mathcal{S}	\subset	\mathcal{S}_U
Requirements	r	\in	\mathcal{R}	\subset	$\mathcal{R}_U = \mathcal{L}_U \cup \mathcal{S}_U$
Slots	i	\in	\mathcal{I}	\subset	\mathcal{I}_U
Variables	x	\in			VARS
Type vars	α	\in			TYVARS
Trait names	t	\in			TRAITNAMES
Class names	c	\in			CLASSNAMES

Figure 5. Naming conventions

as shown in Figure 5. This convention allows entities to be distinguished by the “type” of their name, without any additional book-keeping.

In our model, a program is a series of zero or more declarations followed by an expression. The calculus is organized into three components: a trait language, a class language, and an expression language. The trait language includes expressions forms for all of the features we have discussed: composition of traits, and exclusion, hiding, aliasing, and renaming of methods. The declaration form for traits allows parameterization by type variables. Traits are initially formed with a single method rather than a method suite, which simplifies the semantics; a trait with multiple methods can be constructed via repeated composition.

The class language provides a simple model of single inheritance with no visibility control. The calculus has an empty base class, `nil`; all other classes must be defined via inheritance. Subclass formation takes a super class, a trait expression, and a constructor. In our model, the only methods introduced in a subclass are those provided by the trait; there is no separate notion of method definition for classes. In particular, this means that any required methods of the trait must be provided by the super class. We will sometimes refer to subclass formation as *trait inlining* to emphasize the role that traits play. When the trait being inlined is the result of simple trait formation and concatenation, trait inlining reduces to standard single inheritance.

Subclass formation also allows the introduction of state via a constructor function. State is restricted to *field records* which can be concatenated using the \oplus operator. Class constructors are syntactically constrained to apply the super class constructor to an expression, and concatenate the result with a new field record. Fields referenced in an inlined trait may originate from either the super class or the newly formed subclass.

The expression language is a simple object calculus with first class functions. It has imperative features (object instantiation and field update) which allow us to put the class language to work in a realistic way.

We will need several notational tools: we write $A \uparrow B$ for $A \cap B = \emptyset$; when f is a function, $f[x \mapsto y]$ denotes the function that takes x to y and otherwise behaves the same as f ; the notation $f \downarrow A$ yields the restriction of the function f to the domain A ; the notation $f \setminus x$ is shorthand for $f \downarrow (\text{dom}(f) \setminus \{x\})$ or, if x is a set, $f \downarrow (\text{dom}(f) \setminus x)$. We make heavy use of notation like $\{x_y \mid y \in Y\}$ to describe a collection of elements x_y indexed by a set Y . Such notation allows us to give types to a collection of labels, e.g. $\langle l : \tau_l \mid l \in \mathcal{L} \rangle$. We define the binary operator \uplus , as follows:

$$\frac{\tau_l = \tau'_l \text{ for all } l \in \mathcal{L}_1 \cap \mathcal{L}_2}{\langle l : \tau_l \mid l \in \mathcal{L}_1 \rangle \uplus \langle l : \tau'_l \mid l \in \mathcal{L}_2 \rangle = \langle l : \tau_l \mid l \in \mathcal{L}_1, l : \tau'_l \mid l \in \mathcal{L}_2 \rangle}$$

The \uplus operator joins two possibly overlapping label/type collections. It is only defined when the operands agree on the type of any shared labels. While we group different label/type collections in syntactically distinct categories (e.g., trait types versus object types), we freely use \uplus to join two or more such collections in the same syntactic category.

5. Static semantics

Typing judgments in our calculus are written in terms of an ordered context Γ . Types can inhabit one of three syntactic categories: trait types θ , class types χ , or expression types τ . The most important judgment forms are the following:

$\Gamma \vdash T : \theta$	trait T has type θ
$\Gamma \vdash C : \chi$	class C has type χ
$\Gamma \vdash e : \tau$	expression e has type τ
$\Gamma \vdash \mu : \tau$	method μ has type τ
$\Gamma \vdash \tau_1 <: \tau_2$	τ_1 is a subtype of τ_2

We also make heavy use of well-formedness checks, which take the form $\Gamma \vdash \square \text{ ok}$ where \square is a type form, or $\Gamma \vdash \text{ok}$ to assert that the context Γ is well-formed.

Although traits can be viewed as sophisticated syntactic sugar that is “flattened” to a core class-based language [NDS06], there are advantages to recognizing traits directly. For the static semantics, giving types to traits allows the detection of a number of errors during trait manipulation that would otherwise not be detected until trait inlining; it also makes separate compilation of traits possible.

In order to typecheck a trait, we must know the types of all of its provided methods and of the self-methods, super-methods, and fields that it mentions. This type information amounts to a collection of assumptions about (or constraints on) the classes in which the trait will be inlined. The assumptions are guaranteed to hold of well-typed provided methods, which are syntactically required to specify their own type. The remaining assumptions constitute requirements of the trait, which fall into two categories:

1. Self-method requirements, which can be fulfilled via trait concatenation or trait method aliasing.
2. Super-method and field requirements, which can only be fulfilled when inlining a trait.

A straightforward type system might structure trait types as a collection of typed labels, some of which are marked as required; this is the approach of [FR03]. Unfortunately, this view of trait types forces a conservative typing of method exclusion. Consider the trait `TFoo`:

```
trait TFoo = {
  provides A() : unit { print("Hello!"); }
  provides B() : unit { self.A(); self.C(); }
  requires C : unit -> unit
}
```

The trait `TFoo` **exclude** `A` should require both `A` and `C`, while `TFoo` **exclude** `B` should not have any requirements. But a flat trait type assigning types to labels gives no way to distinguish between excluding `A` and excluding `B`; whenever a method is excluded, it must be conservatively counted as a required method because it may be invoked from another provided method.

To overcome this limitation, we track requirements on a per-provided-method, rather than per-trait, basis. The requirements for a provided method are collected into a set of *inlining assumptions* ρ that represent that method’s view of any class that inlines the trait. To distinguish between super-method and self-method requirements, we have a universe of super-method names, \mathcal{S}_U ; for any $s \in \mathcal{S}_U$ there is a unique $m \in \mathcal{M}_U$ such that $s = \text{super}.m$. A method may require both m and $\text{super}.m$, but for the two requirements to be coherent, their types must be compatible. Since our calculus only supports width subtyping, an overriding method must have the same type as its corresponding super-method. This condition is reflected in the well-formedness judgment for inlining assumptions:

$$\frac{\Gamma \vdash \tau_r \text{ ok for all } r \in \mathcal{R} \quad \tau_m = \tau_{\text{super}.m} \text{ for all } m \text{ with } m \in \mathcal{R}, \text{super}.m \in \mathcal{R}}{\Gamma \vdash \langle r : \tau_r \mid r \in \mathcal{R} \rangle \text{ ok}}$$

Inlining assumptions can be seen as a compact and uniform representation of a pair of object types, τ_{super} and τ_{self} , which are supertypes of the eventual super- and self-object types. To transform a set of inlining assumptions into its corresponding object types, we have the functions `super` and `self`:

$$\begin{aligned} \text{super } \langle r : \tau_r \mid r \in \mathcal{R} \rangle &= \langle m : \tau_{\text{super}.m} \mid \text{super}.m \in \mathcal{R} \rangle \\ \text{self } \langle r : \tau_r \mid r \in \mathcal{R} \rangle &= \langle l : \tau_l \mid l \in \mathcal{R} \cap \mathcal{L}_U \rangle \end{aligned}$$

A trait type θ is a labeled collection of method types with inlining assumptions. A trait formation expression provides a single

method and the initial inlining assumptions for that method. To type trait formation, we ensure that the given inlining assumptions are well-formed and that the given method can be typed under them:

$$\frac{\mu = m(x : \tau_1) : \tau_2 \{e\} \quad \Gamma \vdash \rho \uplus \langle\langle m : \tau \rangle\rangle \text{ ok} \quad \Gamma, \text{super} : \text{super}(\rho), \text{self} : \text{self}(\rho) \vdash \mu : \tau}{\Gamma \vdash \langle\langle \mu; \rho \rangle\rangle : \langle\langle m : \tau_m @ \rho' \rangle\rangle}$$

Notice that a *recursive* method must specify its own type in its inlining assumptions. The \uplus operator ensures that this type agrees with the method's actual type.

Although trait formation only produces single-method traits, in general traits will acquire several methods via composition. Each provided method has its own inlining assumptions, but we also want to consider the inlining assumptions for the trait as a whole. We introduce a function `Inl` that yields the trait inlining assumptions for given a trait type. Trait inlining assumptions are formed using the \uplus operator, which ensures that repeated assumptions about a method or field will have the same type:

$$\frac{\rho = \uplus_{m \in \mathcal{M}} (\rho_m \uplus \langle\langle m : \tau_m \rangle\rangle) \quad \Gamma \vdash \rho \text{ ok}}{\text{Inl}(\langle\langle m : \tau_m @ \rho_m \rangle\rangle_{m \in \mathcal{M}}) = \rho}$$

We also have a well-formedness check for trait types, which simply checks that `Inl` can succeed in producing inlining assumptions for the trait:

$$\frac{\text{Inl}(\theta) \text{ is defined}}{\Gamma \vdash \theta \text{ ok}}$$

The well-formedness check for trait types does a lot of work for the type system by centralizing coherency checking. Typing judgments need only include additional, specialized constraints. For trait composition, the additional constraint is that the given traits are disjoint (*i.e.*, specify a disjointly-named set of provided methods):

$$\frac{\Gamma \vdash T_1 : \theta_1 \quad \Gamma \vdash T_2 : \theta_2 \quad \text{dom}(\theta_1) \uplus \text{dom}(\theta_2) \quad \Gamma \vdash \theta_1 \uplus \theta_2 \text{ ok}}{\Gamma \vdash T_1 + T_2 : \theta_1 \uplus \theta_2}$$

Method exclusion requires that the method to be excluded is actually provided by the trait:

$$\frac{\Gamma \vdash T : \theta \quad m \in \theta \quad \Gamma \vdash \theta \setminus m \text{ ok}}{\Gamma \vdash T \text{ exclude } m : \theta \setminus m}$$

Method aliasing is somewhat more complex. For aliasing m as m' , the typing judgment first ensures that m is a provided method and that m' is not. It then forms a new trait type by joining inlining assumptions for m' to the old trait type. The inlining assumptions for m' are the same as those for m ; in particular, if m invoked itself, m' will invoke m .⁴

$$\frac{\Gamma \vdash T : \theta \quad m \in \theta \quad m' \notin \theta \quad \theta' = \theta \uplus \langle\langle m' : \theta(m) \rangle\rangle \quad \Gamma \vdash \theta' \text{ ok}}{\Gamma \vdash T \text{ alias } m \text{ as } m' : \theta'}$$

Note that m' may be the name of a required method, so aliasing can be used to fulfill trait requirements. The same is true for renaming.

Method renaming allows provided methods, required self-methods, and required super-methods to be renamed. To rename r to r' , the typing judgment checks that r and r' are not field names, that r is mentioned in the inlining assumptions (*i.e.*, it is either provided or required), and that r' is not a provided method. The new trait type is formed in two steps: first, r is renamed to r' in all of the sets of inlining assumptions; then, r is renamed to r' in the resulting trait type, which will only have an effect if r is a

⁴The issue of aliasing a recursive method is detailed in the dynamic semantics.

provided method:

$$\frac{\Gamma \vdash T : \langle\langle m : \tau_m @ \rho_m \rangle\rangle_{m \in \mathcal{M}} \quad r, r' \notin \mathcal{F}_U \quad r' \notin \mathcal{M} \quad r \in \rho_m \text{ for some } \rho_m \in \mathcal{M} \quad \theta = \langle\langle m : \tau_m @ (\rho_m[r'/r]) \rangle\rangle_{m \in \mathcal{M}} \uplus [r'/r] \quad \Gamma \vdash \theta \text{ ok}}{\Gamma \vdash T \text{ rename } r \text{ to } r' : \theta}$$

Typing method hiding is somewhat challenging for our system. We want to completely remove any mention of a method m' that is to be hidden, so that its name may be reused (possibly at a different type). But if we simply remove m' from the trait type, its requirements (which may not yet be fulfilled) will disappear from the trait as well, which is unsound. Our strategy is to remove m' from the trait's type, but transitively record its requirements in the inlining assumptions of any other provided method that invokes m' . In effect, we are doing a one-step path compression on the method call graph. The `hide` function achieves this result:

$$\text{hide}(\rho, m', \rho_{m'}) = \begin{cases} (\rho \uplus \rho_{m'}) \setminus m' & m' \in \rho \\ \rho & \text{otherwise} \end{cases}$$

The type judgement for hiding simply applies `hide` to each provided method, dropping m' :

$$\frac{\Gamma \vdash T : \langle\langle m : \tau_m @ \rho_m \rangle\rangle_{m \in \mathcal{M}} \quad m' \in \mathcal{M} \quad \theta = \langle\langle m : \tau_m @ \text{hide}(\rho_m, m', \rho_{m'}) \rangle\rangle_{m \in \mathcal{M} \setminus m'} \quad \Gamma \vdash \theta \text{ ok}}{\Gamma \vdash T \text{ hide } m' : \theta}$$

Notice that, if m' is not used by any other provided method in the trait, its requirements are dropped altogether. In this case, m' is dead code for the trait, and the method itself is eliminated in the dynamic semantics.

Finally, we have subclass formation. We first type the constructor, trait, and superclass. We then ensure that the trait's inlining assumptions about the class hold. A given type τ_l might be specified as part of the expected `super`-type, the expected `self`-type, and the actual superclass type; the typing judgement will only succeed if these specifications agree on the form of τ_l . Because of this requirement, we need only check the relationships between the various label sets to ensure that the class is well-typed:

$$\frac{\Gamma, x : \tau \vdash e_{\text{cons}} : \tau_{\text{cons}} \quad \Gamma, x : \tau \vdash e_F : \{f : \tau_f \mid f \in \mathcal{F}\} \quad \Gamma \vdash T : \theta \quad \Gamma \vdash C : \tau_{\text{cons}} \rightarrow \{\{l : \tau_l \mid l \in \mathcal{L}_C\}\} \quad \mathcal{F} \uplus \mathcal{L}_C \quad \mathcal{L} = \mathcal{L}_C \cup \mathcal{F} \cup \text{dom}(\theta) \quad \langle\{l : \tau_l \mid l \in \mathcal{L}_{\text{super}}\}\rangle = \text{super}(\text{Inl}(\theta)) \quad \mathcal{L}_{\text{super}} \subset \mathcal{L}_C \quad \langle\{l : \tau_l \mid l \in \mathcal{L}_{\text{self}}\}\rangle = \text{self}(\text{Inl}(\theta)) \quad \mathcal{L}_{\text{self}} \subset \mathcal{L}}{\Gamma \vdash \lambda(x : \tau).(\text{super } e_{\text{cons}}) \oplus e_F \text{ in } T \text{ extends } C : \tau \rightarrow \{\{l : \tau_l \mid l \in \mathcal{L}\}\}}$$

6. Dynamic semantics

We define evaluation with a big-step operational semantics. Evaluation judgments are written in the context of an environment E and a store S . Environments map names to trait, class, and expression values. Stores map addresses to object values, allowing objects to have mutable fields. Declaration evaluation yields a new environment and possibly a new store, while expression evaluation (which can occur in a declaration evaluation) yields an expression value and possibly a new store:

$$\begin{array}{ll} \text{Program evaluation} & E, S \vdash P \longrightarrow ev, E', S' \\ \text{Declaration evaluation} & E, S \vdash D \longrightarrow E', S' \\ \text{Expression evaluation} & E, S \vdash e \longrightarrow ev, S' \end{array}$$

At the core of our calculus is a revised notion of trait values. The standard view of trait values as simple collections of named provided methods is insufficient to support trait privacy, and makes a realistic model of deep renaming difficult to achieve. We adopt

Riecke and Stone’s approach [RS02] and distinguish between internal method names (which we term *slots*) and external method names. In our model, trait values are collections of internally named provided methods, some of which may be externally named as well. To support deep renaming of required methods, we distinguish between internal and external names for them as well; the internal name of a required method is eventually assigned to the method that fulfills that requirement.

More formally, a dictionary ϕ is a finite partial function that maps method names to slots. A method suite value Mv maps slots to method values. A trait value tv is a method suite value together with dictionaries for its provided and required methods:

$$\begin{array}{ll} \phi & ::= \{r \mapsto i \mid r \in \mathcal{R}\} & \text{dictionary} \\ Mv & ::= \{i \mapsto \mu v_i \mid i \in \mathcal{I}\} & \text{method suite value} \\ \mu v & ::= [E; \phi_\mu; \lambda(x : \tau).e; \rho] & \text{method value} \\ tv & ::= \langle Mv; \phi_P; \phi_R \rangle & \text{trait value} \end{array}$$

To prove type soundness we will need to give types to trait values, so we track inlining assumptions in method values.

Method hiding and renaming only affect a trait value’s dictionaries, *i.e.*, its external naming; its method suite remains unchanged.⁵ Notice that a method value μv contains a dictionary ϕ_μ in addition to the standard closure over the lexical environment E . This dictionary, established during evaluation of trait formation, can be thought of as a closure over the trait’s current external name environment. Unlike the trait itself, which has two dictionaries, each method closure contains only the single dictionary ϕ_μ ; from the perspective of a particular method, there is no difference between provided and required methods, because by the time the method is invoked all required methods must have been provided. It is ϕ_μ that allows a method to remain coherent in the presence of method hiding and renaming.

Ultimately, the methods in a trait value will be inlined into a class value. Method dispatch is dictionary-based, but the dictionary used is *dependent on the location of the call* (this is the crux of Riecke and Stone’s approach). More concretely: suppose we have a trait, `TFooBar`, which provides methods `foo` and `bar`. Further, suppose that `foo` invokes `bar`. When the trait is first formed, we assign slots to `foo` and `bar` and record these assignments in the ϕ_μ dictionaries for both methods. We can assume $\phi_\mu = \{\text{foo} \mapsto 1, \text{bar} \mapsto 2\}$ for both methods.⁶ If we then rename `bar` to `baz`, the dictionary for the trait itself is changed to reflect this ($\phi_P = \{\text{foo} \mapsto 1, \text{baz} \mapsto 2\}$), but the dictionaries for `foo` and `bar` remain the same. Suppose we inline `TFooBar` into a class and instantiate an object `obj`. We are able to invoke `obj.baz` using a dictionary giving an external view of the object (similar to ϕ_P). But if we invoke `obj.foo`, what happens when we reach the call to `self.bar`, which no longer exists from the external viewpoint? The ϕ_μ dictionary associated with `foo` is used to discover the *slot* for `bar`, which will be the same as the slot for `baz`.

The previous example glosses over a few evaluation details, which are explained below. The main idea to keep in mind is the motivation for all the dictionary juggling: we need to know what actual method to invoke, given a method name and context, and we need to do this in the face of aliasing, renaming, hiding, and excluding. Figure 6 shows the previous example and others in more formal detail; it should be read in parallel with the evaluation rules. The remainder of this section gives a complete description of evaluation for traits and classes and describes the nonstandard portions of expression evaluation. Some additional rules for expression eval-

⁵ One of the rules for renaming changes the closures in the suite, but it leaves the slot assignments of the suite intact.

⁶ In this example and others, we let $\mathcal{I}_U = \mathbb{Z}^+$, but formally \mathcal{I}_U is held abstract.

We define

$$\begin{array}{ll} \mu_{\text{foo}} & =_{\text{def}} \text{foo}(x : \text{unit}) : \text{unit} \{ \text{self.bar}() \} \\ \mu_{\text{bar}} & =_{\text{def}} \text{bar}(x : \text{unit}) : \text{unit} \{ \dots \} \end{array}$$

and derive

$$\begin{array}{l} \emptyset, \emptyset \vdash \text{TFoo} = \langle \mu_{\text{foo}}; \langle \langle \text{bar} : \text{unit} \rightarrow \text{unit} \rangle \rangle \rangle \longrightarrow E_1, \emptyset \\ E_1, \emptyset \vdash \text{TBar} = \langle \mu_{\text{bar}}; \langle \langle \rangle \rangle \rangle \longrightarrow E_2, \emptyset \\ E_2, \emptyset \vdash \text{TFooBar} = \text{TFoo} + \text{TBar} \longrightarrow E_3, \emptyset \end{array}$$

where $E_3 =$

$$\left\{ \begin{array}{l} \text{TFoo} \mapsto \langle \{1 \mapsto \mu v_{\text{foo}}\}; \{\text{foo} \mapsto 1\}; \{\text{bar} \mapsto 2\} \rangle, \\ \text{TBar} \mapsto \langle \{1 \mapsto \mu v_{\text{bar}}\}; \{\text{bar} \mapsto 1\}; \emptyset \rangle \\ \text{TFooBar} \mapsto \langle \langle \{1 \mapsto \mu v_{\text{foo}}, 2 \mapsto \mu v_{\text{bar}}\}; \{\text{foo} \mapsto 1, \text{bar} \mapsto 2\}; \emptyset \rangle \rangle \end{array} \right\}$$

so that

$$\begin{array}{l} E_3 \vdash \text{TFooBar} \text{ alias } \text{foo} \text{ as } \text{fiz} \\ \longrightarrow \langle \{1 \mapsto \mu v_{\text{foo}}, 2 \mapsto \mu v_{\text{bar}}, 3 \mapsto \mu v_{\text{foo}}\}; \{\text{foo} \mapsto 1, \text{bar} \mapsto 2, \text{fiz} \mapsto 3\}; \emptyset \rangle \end{array}$$

$$\begin{array}{l} E_3 \vdash \text{TFooBar} \text{ rename } \text{bar} \text{ to } \text{baz} \\ \longrightarrow \langle \{1 \mapsto \mu v'_{\text{foo}}, 2 \mapsto \mu v'_{\text{bar}}\}; \{\text{foo} \mapsto 1, \text{baz} \mapsto 2\}; \emptyset \rangle \end{array}$$

$$\begin{array}{l} E_3 \vdash \text{TFooBar} \text{ exclude } \text{bar} \\ \longrightarrow \langle \{1 \mapsto \mu v_{\text{foo}}\}; \{\text{foo} \mapsto 1\}; \{\text{bar} \mapsto 2\} \rangle \\ = E_3(\text{TFoo}) \end{array}$$

$$\begin{array}{l} E_3 \vdash \text{TFooBar} \text{ hide } \text{bar} \\ \longrightarrow \langle \{1 \mapsto \mu v''_{\text{foo}}, 2 \mapsto \mu v''_{\text{bar}}\}; \{\text{foo} \mapsto 1\}; \emptyset \rangle \end{array}$$

$$\begin{array}{l} E_3 \vdash (\text{TFooBar} \text{ hide } \text{bar}) + \text{TBar} \\ \longrightarrow \langle \{1 \mapsto \mu v''_{\text{foo}}, 2 \mapsto \mu v''_{\text{bar}}, 3 \mapsto \mu v_{\text{bar}}\}; \{\text{foo} \mapsto 1, \text{bar} \mapsto 3\}; \emptyset \rangle \end{array}$$

$$\begin{array}{l} E_3 \vdash \text{TFooBar} \text{ hide } \text{foo} \\ \longrightarrow \langle \{2 \mapsto \mu v_{\text{bar}}\}; \{\text{bar} \mapsto 2\}; \emptyset \rangle \\ \approx E_3(\text{TBar}) \end{array}$$

Figure 6. Evaluation examples

uation appear in the appendix. The complete formal system can be found in an extended version of this paper [RT06].

6.1 Trait evaluation

Trait evaluation judgments have the form $E \vdash T \longrightarrow tv$, since trait evaluation does not use or modify the store. We will often write

$$E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle$$

to assert that T evaluates to $\langle Mv; \phi_P; \phi_R \rangle$ with $\text{dom}(Mv) = \mathcal{I}$, $\text{dom}(\phi_P) = \mathcal{M}$, and $\text{dom}(\phi_R) = \mathcal{R}$; we always have that $\mathcal{M} \uplus \mathcal{R}$. For slot manipulation, we will use a function NS (“new slots”) which takes a set of method names and a set of slots, and a function FS (“fresh slots”) which takes two sets of slots. NS yields a dictionary mapping each of the given names to a unique, new slot not in the given set of slots. FS yields a translation function φ which maps each of the slots in its first parameter to a unique slot not contained in its second parameter. In other words,

$$\frac{\phi = \text{NS}(\mathcal{M}, \mathcal{I})}{\text{dom}(\phi) = \mathcal{M} \quad \text{rng}(\phi) \uplus \mathcal{I} \quad \phi \text{ is one-to-one}} \quad \frac{\varphi = \text{FS}(\mathcal{I}_1, \mathcal{I}_2)}{\text{dom}(\varphi) = \mathcal{I}_1 \quad \text{rng}(\varphi) \uplus \mathcal{I}_2 \quad \varphi \text{ is one-to-one}}$$

For example, we might have

$$\begin{aligned} \text{NS}(\{m_1, m_2\}, \{1, 2, 3\}) &= \{m_1 \mapsto 4, m_2 \mapsto 5\} \\ \text{FS}(\{1, 4\}, \{1, 2, 3\}) &= \{1 \mapsto 4, 4 \mapsto 5\} \end{aligned}$$

Evaluating a trait formation expression to a trait value establishes the initial slot assignment for the provided method and each required (super- or self-) method using NS. Slots are not established for field requirements, which cannot be renamed in our model. The resulting dictionary is used as the external name closure in the constructed method values:

$$\frac{\begin{array}{l} \phi_P = \text{NS}(\{m\}, \emptyset) \quad \phi_R = \text{NS}(\text{dom}(\rho) \setminus \mathcal{F}_U, \{\phi_P(m)\}) \\ Mv = \{\phi_P(m) \mapsto [E; \phi_P \cup \phi_R; \lambda(x : \tau).e; \rho]\} \end{array}}{E \vdash \langle m(x : \tau_1) : \tau_2 \{e\}; \rho \rangle \longrightarrow \langle Mv; \phi_P; \phi_R \rangle}$$

Notice that the dictionaries ϕ_P and ϕ_R are joined as a single dictionary in the method value. An alternative presentation of trait values could maintain a single dictionary at the trait level, and determine which methods were actually provided by examining $\text{dom}(Mv)$. For clarity and simplicity, we separate the dictionaries and maintain the invariants $\text{rng}(\phi_P) \subseteq \text{dom}(Mv)$ and $\text{dom}(Mv) \cap \text{rng}(\phi_R)$. The set of slots used by a trait is $\text{dom}(Mv) \cup \text{rng}(\phi_R)$.

Aliasing a method does not simply map a new external name to the existing internal name for the method; if we alias m as m' , we want the two methods to share implementations but to otherwise be independent, so that in particular we may later exclude m without impacting m' . Another concern is recursion: if m invokes itself, and we alias m as m' , what should happen to the recursive call for m' ? To see why this is important, consider that after aliasing m as m' , we could exclude m from the trait and replace it with some other method. This is a somewhat thorny issue, because if we choose to have m' recurse on itself rather than invoking m , we have only dealt with *direct* recursion; if m recurses on itself indirectly via some other method, m' would still end up invoking m via that same method. To keep the semantics simple and consistent, aliasing m as m' will leave invocations of m as still calling m .

Aliasing has two cases, which we handle with different rules. A method can be aliased to fulfill a method requirement, in which case the slot assigned for the required method is used for the alias, and the requirement is removed from ϕ_R :

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ m \in \mathcal{M} \quad m' \notin \mathcal{M} \quad m' \in \mathcal{R} \\ \phi'_P = \phi_P[m' \mapsto \phi_R(m')] \\ Mv' = Mv[\phi_R(m') \mapsto Mv(\phi_P(m))] \end{array}}{E \vdash T \text{ alias } m \text{ as } m' \longrightarrow \langle Mv'; \phi'_P; \phi_R \setminus m' \rangle}$$

If the aliasing operation does not fulfill a method requirement, we use NS to establish a new internal name for the alias:

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ m \in \mathcal{M} \quad m' \notin \mathcal{M} \quad m' \notin \mathcal{R} \\ \phi'_P = \phi_P \cup \text{NS}(\{m'\}, \mathcal{I} \cup \text{rng}(\phi_R)) \\ Mv' = Mv[\phi'_P(m') \mapsto Mv(\phi_P(m))] \end{array}}{E \vdash T \text{ alias } m \text{ as } m' \longrightarrow \langle Mv'; \phi'_P; \phi_R \rangle}$$

As with aliasing, provided methods may be renamed to fulfill required methods. In addition, required methods may be renamed. We give three rules for renaming. To rename a provided method without fulfilling a method requirement, we remove its old external name from the ϕ_P dictionary and insert a new mapping from the new name to the existing slot. We also rename the method in the

inlining assumptions for each provided method:

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ r \in \mathcal{M} \quad r' \notin \mathcal{M} \quad r' \notin \mathcal{R} \\ Mv = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}}\} \\ \phi'_P = (\phi_P \setminus r)[r' \mapsto \phi_P(r)] \\ Mv' = \{i \mapsto [E_i; \phi_i; e_i; \rho_i[r'/r]]^{i \in \mathcal{I}}\} \end{array}}{E \vdash T \text{ rename } r \text{ to } r' \longrightarrow \langle Mv'; \phi'_P; \phi_R \rangle}$$

The rule for renaming required methods is similar, but works on ϕ_R :

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ r \in \mathcal{R} \quad r' \notin \mathcal{M} \quad r' \notin \mathcal{R} \\ Mv = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}}\} \\ \phi'_R = (\phi_R \setminus r)[r' \mapsto \phi_R(r)] \\ Mv' = \{i \mapsto [E_i; \phi_i; e_i; \rho_i[r'/r]]^{i \in \mathcal{I}}\} \end{array}}{E \vdash T \text{ rename } r \text{ to } r' \longrightarrow \langle Mv'; \phi_P; \phi'_R \rangle}$$

Renaming a provided method to fulfill a method requirement is more complex. Unlike the situation for aliasing, we are not establishing a new, independent method, and thus the slot assignment for the original method must be retained. At the same time, the required method being fulfilled has its own slot assignment which must also be retained. Both of these assignments represent commitments made to the ϕ_μ dictionaries in the method closures for the trait. To fulfill the commitments, we slightly modify the promise by altering the target slots for ϕ_μ ; we arbitrarily choose to drop the required method's slot in favor of the provided method's slot. The modification is performed by composing a translation function φ with the original ϕ_μ dictionaries. The translation is the identity function on slots except at the required method's slot, where it maps to the provided method's slot:

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ r \in \mathcal{M} \quad r' \notin \mathcal{M} \quad r' \in \mathcal{R} \\ Mv = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}}\} \\ \phi'_P = (\phi_P \setminus r)[r' \mapsto \phi_P(r)] \\ \varphi = \{i \mapsto i, \phi_R(r') \mapsto \phi_P(r)\} \\ Mv' = \{i \mapsto [E_i; \varphi \circ \phi_i; e_i; \rho_i[r'/r]]^{i \in \mathcal{I}}\} \end{array}}{E \vdash T \text{ rename } r \text{ to } r' \longrightarrow \langle Mv'; \phi'_P; \phi_R \setminus m' \rangle}$$

To support hiding and excluding methods, and to properly type trait values, we need to drop unreachable hidden methods and unused required methods. The auxiliary judgement form $tv \hookrightarrow tv'$ rewrites a trait value after “dead-code elimination.” The judgement collects all of the slots mentioned in the inlining assumptions of the trait's provided methods, using the dictionary stored with each method to convert from method names to slots. The method suite and required method dictionary are then restricted to include only the mentioned slots; the notation $\phi_R^{-1}(\mathcal{I})$ signifies the inverse image of \mathcal{I} under ϕ_R :

$$\frac{\begin{array}{l} Mv = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \text{dom}(Mv)}\} \quad \mathcal{I}_P = \text{rng}(\phi_P) \\ \mathcal{I} = \mathcal{I}_P \cup \left(\bigcup_{i \in \mathcal{I}_P} \phi_i(\text{dom}(\rho_i)) \right) \quad \mathcal{R} = \phi_R^{-1}(\mathcal{I}) \end{array}}{\langle Mv; \phi_P; \phi_R \rangle \hookrightarrow \langle Mv \downarrow \mathcal{I}; \phi_P; \phi_R \downarrow \mathcal{R} \rangle}$$

Hiding a method removes it from the provided method dictionary. It also updates the inlining assumptions of all provided methods, using the hide function from the static semantics:

$$\frac{\begin{array}{l} E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ m \in \mathcal{M} \quad j = \phi_P(m) \quad Mv = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}}\} \\ Mv' = \{i \mapsto [E_i; \phi_i; e_i; \text{hide}(\rho_i, m, \rho_j)]^{i \in \mathcal{I}}\} \\ \langle Mv'; \phi_P \setminus m; \phi_R \rangle \hookrightarrow tv \end{array}}{E \vdash T \text{ hide } m \longrightarrow tv}$$

A method m hidden by this operation may still be available in the method suite, via its internal name: previously established methods can gain access to m by looking up the slot for m in the ϕ_μ dictionary bundled with their closure. If m is not used elsewhere in the trait, however, it is dropped.

To exclude a method m from a trait, it must be removed from both ϕ_P and Mv . Because other methods in the trait may reference m , we add m to ϕ_R , maintaining the internal name for m ; dead-code elimination will remove this spurious requirement, and possibly others, if m is not actually needed:

$$\frac{E \vdash T \longrightarrow \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \quad m \in \mathcal{M} \quad \langle Mv \setminus \phi_P(m); \phi_P \setminus m; \phi_R[m \mapsto \phi_P(m)] \rangle \hookrightarrow tv}{E \vdash T \text{ exclude } m \longrightarrow tv}$$

The most complex trait operation is composition. The two composed traits must have disjoint external names for their provided methods, but there may be considerable overlap in their internal naming, so we must adjust slot assignments accordingly. We use a technique similar to the one described for renaming, and treat the slot assignments of one of the traits as authoritative, creating a translation φ to adjust the other trait's dictionaries:

$$\frac{\begin{array}{l} E \vdash T_1 \longrightarrow \langle \mathcal{I}_1 Mv_1; \mathcal{M}_1 \phi_{P_1}; \mathcal{R}_1 \phi_{R_1} \rangle \\ E \vdash T_2 \longrightarrow \langle \mathcal{I}_2 Mv_2; \mathcal{M}_2 \phi_{P_2}; \mathcal{R}_2 \phi_{R_2} \rangle \\ \mathcal{M}_1 \uparrow \mathcal{M}_2 \quad Mv_2 = \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}_2}\} \\ \varphi_P = \{\phi_{P_2}(m) \mapsto \phi_{R_1}(m) \mid m \in \mathcal{M}_2 \cap \mathcal{R}_1\} \\ \varphi_R = \{\phi_{R_2}(m) \mapsto \phi_{P_1}(m) \mid m \in \mathcal{R}_2 \cap \mathcal{M}_1\} \\ \mathcal{I}'_1 = \mathcal{I}_1 \cup \text{rng}(\phi_{R_1}) \quad \mathcal{I}'_2 = \mathcal{I}_2 \cup \text{rng}(\phi_{R_2}) \\ \varphi_F = \text{FS}(\mathcal{I}'_2 \setminus \text{dom}(\varphi_R \cup \varphi_P), \mathcal{I}'_1) \\ \varphi = \varphi_P \cup \varphi_R \cup \varphi_F \quad \phi_P = \phi_{P_1} \cup (\varphi \circ \phi_{P_2}) \\ \phi_R = (\phi_{R_1} \cup (\varphi \circ \phi_{R_2})) \setminus (\mathcal{M}_1 \cup \mathcal{M}_2) \\ Mv = Mv_1 \cup \{\varphi(i) \mapsto [E_i; \varphi \circ \phi_i; e_i; \rho_i]^{i \in \mathcal{I}_2}\} \end{array}}{E \vdash T_1 + T_2 \longrightarrow \langle Mv; \phi_P; \phi_R \rangle}$$

The construction of the translation φ is performed in three steps: we construct φ_P for the methods provided in T_2 that fulfill methods required by T_1 ; we construct φ_R for the methods required in T_2 that are provided by T_1 ; finally, we construct φ_F to map all remaining slot assignments from T_2 to fresh slots that do not occur in T_1 . A new method suite Mv joins the method suite from T_1 and an adjusted method suite for T_2 that reflects the translated slot assignments.

6.2 Class evaluation

Class evaluation results in an evaluated constructor, a method suite, and a dictionary into that method suite:

$$\begin{array}{ll} cv ::= \langle \lambda v; Mv; \phi_C \rangle & \text{class value} \\ \lambda v ::= [E; \lambda(x : \tau).e] & \text{function value} \end{array}$$

The judgment form for class evaluation is written $E \vdash C \longrightarrow cv$; as with traits, the store is not used when evaluating a class expression. We evaluate **nil** to the empty class value, writing $\{\}$ for the empty field record:

$$\frac{}{E \vdash \text{nil} \longrightarrow \langle \{\emptyset; \lambda x.\{\}\}; \emptyset; \emptyset \rangle}$$

Handling inheritance requires us to deal with super-inocations. Since class methods may be overridden, and hence no longer accessible from the class's method suite, we bind super-inocations to new, hidden provided methods. The judgment form $cv \vdash tv \Longrightarrow Mv$ extends the method suite in tv to a new method suite that provides the relevant super-methods from cv :

$$\frac{Mv' = Mv \cup \{\phi_R(s) \mapsto Mv_C(\phi_C(s)) \mid s \in \text{dom}(\phi_R) \cap S_U\}}{\langle \lambda v_{\text{super}}; Mv_C; \phi_C \rangle \vdash \langle Mv; \phi_P; \phi_R \rangle \Longrightarrow Mv'}$$

Note that this rewriting does not create any slots; since super-methods are treated as requirements, we simply use the slots from the required method dictionary.

Evaluation of inheritance is similar to evaluation of trait composition: we are reconciling two method suites with incompatible slot assignments. In this case, however, a method with the same external name may be provided by both the trait and the superclass. We retain the class's slot assignment for the method, but use the trait's implementation, thereby overriding the method:

$$\frac{\begin{array}{l} E \vdash T \longrightarrow tv \quad tv = \langle \mathcal{I} Mv; \mathcal{M} \phi_P; \mathcal{R} \phi_R \rangle \\ E \vdash C \longrightarrow cv \quad cv = \langle \lambda v_{\text{super}}; Mv_C; \phi_C \rangle \\ cv \vdash tv \Longrightarrow \{i \mapsto [E_i; \phi_i; e_i; \rho_i]^{i \in \mathcal{I}'}\} \\ \varphi_P = \{\phi_P(m) \mapsto \phi_C(m) \mid m \in \mathcal{M} \cap \text{dom}(\phi_C)\} \\ \varphi_R = \{\phi_R(m) \mapsto \phi_C(m) \mid m \in \mathcal{R} \cap \text{dom}(\phi_C)\} \\ \varphi_F = \text{FS}(\mathcal{I}' \setminus \text{dom}(\varphi_P), \text{dom}(Mv_C)) \\ \varphi = \varphi_P \cup \varphi_R \cup \varphi_F \\ Mv'_T = \{\varphi(i) \mapsto [E_i; \varphi \circ \phi_i; e_i; \rho_i]^{i \in \mathcal{I}'}\} \\ Mv'_C = (Mv_C \setminus \text{rng}(\varphi_P)) \cup Mv'_T \\ \text{super} \notin \text{dom}(E) \quad E_{\text{cons}} = E[\text{super} \mapsto \lambda v_{\text{super}}] \\ \lambda v_{\text{cons}} = [E_{\text{con}}; \lambda(x : \tau).(super \ e_{\text{cons}}) \oplus e_F] \end{array}}{E \vdash \lambda(x : \tau).(super \ e_{\text{cons}}) \oplus e_F \text{ in } T \text{ extends } C \longrightarrow \langle \lambda v_{\text{cons}}; Mv'_C; \phi_C \cup (\varphi \circ \phi_P) \rangle}$$

In reading the above rule, recall that $\text{rng}(\phi_R) \uparrow \mathcal{I}$. We also have that $\text{dom}(\varphi_R) \uparrow \mathcal{I}'$, because no super-method name (**super**. m) can appear in ϕ_C .

6.3 Object instantiation and method dispatch

Expression evaluation is written $E, S \vdash e \longrightarrow ev, S'$. Most of the rules for expression evaluation are standard, but we describe the rules related to object instantiation and method dispatch to highlight the role that ϕ_μ dictionaries play. The remaining expression evaluation rules can be found in Appendix B.

An object value is a field record value paired with a method suite:

$$\begin{array}{ll} ov ::= \langle fv; Mv \rangle & \text{object value} \\ fv ::= \{f = ev_f \mid f \in \mathcal{F}\} & \text{field record value} \\ ev ::= (a, \phi) & \text{object reference} \\ & | \lambda v & \text{function value} \\ & | fv & \text{field record value} \\ & | () & \text{unit value} \end{array}$$

Evaluating an object instantiation takes a class and a constructor parameter and updates the store to map a new address a to a new object value. The evaluation yields the address a paired with the class's dictionary ϕ_C :

$$\frac{\begin{array}{l} E \vdash C \longrightarrow \langle [E_F; \lambda x.e_F]; Mv; \phi_C \rangle \\ E, S \vdash e \longrightarrow ev, S_1 \\ E_F[x \mapsto ev], S_1 \vdash e_F \longrightarrow fv, S_2 \\ a \notin \text{dom}(S_2) \end{array}}{E, S \vdash \text{new } C \ e \longrightarrow (a, \phi_C), S_2[a \mapsto \langle fv; Mv \rangle]}$$

To evaluate a method dispatch $e.m$, we first evaluate e to an address a and dictionary ϕ . We look up the object value associated with the address, then use the dictionary to index into the object's method suite at m , yielding the closure for the method m . The evaluation results in a function value equivalent to m 's closure: the environment of the function value takes *self* to (a, ϕ_μ) so that self-inocations within m have m 's view of the class:

$$\frac{\begin{array}{l} E, S \vdash e \longrightarrow (a, \phi), S' \quad S'(a) = \langle fv; Mv \rangle \\ Mv(\phi(m)) = [E_m; \phi_m; \lambda(x : \tau).e_m; \rho_m] \end{array}}{E, S \vdash e.m \longrightarrow [E_m[\text{self} \mapsto (a, \phi_m)]; \lambda(x : \tau).e_m], S'}$$

To support super-method invocation, we invoke the super-method name in the context of `self`:

$$\frac{E, S \vdash \mathbf{self}.\langle \mathbf{super}.m \rangle \longrightarrow ev, S'}{E, S \vdash \mathbf{super}.m \longrightarrow ev, S'}$$

This approach works because super-method names are indexed in the dictionary for each method, and super-method implementations are provided, with the appropriate slot, during class formation. Notice that this rule does *not* apply to the constructor for a class, which does not invoke a super-method but invokes “super” itself.

Finally, the evaluation of `self` is a simple lookup in the current environment.

$$\frac{}{E, S \vdash \mathbf{self} \longrightarrow E(\mathit{self}), S}$$

7. Type soundness

A detailed proof of type soundness and the run-time typing rules for the system can be found in the tech report version of this paper [RT06]. Here we give a brief overview of the issues and theorems related to type soundness.

One important issue for type soundness is typing for link-time and run-time values. We must give types to trait values, class values, and expression values, including object addresses. However, because stores may be cyclic, we cannot type addresses via a recursive examination. We follow the standard technique and introduce a store typing Σ , which is a map from object addresses to object types. The intuition behind this approach is that well-typed programs will always store values of the same type in a given address, so we need only type locations when they are first introduced. Run-time typing judgments are given in terms of store typings, and thus have the form $\Sigma \vdash \square : \square$. Since environments and stores are also runtime values, we type them as well: the judgment $\Sigma \vdash E : \Gamma$ says that environment E has type Γ , while $\vdash S : \Sigma$ says that store S has type Σ .

Type preservation is easy to state for a big-step semantics, but the difficulty of stating a progress property is usually held as a drawback of the big-step style. The difficulty with progress is that, for a big-step semantics, non-termination and `WRONG` are indistinguishable. In particular, if we do not have \emptyset , $\emptyset \vdash P \longrightarrow ev, E, S$, it could either be that P diverges or that P goes wrong. To overcome this problem, we follow [FR04] and introduce a *height function* $H_{E, S}$, which gives the height of the derivation tree for a program under E and S . We have the following definition:

DEFINITION 7.1 (Divergence). *We say a program P diverges if there is no n such that $H_{\emptyset, \emptyset}(P) = n$.*

If $E, S \vdash P \longrightarrow ev$ then $H_{E, S}(P) = n$, where n is the height of the derivation tree for the judgment. If P diverges in the context of E and S , then $H_{E, S}(P)$ diverges as well. Most importantly, if P does not evaluate to any ev under E and S , then $H_{E, S}(P)$ *converges*, and measures the height of the evaluation derivation up to the point that P went wrong. For example,

$$H_{E, S}(t = (\bar{\alpha})T; P) = 1 + \begin{cases} H_{E', S}(P) & \text{if } E, S \vdash t = (\bar{\alpha})T \longrightarrow E', S \\ 1 & \text{otherwise} \end{cases}$$

Thus, the height function allows us to distinguish between non-termination and stuck states, without having explicit `WRONG` transitions. Using this height function, we can state and prove the following soundness result [RT06]:

THEOREM 7.1 (Type soundness). *If $\emptyset \vdash_P P : \tau$ then either P diverges or there exist a store typing Σ , a store S , a context Γ , an*

environment E , a type τ' , and an expression value ev such that $\Sigma \vdash E : \Gamma$ and $\vdash S : \Sigma$ and $\emptyset, \emptyset \vdash P \longrightarrow ev, E, S$ and $\Sigma \vdash ev : \tau'$ and $\epsilon \vdash \tau' <: \tau$.

This result says that a well-typed, *terminating* program P does not go wrong. In particular, P will evaluate to a result that improves on its static type. While this is a weaker statement than soundness for a small-step semantics, a characterization of terminating programs is sufficient for our purposes with this calculus.

8. Related work

Traits were originally proposed as a mechanism for `SMALLTALK` by Schärli *et al.* [SDNB03]. In addition to studying the language design and methodological issues, they also developed a formal model for their system [SDN⁺02]. The most important difference between our system and this original work is that we are working in a strongly-typed setting, instead of an untyped language. Another major point of difference is that we have abandoned the *flattening* property [NDS06] in favor of supporting deep renaming and private methods in traits.⁷ In this sense our system is not a conservative extension of traditional class-based designs, but the meta-programming techniques enabled by our system provide an argument for a dictionary-based semantics. Further argument for this approach can be found in Riecke and Stone’s paper [RS02].

The introduction of traits for `SMALLTALK` has prompted a flurry of work on traits for statically-typed languages. Fisher and Reppy developed the first formal model of traits in a statically typed setting [FR04]. This model was subsequently extended to support polymorphic traits (key for examples such as the synchronized readers) and stateful objects [FR03]. This extended trait calculus was the starting point for the system in this paper. We have made a number of refinements in the static semantics. Our calculus tracks method requirements on a per-method basis, which provides more accuracy when excluding methods. We have also unified the handling of methods and super methods in the requirements by introducing a separate namespace for super methods (*e.g.* `super.m`). While a simple trick, this technique streamlines the static semantics significantly. We have reformulated the link-time and dynamic semantics of method binding using Riecke–Stone dictionaries, which allows the support for the deep operations of renaming and hiding at the trait level. These last two are our most significant additions to the Fisher–Reppy calculus.

Smith and Drossopoulou recently described a family of three different extensions of `JAVA` with traits [SD05]. The first of these, *Chai*₁, defers all checking until traits have been included in a class. The second, *Chai*₂, adds trait types and is similar in expressiveness to the Fisher–Reppy trait calculus, with a couple of exceptions. Like `JAVA`, *Chai*₂ uses nominal subtyping, instead of structural typing, and does not have polymorphic traits. The differences between our work the Fisher–Reppy calculus apply to *Chai*₂ as well. *Chai*₃ extends *Chai*₂ by allowing traits to be replaced at runtime. This feature is orthogonal to the focus of our work, but could be added to our system.

Another proposed design of traits for `JAVA` is *Featherweight Trait JAVA* (FTJ) [LS04], which adds traits to Featherweight `JAVA` [IPW01]. This system is fairly similar to the Fisher–Reppy trait calculus (with some technical differences), and does not support either deep renaming or private methods at the trait level.

There are strong similarities between traits and *mixins* [BC90, FKF98, OAC⁺04], which are another mechanism designed to give many of the benefits of multiple inheritance without the complica-

⁷Strictly speaking, one has to reformulate the flattening property for our system, since we do not have a mechanism for defining methods directly in a class.

tions. The main difference between mixins and traits is that mixins force a linear order in their composition (it is this order that avoids the complexities of the diamond property). This linear order introduces fragility problems and may make code maintenance more difficult [SDNB03]. Mixin mechanisms must also deal with constructor functions, which can be another source of fragility, since it is hard to predict what the interface of the super-class constructor will be. A solution to this problem is to define the mixin's constructors at the point of mixin application [ALZ03]. This issue does not affect traits, since they are not defining or initializing object state. Personalities are another trait-like mechanism designed for JAVA, although they are much more limited in their expressiveness [Bla98] and they do not have a formal model. The language Scala has a special form of abstract class called a trait class [OAC⁺04]. Trait classes are used as mixins and also to support family polymorphism, but they do not support the trait operations such as exclusion, or our deep renaming.

Bracha's Jigsaw framework is often cited as the first formal account of mixins [Bra92]. Like our calculus, and unlike the other trait systems discussed above, Jigsaw supports deep renaming (Bracha calls it global renaming) and method hiding. His system also has a static binding mechanism (called freezing). Bracha gave a dynamic semantics and a type system for Jigsaw, but did not prove type soundness. While it is possible that our metaprogramming idiom could be applied to mixins in the Jigsaw framework, we believe that traits are a better fit.

Examples like the `TSync` trait closely resemble the use of aspects [KLM⁺97] to specify "cross-cutting" concerns. While traits have always had some overlap with aspects, trait-based metaprogramming brings the two even closer. Both traits and aspects are specified outside of the classes to which they are applied; the primary difference between the two is how they are applied to classes. Aspects specify points of application via *pointcuts*, which pick out *join points* in the target classes; hence, aspects are in control of their own application. On the other hand, traits are completely inert unless they are explicitly inlined during class formation, leaving control to the class implementor.

9. Conclusion

Traits provide a promising mechanism for constructing class hierarchies from reusable components. We have introduced two new trait operations, method *hiding* and *renaming*, which are the *deep* counterparts to method exclusion and aliasing. These operations provide new ways of resolving conflicts in trait composition. Furthermore, they support privacy at the trait level and trait-based metaprogramming. Our formal model gives a detailed semantics to these operations in a statically typed setting, while also improving the granularity of the type system over previous calculi. There are several remaining questions for a concrete language design built around our calculus, which we briefly raise:

- In a language with a rich module system, it is not yet clear how traits should interact with features such as signatures and functors. One can imagine using signature ascription to implement method hiding in traits (as done in `MOBY` at the class level [FR99]).
- While the type system we have presented in this paper is flexible, it is also verbose, since each method provided by a trait specifies its own requirements. This granular type information can be inferred from a trait definition, but introducing traits into a module system will require a programmer to explicitly write down trait signatures, and it is not clear what form such signatures should take.

- The most important questions posed by our work regard the design of a trait metalanguage. At the least, such a language should include an explicit abstraction mechanism for trait method names, but the design space is large and essentially unexplored. We plan to first implement the trait calculus of this paper in the language `MOBY`, which will allow us to gain experience with trait programming and manual trait-based metaprogramming. We hope that programming experience will lead to the recognition of patterns and idioms that can then be codified into a trait metalanguage.

Acknowledgments We would like to thank Kathleen Fisher, who provided helpful feedback on an early draft and insight into the type theory for the paper. The anonymous reviewers gave thorough feedback on the formal system and good insights on the presentation; their work is appreciated.

References

- [ALZ03] Ancona, D., G. Lagorio, and E. Zucca. Jam—designing a java extension with mixins. *ACM Transactions on Programming Languages and Systems*, **25**(5), September 2003, pp. 641–712.
- [BC90] Bracha, G. and W. Cook. Mixin-based inheritance. In *ECOOP'90*, New York, NY, October 1990. ACM, pp. 303–311.
- [Bla98] Blando, L. Designing and programming with personalities. Master's dissertation, Northeastern University, Boston, MA, December 1998. Available as Technical Report NU-CCS-98-12.
- [Bra92] Bracha, G. *The Programming Language Jigsaw: Mixins, Modularity and Multiple Inheritance*. Ph.D. dissertation, University of Utah, March 1992.
- [BSD03] Black, A. P., N. Schärli, and S. Ducasse. Applying traits to the Smalltalk collection classes. In *OOPSLA'03*, New York, NY, October 2003. ACM, pp. 47–64.
- [FKF98] Flatt, M., S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *POPL'98*, New York, NY, January 1998. ACM, pp. 171–183.
- [FR99] Fisher, K. and J. Reppy. The design of a class mechanism for Moby. In *PLDI'99*, New York, NY, May 1999. ACM, pp. 37–49.
- [FR03] Fisher, K. and J. Reppy. Statically typed traits. *Technical Report TR-2003-13*, Dept. of Computer Science, U. of Chicago, Chicago, IL, December 2003.
- [FR04] Fisher, K. and J. Reppy. A typed calculus of traits. In *FOOL11*, January 2004.
- [IPW01] Igarashi, A., B. C. Pierce, and P. Wadler. Featherweight java: a minimal core calculus for java and gj. *ACM Transactions on Programming Languages and Systems*, **23**(3), 2001, pp. 396–450.
- [KLM⁺97] Kiczales, G., J. Lamping, A. Menhdhekar, C. Maeda, C. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-oriented programming. In M. Akşit and S. Matsuoka (eds.), *ECOOP'97*, vol. 1241 of *LNCS*, pp. 220–242. Springer-Verlag, New York, NY, 1997.
- [LS04] Liquori, L. and A. Spiwack. Featherweight-trait java: A trait-based extension for FJ. *Technical Report RR-5247*, Institut National de Recherche en Informatique et en Automatique, June 2004.
- [NDS06] Nierstrasz, O., S. Ducasse, and N. Schärli. Flattening traits. *Journal of Object Technology*, **5**(3), May 2006. (to appear).
- [OAC⁺04] Odersky, M., P. Altherr, V. Cremet, B. Emir, S. Micheloud, N. Mihaylov, M. Schinz, E. Stenman, and M. Zenger. *The Scala Language Specification (Draft)*. Programming Methods Laboratory, EPFL, Switzerland, February 2004. Available from lamp.epfl.ch/scala.

- [Qui04] Quitslund, P. J. Java traits — improving opportunities for reuse. *Technical Report CSE 04-005*, OGI School of Science & Engineering, September 2004.
- [RS02] Riecke, J. G. and C. A. Stone. Privacy via subsumption. *Information and Computation*, 172(1), January 2002, pp. 2–28. A preliminary version appeared in FOOL5.
- [RT06] Reppy, J. and A. Turon. A foundation for trait-based metaprogramming (extended version). *Technical report*, Dept. of Computer Science, U. of Chicago, Chicago, IL, 2006.
- [SD05] Smith, C. and S. Drossopoulou. Chai: Traits for Java-like languages. In *ECOOP'05*, LNCS, New York, NY, July 2005. Springer-Verlag.
- [SDN⁺02] Schärli, N., S. Ducasse, O. Nierstrasz, R. Wuyts, and A. Black. Traits: The formal model. *Technical Report CSE 02-013*, OGI School of Science & Engineering, November 2002. (revised February 2003).
- [SDNB03] Schärli, N., S. Ducasse, O. Nierstrasz, and A. Black. Traits: Composable units of behavior. In *ECOOP'03*, vol. 2743 of LNCS, New York, NY, July 2003. Springer-Verlag, pp. 248–274.
- [Str94] Stroustrup, B. *The Design and Evolution of C++*. Addison-Wesley, Reading, MA, 1994.

A. Additional typing judgments

Subtyping:

$$\boxed{\Gamma \vdash \tau_1 <: \tau_2}$$

$$\frac{\Gamma \vdash \tau \text{ ok}}{\Gamma \vdash \tau <: \tau} \quad \frac{\Gamma \vdash \langle l : \tau_l \text{ }^{l \in \mathcal{L}_1} \rangle \text{ ok} \quad \mathcal{L}_2 \subset \mathcal{L}_1}{\Gamma \vdash \langle l : \tau_l \text{ }^{l \in \mathcal{L}_1} \rangle <: \langle l : \tau_l \text{ }^{l \in \mathcal{L}_2} \rangle}$$

$$\frac{\Gamma \vdash \tau'_2 <: \tau'_1 \quad \Gamma \vdash \tau''_1 <: \tau''_2}{\Gamma \vdash \tau'_1 \rightarrow \tau''_1 <: \tau'_2 \rightarrow \tau''_2}$$

Expression typing:

$$\boxed{\Gamma \vdash e : \tau}$$

$$\frac{\Gamma \vdash \text{ok} \quad x \in \text{dom}(\Gamma)}{\Gamma \vdash x : \Gamma(x)} \quad \frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda(x : \tau).e : \tau_1 \rightarrow \tau_2}$$

$$\frac{\Gamma \vdash e_1 : \tau' \rightarrow \tau \quad \Gamma \vdash e_2 : \tau'}{\Gamma \vdash e_1 e_2 : \tau} \quad \frac{\Gamma \vdash C : \tau \rightarrow \{l : \tau_l \text{ }^{l \in \mathcal{L}}\} \quad \Gamma \vdash e : \tau}{\Gamma \vdash \mathbf{new} C e : \langle l : \tau_l \text{ }^{l \in \mathcal{L}} \rangle}$$

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash \mathbf{self} : \Gamma(\mathbf{self})} \quad \frac{\Gamma \vdash \Gamma(\mathbf{super}) <: \langle m : \tau_m \rangle}{\Gamma \vdash \mathbf{super}.m : \tau_m}$$

$$\frac{\Gamma \vdash \text{ok}}{\Gamma \vdash () : \mathbf{unit}} \quad \frac{\Gamma \vdash e_1 : \tau \quad \Gamma \vdash \tau <: \langle f : \tau_f \rangle \quad \Gamma \vdash e_2 : \tau_f}{\Gamma \vdash e_1.f := e_2 : \mathbf{unit}}$$

$$\frac{\Gamma \vdash \text{ok} \quad \Gamma \vdash e_f : \tau_f \quad \forall f \in \mathcal{F}}{\Gamma \vdash \{f = e_f \text{ }^{f \in \mathcal{F}}\} : \{f : \tau_f \text{ }^{f \in \mathcal{F}}\}}$$

$$\frac{\Gamma \vdash e_1 : \{f = e_f \text{ }^{f \in \mathcal{F}_1}\} \quad \Gamma \vdash e_2 : \{f = e_f \text{ }^{f \in \mathcal{F}_2}\} \quad \mathcal{F}_1 \uplus \mathcal{F}_2}{\Gamma \vdash e_1 \oplus e_2 : \{f = e_f \text{ }^{f \in \mathcal{F}_1 \cup \mathcal{F}_2}\}}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash \tau <: \langle l : \tau_l \rangle}{\Gamma \vdash e.l : \tau_l} \quad \frac{\Gamma \vdash e : \tau \quad \Gamma \vdash \tau <: \tau'}{\Gamma \vdash e : \tau'}$$

Declaration typing:

$$\boxed{\Gamma \vdash D \Rightarrow \Gamma'}$$

$$\frac{t \notin \text{dom}(\Gamma) \quad \Gamma, \bar{\alpha} \vdash T : \theta}{\Gamma \vdash t = (\bar{\alpha})T \Rightarrow \Gamma, t : \Lambda(\bar{\alpha}).\theta} \quad \frac{c \notin \text{dom}(\Gamma) \quad \Gamma \vdash C : \chi}{\Gamma \vdash c = C \Rightarrow \Gamma, c : \chi}$$

$$\frac{x \notin \text{dom}(\Gamma) \quad \Gamma \vdash e : \tau}{\Gamma \vdash x = e \Rightarrow \Gamma, x : \tau}$$

Program typing:

$$\boxed{\Gamma \vdash_P P : \tau}$$

$$\frac{\Gamma \vdash D \Rightarrow \Gamma' \quad \Gamma' \vdash_P P : \tau}{\Gamma \vdash_P D; P : \tau} \quad \frac{\Gamma \vdash e : \tau}{\Gamma \vdash_P e : \tau}$$

B. Additional evaluation rules

Program evaluation:

$$\boxed{E, S \vdash P \longrightarrow ev, E', S'}$$

$$\frac{E, S \vdash D \longrightarrow E', S'}{E', S' \vdash P \longrightarrow ev, E'', S''} \quad \frac{E, S \vdash e \longrightarrow ev, S'}{E, S \vdash e \longrightarrow ev, E, S'}$$

Declaration evaluation:

$$\boxed{E, S \vdash D \longrightarrow E', S'}$$

$$\frac{E \vdash T \longrightarrow tv \quad t \notin \text{dom}(E)}{E, S \vdash t = (\bar{\alpha})T \longrightarrow E[t \mapsto (\bar{\alpha})tv], S}$$

$$\frac{E \vdash C \longrightarrow cv \quad c \notin \text{dom}(E)}{E, S \vdash c = C \longrightarrow E[c \mapsto cv], S}$$

$$\frac{E, S \vdash e \longrightarrow ev \quad x \notin \text{dom}(E)}{E, S \vdash x = e \longrightarrow E[x \mapsto ev], S'}$$

Expression evaluation:

$$\boxed{E, S \vdash e \longrightarrow ev, S'}$$

$$\frac{x \in E}{E, S \vdash x \longrightarrow E(x), S} \quad \frac{E, S \vdash \lambda(x : \tau).e \longrightarrow [E; \lambda(x : \tau).e], S}{E, S \vdash \lambda(x : \tau).e \longrightarrow [E'; \lambda(x : \tau).e], S_1}$$

$$\frac{E, S_1 \vdash e_1 \longrightarrow [E'; \lambda(x : \tau).e], S_1 \quad E, S_1 \vdash e_2 \longrightarrow ev_2, S_2 \quad E'[x \mapsto ev_2], S_2 \vdash e \longrightarrow ev, S_3}{E, S \vdash e_1 e_2 \longrightarrow ev, S_3}$$

$$\frac{E, S \vdash e \longrightarrow (a, \phi), S' \quad S'(a) = \langle f v; M v \rangle \quad f v(f) = ev_f}{E, S \vdash e.f \longrightarrow ev_f, S'}$$

$$\frac{E, S \vdash e_1 \longrightarrow (a, \phi), S_1 \quad E, S_1 \vdash e_2 \longrightarrow ev_2, S_2 \quad S_2(a) = \langle \{f = ev_f \text{ }^{f \in \mathcal{F}}\}; M v \rangle}{E, S \vdash e_1.f := e_2 \longrightarrow ev_2, S_3}$$

$$\frac{E, S \vdash e_1 \longrightarrow (a, \phi), S_1 \quad E, S_1 \vdash e_2 \longrightarrow ev_2, S_2 \quad S_2(a) = \langle \{f = ev_f \text{ }^{f \in \mathcal{F} \setminus f'}\}; M v \rangle}{E, S \vdash e_1.f := e_2 \longrightarrow ev_2, S_3}$$

$$\frac{E, S \vdash e_1 \longrightarrow \{f = ev_f \text{ }^{f \in \mathcal{F}_1}\}, S_1 \quad E, S_1 \vdash e_2 \longrightarrow \{f = ev_f \text{ }^{f \in \mathcal{F}_2}\}, S_2}{E, S \vdash e_1 \oplus e_2 \longrightarrow \{f = ev_f \text{ }^{f \in \mathcal{F}_1 \cup \mathcal{F}_2}\}, S_2}$$

$$E, S \vdash e_1 \longrightarrow ev_1, S_1$$

⋮

$$E, S_{n-1} \vdash e_n \longrightarrow ev_n, S_n$$

$$\frac{E, S \vdash \{f_1 = e_1, \dots, f_n = e_n\} \longrightarrow \{f_1 = ev_1, \dots, f_n = ev_n\}, S_n}{E, S \vdash \{ \} \longrightarrow \{ \}, S} \quad \frac{E, S \vdash () \longrightarrow (), S}{E, S \vdash () \longrightarrow (), S}$$

$$\boxed{E, S \vdash \{ \} \longrightarrow \{ \}, S} \quad \boxed{E, S \vdash () \longrightarrow (), S}$$