

# Tribe: More Types for Virtual Classes\*

Dave Clarke<sup>1</sup>, Sophia Drossopoulou<sup>2</sup>, James Noble<sup>3</sup>, and Tobias Wrigstad<sup>4</sup>

<sup>1</sup> CWI, Amsterdam, The Netherlands

<sup>2</sup> Imperial College London, UK

<sup>3</sup> Victoria University of Wellington, Wellington, NZ

<sup>4</sup> Stockholm University, Stockholm, Sweden

**Abstract.** Beginning with BETA, a range of programming language mechanisms have been developed to allow inheritance in the presence of mutually dependent classes. This paper presents Tribe, a type system which generalises and simplifies other formalisms of such mechanisms, by treating issues which are inessential for soundness, such as the precise details of dispatch and path initialisation, as orthogonal to the core formalism. Tribe can support path types dependent simultaneously on both classes and objects, which is useful for writing library code, and ubiquitous access to an object's family (= `owner`), which offers family polymorphism without the need to drag around family arguments. Languages based on Tribe will be both simpler and more expressive than existing designs, while having a simpler type system, serving as a useful basis for future language designs.

## 1 Introduction

Languages and formalisms such as BETA [13], gBeta [5], Caesar [15], Scala [8], Jx [16], .FJ [11], Concord [12] and *vc* [7] offer advanced notions of inheritance which overcome many of the weaknesses of standard single and multiple inheritance. In particular, they feature a notion of *family polymorphism* in which a group of mutually dependent classes can be inherited together in such a way that the relationship between the classes is preserved through inheritance. *Virtual classes* are one means for achieving family polymorphism. A virtual class is a nested class which can be overridden like a method. A key advantage of overriding a class definition, as opposed to extending it, is that the class name remains the same. This enables any code which operates on one family to also work for extensions of that family.

In this paper, we introduce Tribe, a general yet simple system which incorporates these mechanisms into a single seamless framework, extends the language of types, the notion of subtypes, offers powerful inheritance, and tackles method overriding.

In this section we illustrate the basics of family polymorphism, and outline our extensions to existing systems. In the following example, nested classes are

---

\* Work partially supported by a gift from Microsoft Research, , by the Royal Society of New Zealand Marsden Fund, and the EU grant MOBIUS.

used to express that the family `Graph` has member (virtual) classes `Node` and `Edge`.

```
class Graph {
  class Node {
    Edge connect(Node other) { return new Edge(this, other); }
  }
  class Edge {
    Node from, to;
    Edge(Node f, Node t) { from = f; to = t; }
  }
}
```

Subclassing enables the construction of new families from existing ones. In the example below, the family `ColouredGraph` inherits member classes `Node` and `Edge`, and can extend them using a mechanism known as *further binding*. Here, a new field `nodeColour` is added to `Node` inherited from `Graph`.

```
class ColouredGraph extends Graph {
  class Node {
    Colour nodeColour;
  }
}
```

*Distinguishing the Families* Grouping classes into families raises the question as to whether mixing objects from different families should be allowed. In our example the intention is that coloured nodes should be connected to coloured nodes only. However:

```
Graph.Node n = new Graph.Node();
Graph.Node cn = new ColouredGraph.Node(); // dubious subsumption

n.connect(cn); // mixes two kinds of nodes
```

Through the—*prima facie* obvious—subtype relation `ColouredGraph.Node ≤ Graph.Node`, we can create an edge between a node and a coloured node. If the coloured node defines additional methods not found in `node`, this dubious use of subsumption could lead to “message not understood” errors.

A number of approaches have been devised to address this problem, by keeping better track of the family through the type system. The first approach, used in Concord, .FJ, and Jx, [12, 11, 16, 15] uses types (e.g. `Graph.Node`) based on a static notion of a type family, typically a *class family*, and achieves its goal by eliminating relationships such as `ColouredGraph.Node ≤ Graph.Node`. The second approach, used in *vc* [7] and *vObj/Scala* [17, 8], keeps close track of which *object family* the nested classes belong to by using paths in types. An *object family type* such as `x.f.C` refers to the class `C` which is nested within the object at the end of path `x.f`. Paths have to be invariant, and so consist of chains of final fields or variables and some class. To illustrate how the error in the above example is caught, consider the following:

```

Graph g      = new Graph();
g.Node n    = new g.Node();
ColouredGraph cg = new ColouredGraph();
cg.Node cn  = new cg.Node();

n.connect(cn); // Type Error !!!

```

The type error occurs because the type system is unable to derive a relationship between `g.Node` and `cg.Node`, as it (correctly) cannot derive that `g` and `cg` are the same family.

Note, that in contrast to Jx, .FJ, and Concord, our approach, *vc* and *vObj/Scala* can distinguish nodes coming from different coloured graphs, *e.g.*,

```

ColouredGraph cg1, cg2;
cg1.Node cn1, cn3; cg2.Node cn2;

cn1.connect(cn3); // Type Correct
cn2.connect(cg3); // Type Error!!!

```

*Family Parameters* When writing library code which operates on objects of nested classes, it is sometimes necessary to pass around an object to represent the family. For example, in a *vc*-like setting, in a library outside the `Graph` family, to copy edges, one would have:

```

class Library {
  g.Edge copyEdge(Graph g, g.Edge e) {
    g.Node from = e.from;
    g.Node to   = e.to;
    return new g.Edge(from, to); }
}

```

where the parameter `g` is in some sense superfluous, as it serves no other purpose, than to express the type of the parameter `e`, and the result.

Some languages have types to cater for cases where the family object is not needed. For example, Scala [8] has projection types `Graph#Node`, and Jx [16] has types such as `Graph.Node`. The natural subtype relation  $g.Node \leq Graph.Node$  is valid and has a natural reading: while `g.Node` can be read as saying a `Node` from family `g`, `Graph.Node` can be read as a node from *some* graph family.

Both *vc* and Scala (and Java) have types which are used to refer to the surrounding instance of a given class (from within that class). *vc* uses types of the form `this.out.out.out.C` to refer the surrounding `C` instance (Scala's syntax is different, and closer to Java's). Types are further generalised to have the form `this.out.out.out.f.g.h.C`, which access some path of some surrounding class.

Tribe extends this approach by supporting types that cater for the case where the family object is not needed using the `owner` keyword, and also supports path types to talk about types depending on it. Thus, it can avoid passing superfluous family parameters:

```

class Library{
  int distance(Graph.Node n1, n1.owner.Node n2) { ... }

  e.owner.Edge copyEdge(Graph.Edge e) {
    e.owner.Node from = e.from;
    e.owner.Node to   = e.to;
    new e.owner.Edge(from, to);
  }
}

```

This kind of code gives the benefits of family polymorphism without having to pass the family around. It is also quite robust to change. For example, a method may originally be designed to only require a `Graph.Node`. Subsequent changes may require it to use other features of `Graph`. Without types such as `n1.owner.Node` this possible only by adding an extra argument to pass in a family object.

*Tribe Types* We have shown a number of different kinds of types that appear (and ought to appear) in different systems. All are based on some restriction of the following (where *ff* are final fields including “owner”):

$$T ::= (x \mid \text{this} \mid C).(C \mid \text{ff})^*$$

In this paper we eliminate any restrictions, and make the type system both simpler and more expressive. Following this grammar, Tribe Types have a very natural reading, and subtyping is general and natural, as the following example illustrates. Consider the following code:

```

class Musician {
  class Instrument { }
}
class Guitarist extends Musician {
  final Guitar axe;
  class Guitar extends Instrument {
    final String first, second, third, ...;
    class String { }
  }
}
final Guitarist slash;

```

The following are some of the possible types our system can describe:

Type	Interpretation
<code>slash</code>	ex-Guns 'n' Roses' guitarist Slash
<code>Guitarist</code>	Some guitarist
<code>slash.axe</code>	Slash's main guitar
<code>slash.Guitar</code>	one of Slash's guitars
<code>Guitarist.axe</code>	the main guitar of some guitarist
<code>Guitarist.Guitar</code>	some guitarist's guitar
<code>slash.axe.String</code>	some string of Slash's main guitar
<code>slash.Guitar.first</code>	the first string of one of Slash's guitars
<code>Guitarist.Guitar.String</code>	some string of some guitarist's guitar

The natural reading of types in Tribe also extends in an obvious manner to subtyping. For the above code snippet, the following subtyping relations hold:

$$\begin{array}{ccc}
 \text{slash.} \mathbf{axe} & \leq & \mathbf{Guitarist.} \mathbf{axe} \\
 | \wedge & & | \wedge \\
 \text{slash.} \mathbf{Guitar} & \leq & \mathbf{Guitarist.} \mathbf{Guitar} \\
 | \wedge & & | \wedge \\
 \text{slash.} \mathbf{Instrument} & \leq & \mathbf{Guitarist.} \mathbf{Instrument} \leq \mathbf{Musician.} \mathbf{Instrument}
 \end{array}$$

*Formalising Tribe* Developing sound static type systems for the mechanisms described thus far is a challenge to which many researchers have risen (see citations above). Extending such languages and type systems often requires considerable work due to their subtle nature. For instance, Ernst, Ostermann and Cook’s formalisation of virtual classes was the first account establishing their soundness [7].

This paper presents a formalisation of virtual classes which is both simpler and more general than *vc*. The changes we have made are as follows:

- The notions of type and hence subtyping have been generalised.
- We are non-committal as to exactly which method is dispatched in the case of ambiguity. Requiring that all candidates satisfy the desired typing constraint is sufficient for soundness. Any more precise choice is a *refinement* of our approach, and thus also sound.
- Correct final field initialisation and use is crucial for soundness of path-based types, but dealing with it can significantly increase the complexity of a type system. Rather than cluttering the type system with tests for handling this properly, we simply treat both uninitialised final fields and attempts at reassigning final fields as errors. Dealing with finals properly is known and orthogonal to the rest of the system.
- We present both classes and methods, rather than unify them as in BETA.
- Class tables are statically resolved, whereas in *vc* they are dynamically computed.
- We adopt a small-step semantics rather than a large-step semantics.

The culmination of all of these changes is a simpler formalism and simpler proofs. As a result, we hope that our approach gives a simpler and more general model of the main issues underlying virtual classes, which will serve as the basis for future developments.

We also present a way of dealing with method invocation in the presence of overriding and multiple inheritance, and outline two novel constructs—over-the-top types and adoption—that allow cross-family subclassing while retaining the advantages of virtual classes.

## 2 The Tribe Programming Language

The syntax of Tribe is given in Figure 1, where  $x$  ranges over variables,  $f$  over all field names,  $ff$  over final field names,  $m$  over method names,  $C$  over class

names, and  $\iota$  over addresses. For convenience, all variables are final. At times we use the syntactic category  $ff$  to indicate finality, other times we use the `final` keyword, and sometimes both.

---

<code>class</code>	<code>::= class C extends <math>\bar{C}</math> { <i>cnstr fld* class* mthd*</i> }</code>	CLASSES
<code>cnstr</code>	<code>::= C(<math>\overline{T x}</math>) { <code>this</code>.<i>ff</i> := <i>e</i>; }</code>	CONSTRUCTORS
<code>fld</code>	<code>::= T f;   final T <i>ff</i>;</code>	FIELDS
<code>mthd</code>	<code>::= T <i>md</i> (<math>\overline{T x}</math>) { <i>e</i> }</code>	METHODS
<code>var</code>	<code>::= <code>this</code>   <i>x</i>   <math>\circ</math>   <math>\iota</math></code>	TARGETS
<code>T</code>	<code>::= var.(C   <i>ff</i>)*</code>	TYPES
<code>p</code>	<code>::= var.<i>ff</i>*</code>	PATHS
<code>e</code>	<code>::= null   <span style="background-color: #cccccc;">error</span>   <i>p</i>   <i>e.f</i>   <i>p.f</i> := <i>e</i></code>	EXPRESSIONS
	<code>  <i>p.m</i>(<math>\bar{p}</math>)   new <i>p.C</i>(<math>\bar{p}</math>)   final T <i>x</i> := <i>e</i>; <i>e</i>   <i>e</i>; <i>e</i></code>	

---

Fig. 1. Syntax of Tribe, with run-time entities in grey.

A class consists of a collection of fields, methods and nested classes. Classes inherit from multiple other classes.  $\circ$  denotes the top-level collection of classes.

Fields can either be final or not. Final fields can be used to form paths in order to refer to specific families. Each object has a special field called `owner` (part of the  $ff$  syntax) which refers to the object’s surrounding object (or  $\circ$  for instances of top level classes).

Types in Tribe are formed out of paths, and may, but need not, include class names. Types which do not include a class name are singleton types, referring to a single object. For example, the type of `this` is singleton type `this`, and the type of `owner` is singleton type `this.owner` (effectively, `owner`). A type  $T.C$  describes all the objects of class  $C$  nested inside an object of type  $T$  (a class family). A type  $T.ff$  describes all the objects which may be referred to through the field  $ff$  of an object of type  $T$  (an object family). Run-time types may contain addresses.

Expressions are as usual. Following the syntax from Jx [16], the “let” expression `final T x := e1; e2` introduces the final variable  $x$  with value  $e_1$  which can be used within  $e_2$ . Run-time expressions may contain addresses. Constructors support field initialisation, which is required because of the existence of final fields. In order to keep the language description simple, we do not support calls to the superclass constructor; instead, we allow a constructor to initialise all fields, including inherited fields. For the same reason, we do not statically check that *all* fields have been initialised. As a consequence, uninitialised fields may be accessed in a running program, or final fields may be assigned twice; in both cases we raise errors. As for null pointer exceptions, such errors do not affect type safety. We consider these simplifications to be legitimate, since techniques exist for ensuring that final fields’ initialisation and superclass constructor calls take place exactly once [14, 10, 9, 20].

---

$\frac{\text{(DEF-OBJ)}}{\frac{P(A) \text{ defined}}{\vdash A.\text{Object } cls}}$	$\frac{\text{(DEF-PROG)}}{\frac{\text{class } C \dots \in P(A)}{\vdash A.C \text{ cls}}}$	$\frac{\text{(DEF-INH)}}{\frac{\vdash A' \sqsubseteq_i A \quad \vdash A.C \text{ cls}}{\vdash A'.C \text{ cls}}}$
$\frac{\text{(SUBCL-PROG)}}{\frac{\text{class } C' \text{ extends } \bar{C} \dots \in P(A)}{\vdash A.C' \sqsubseteq_s A.C_i}}$	$\frac{\text{(SUBCL-INH)}}{\frac{\vdash A' \sqsubseteq_i A \quad \vdash A.C' \sqsubseteq_s A.C}{\vdash A'.C' \sqsubseteq_s A'.C}}$	
$\frac{\text{(FURTHER-BIND)}}{\frac{\vdash A \sqsubseteq_i A' \quad \vdash A'.C \text{ cls}}{\vdash A.C \sqsubseteq_f A'.C}}$	$\frac{\text{(INH)}}{\frac{\vdash A \sqsubseteq_s A' \vee \vdash A \sqsubseteq_f A'}{\vdash A \sqsubseteq_i A'}}$	

---

**Fig. 2.** Classes and Subclassing and Further Binding Resolution

### 3 Class Tables

The semantics of Tribe is defined in terms of class tables, CT, which map class identifiers to the members (*i.e.*, fields, methods and nested class constructors) available in the class. In contrast to calculi such as .FJ [11], class tables represent more than the contents of the program. Following approaches used for the semantics of mixins and traits [2, 19], a class table represent a “flattened” version of the program, where a class not only contains the members directly declared in the class itself, but also all the members inherited by the class.

#### 3.1 Subclasses and Further Binding

Classes appear within other classes, and the same identifier may be used to describe classes within several classes, *e.g.*, `Edge` is defined both within `Graph` and within `ColourGraph`. We therefore distinguish between *absolute classes*, which are absolute paths starting from the program’s root, *e.g.*, `o.ColourGraph.Node`, and *classes*, which consist of a class identifier, *e.g.*, `Node`.

The syntax of absolute class names, where `o` corresponds to the root of the program, is:

$$A ::= o \mid A.C$$

We use  $P$  to denote the program text.  $P(A)$  correspond to the code which is inside absolute class  $A$ .  $P(o)$  corresponds to the entire program. To reduce clutter, we assume that  $P$  is global across the type system, and can be used wherever required.

We introduce the following judgements, defined by the rules in Figure 2:

$$\begin{array}{lll} \vdash A \text{ cls} & A \text{ is a class in the program} & \vdash A \sqsubseteq_s A' \quad A \text{ is a direct subclass of } A' \\ \vdash A \sqsubseteq_f A' & A \text{ directly further binds } A' & \vdash A \sqsubseteq_i A' \quad A \text{ directly inherits from } A' \end{array}$$

We demonstrate the above judgements in terms of the following nest of classes, where, for simplicity, we show class nesting but drop class contents:

```

class A
  class B
    class C
      class C' extends C
    class B' extends B
  class A' extends A

```

Rule (DEF-OBJ) defines Object as nested in every class, *e.g.*,  $\vdash \circ.A.B.Object\ cls$ . Facts expressed directly in the program are reflected through rules (DEF-PROG) and (SUBCL-PROG), *e.g.*,  $\vdash \circ.A.B\ cls$ , and  $\vdash \circ.A.B'\ cls$ , and  $\vdash \circ.A' \sqsubseteq_s \circ.A$ . Subclassing implies inheritance (INH), therefore  $\vdash \circ.A' \sqsubseteq_i \circ.A$ . Rules (SUBCL-INH) and (DEF-INH) express that inheritance implies “copying” nested classes and their relationships; thus,  $\vdash \circ.A'.B'\ cls$ , and  $\vdash \circ.A'.B' \sqsubseteq_s \circ.A'.B$ . By repeated application of above rules we obtain, *e.g.*, that  $\vdash \circ.A'.B'.C' \sqsubseteq_s \circ.A'.B'.C$ , even though classes  $\circ.A'.B'.C'$  and  $\circ.A'.B'$  do *not* appear in the program.

The following lemma guarantees that 1) the subclassing relationship holds only between classes which are (derived to be) nested within the same class; 2) further binding implies that the two types are identical up to subclassing at some point in their path; and 3) that inheritance of nested classes implies inheritance of their surrounding classes.

**Lemma 1.**

1. If  $\vdash A' \sqsubseteq_s A$ , then  $\exists A'', C, C'$  such that  $A' = A''.C'$  and  $A = A''.C$ .
2. If  $\vdash A' \sqsubseteq_f A$ , then  $\exists A'', C, C', \bar{C} \neq \epsilon$  such that  $A' = A''.C'.\bar{C}$ ,  $A = A''.C.\bar{C}$ , and  $\vdash A''.C' \sqsubseteq_s A''.C$ .
3.  $\vdash A.C \sqsubseteq_i^* A'.C$  if and only if  $\vdash A \sqsubseteq_i^* A'$  and  $\vdash A'.C\ cls$ .

### 3.2 Class Table Construction

Class tables, CT, map absolute class names  $As$  to tuples  $(fs, ms, cn)$  representing the fields, methods, and nested class constructors that are either directly present in the class or inherited. Class tables are constructed using the operation  $\oplus$  on tuples, which gives priority to the first argument.

**Definition 1** ( $\oplus$ ). For functions  $g$  and  $g'$ , define function  $g \oplus g'$  as:

$$(g \oplus g')(x) = \begin{cases} g(x), & \text{if } g(x) \text{ is defined;} \\ g'(x) & \text{otherwise.} \end{cases}$$

Define  $\oplus$  for tuples as:

$$(fs, ms, cn) \oplus (fs', ms', cn') = (fs \oplus fs', ms \oplus ms', cn \oplus cn')$$

The construction of class tables is defined as follows.

**Definition 2.** For a program  $P$ , we define the class table  $CT$  as

$$\begin{aligned}
BT(\circ) &= (\emptyset, \emptyset, \{cnstr \mid \text{class } \dots \{ cnstr \dots \} \in P(\circ)\}) \\
BT(A.C) &= (flds, mthds, cnstrs) \\
&\quad \text{where class } C \text{ extends } \overline{C'} \{ \dots \text{ classes } flds \text{ mthds } \} \in P(A) \\
&\quad \text{and } cnstrs = \{cnstr \mid \text{class } \dots \{ cnstr \dots \} \in \text{classes}\} \\
CT(A.Object) &= (\{\text{final owner} : \text{this.owner}\}, \emptyset, \emptyset) \\
CT(A) &= BT(A) \oplus CT(\overline{A}), \text{ where } \overline{A} = \{A' \mid \vdash A \sqsubseteq_i A'\}
\end{aligned}$$

The auxiliary function  $BT$  collects the fields and methods directly present in a class, as well as the constructors defined in directly enclosed classes. Note, that the treatment of constructors differs from that of methods and fields. This is so, because for a class  $C$  defined within  $A$ , objects of class  $A.C$  can only be constructed within  $A$  objects; therefore, the  $C$  constructor is in some sense a special method of  $A$ .

The definition of  $CT(A.Object)$  has the effect that every object has a field called `owner` which refers to an object of its surrounding class (even when the class is at the top-level). The definition of  $CT(A)$  collects all the inherited fields, methods and constructors, using those defined in  $BT$  as overriding definitions.

Class tables offer a *flattened view* of the program which directly classes with the members inherited from superclasses or further bound classes. This approach, first used in the study of mixins [2], separates the mechanisms that *produce* inheritance (in our cases subclasses and further binding), from the *effects* of inheritance (*i.e.*, the copying and potential overriding of members. This simplifies the model considerably, simplifies the proofs, and demonstrates easily how Tribe could be implemented.

*Remark 1.* Observe that we are not too concerned with the order in which methods are inherited. Our system ensures that *all* inherited methods work soundly; the decision about which one to use is an orthogonal issue to soundness.

The following shorthands will be used throughout the remainder of the paper.

**Definition 3.** Given a class table  $CT$ , define the following shorthands:

- $\text{fields}(A) = fst(CT(A))$ .
- $\text{finals}(A) = \{ff : T \mid \text{final } ff : T \in fst(CT(A))\}$ .
- $\text{methods}(A) = snd(CT(A))$ .
- $\text{constructors}(A) = thd(CT(A))$ .
- $(fs, ms, cn) \sqsubseteq (fs', ms', cn')$  iff there exist  $fs'', ms'',$  and  $cn''$  such that  $fs = fs'' \oplus fs', ms = ms'' \oplus ms',$  and  $cn = cn'' \oplus cn'$ .

Finally, define  $\sqsubseteq_i^*$  as the reflexive transitive closure of  $\sqsubseteq_i$ . The following lemma guarantees that 1) class table entries exist only for well formed absolute classes, 2) each class table entry has an `owner` field of appropriate type, inheritance implies that the class table entry expands that of the inherited class.

---

$\mathbf{v} ::= \circ \mid \iota$	NON-NULL VALUES
$v ::= \mathbf{v} \mid \mathbf{null}$	VALUES
$E ::= [-] \mid E.f \mid E.f := e \mid \mathbf{v}.f := E$	REDUCTION CONTEXT
$\mid E.m(\bar{e}) \mid \mathbf{v}.m(\bar{v}, E, \bar{e}) \mid \mathbf{new} E.C(\bar{e})$ $\mid \mathbf{new} \mathbf{v}.C(\bar{v}, E, \bar{e}) \mid E; e$	
$err ::= \mathbf{null} \mid \mathbf{error}$	ERROR VALUES
$N ::= err.f \mid err.f := e \mid err.ff := e$	ERROR CONTEXTS
$\mid v.ff := \mathbf{error} \mid err.m(\bar{e}) \mid \mathbf{new} err.C(\bar{e})$	

---

**Fig. 3.** Dynamic Expressions

**Lemma 2.** *The following properties hold for a class table CT:*

1.  $CT(A) \neq \perp$  if and only if  $\vdash A$  cls.
2. If  $CT(A) \neq \perp$ , then  $\mathbf{Object} \in \mathit{classes}(A)$  and  $\mathbf{owner} : \mathbf{this.owner} \in \mathit{finals}(A)$ .
3. If  $\vdash A \sqsubseteq_i^* A'$ , then  $CT(A) \sqsubseteq CT(A')$ .

## 4 Dynamic Semantics

The semantics of Tribe is presented as a small-step reduction relation of the form  $H, e \rightsquigarrow H', e'$ , which states that configuration  $H, e$  reduces (in one step) to  $H', e'$ , where  $H$  and  $H'$  are heaps and  $e$  and  $e'$  are expressions. The syntax of heaps  $H$ , and objects  $o$ :

$$H ::= \emptyset \mid \iota \mapsto o, H \qquad o ::= [\bar{f} \mapsto \bar{v}]_A$$

An object consists of the name of its absolute class,  $A$ , and values for all of its fields, denoted  $\bar{f} \mapsto \bar{v}$ . The fields also contain a value for the **owner** field.

Define  $H(\iota)$ ,  $H + \iota \mapsto o$  and  $H[\iota \mapsto o]$  as look-up, extend, and update of a heap, and analogously for  $o$ . Also let  $H(\iota).f$  denote  $H(\iota)(f)$  and  $H(\iota).f := v$  denote  $H[\iota \mapsto (H(\iota)[f \mapsto v])]$ , as shorthands for accessing and updating the field of some object in the heap.

Figure 3 contains additional syntax used in the operational semantics. A reduction context,  $E[-]$ , is defined to be an expression with a hole in it, in the standard manner. We also have *error contexts*,  $N$ , (called null contexts in Jx's semantics [16]) to gracefully handle dereferencing of **null**. For convenience a **null-dereference** evaluates immediately to **error**. We apply the same treatment to errors arising from assigning to already initialised final fields and to errors resulting from accessing a final field which has not yet been initialised. Such expressions are not stuck, and the value they reduce to can have any type, thus their presence does not interfere with subject reduction and progress lemmas.

The rules for the operational semantics are given in Figure 4.

(EVAL-FIELD) returns the value of an object's field. Both (EVAL-FIELDASGN) and (EVAL-FINALFIELDASGN) update fields. (EVAL-FINALFIELDASGN) is only applicable if the field contains **error**, indicating that it is uninitialised. (EVAL-FINALFIELDASGN-ERROR) traps the case when an attempt to initialise an already

---

$\frac{\text{(EVAL-FIELDASGN)} \quad \begin{array}{l} H(\iota) = [\dots]_{\mathbf{A}} \\ f \in \text{nonfinals}(\mathbf{A}) \\ H' = H(\iota).f := v \end{array}}{H, \iota.f := v \rightsquigarrow H', v}$	$\frac{\text{(EVAL-FINALFIELDASGN)} \quad \begin{array}{l} H(\iota) = [\dots]_{\mathbf{A}} \quad \overline{ff} \in \text{finals}(\mathbf{A}) \\ H(\iota).\overline{ff} = \mathbf{error} \\ H' = H(\iota).f := v \end{array}}{H, \iota.\overline{ff} := v \rightsquigarrow H', v}$	$\frac{\text{(EVAL-FINALFIELDASGN-ERROR)} \quad \begin{array}{l} H(\iota) = [\dots]_{\mathbf{A}} \\ \overline{ff} \in \text{finals}(\mathbf{A}) \\ H(\iota).\overline{ff} \neq \mathbf{error} \end{array}}{H, \iota.\overline{ff} := v \rightsquigarrow H, \mathbf{error}}$	
$\frac{\text{(EVAL-NEW)} \quad \begin{array}{l} \mathbf{v} = \circ = \mathbf{A} \vee \mathbf{v} = \iota' \wedge H(\iota') = [\dots]_{\mathbf{A}} \\ \mathbf{C}(\overline{T} \overline{x})\{e\} \in \text{constructors}(\mathbf{A}) \\ \text{dom}(\text{nonfinals}(\mathbf{A}.\mathbf{C})) = \overline{f} \quad \text{dom}(\text{finals}(\mathbf{A}.\mathbf{C})) = \overline{ff} \\ \iota \notin \text{dom}(H) \quad H' = H + \iota \mapsto [\mathbf{owner} \mapsto \mathbf{v}, \overline{f} \mapsto \mathbf{null}, \overline{ff} \mapsto \mathbf{error}]_{\mathbf{A}.\mathbf{C}} \end{array}}{H, \mathbf{new} \ \mathbf{v}.\mathbf{C}(\overline{v}) \rightsquigarrow H', e[\iota/\mathbf{this}, \overline{v}/\overline{x}]; \iota}$			
$\frac{\text{(EVAL-METH)} \quad \begin{array}{l} H(\iota) = [\dots]_{\mathbf{A}} \quad T \ m(\overline{T} \ \overline{z})\{e\} \in \text{methods}(\mathbf{A}) \end{array}}{H, \iota.m(\overline{v}) \rightsquigarrow H, e[\iota/\mathbf{this}, \overline{v}/\overline{z}]}$		$\frac{\text{(EVAL-LET)} \quad \begin{array}{l} H, \mathbf{final} \ T \ x := v; e \rightsquigarrow H, e[v/x] \end{array}}{H, \mathbf{final} \ T \ x := v; e \rightsquigarrow H, e[v/x]}$	
$\frac{\text{(EVAL-FIELD)} \quad \begin{array}{l} H(\iota).f = v \end{array}}{H, \iota.f \rightsquigarrow H', v}$	$\frac{\text{(EVAL-SEQ)} \quad \begin{array}{l} H, v; e \rightsquigarrow H, e \end{array}}{H, v; e \rightsquigarrow H, e}$	$\frac{\text{(EVAL-CONTEXT)} \quad \begin{array}{l} H, e \rightsquigarrow H', e' \end{array}}{H, E[e] \rightsquigarrow H', E[e']}$	$\frac{\text{(EVAL-ERROR)} \quad \begin{array}{l} H, E[N] \rightsquigarrow H, \mathbf{error} \end{array}}{H, E[N] \rightsquigarrow H, \mathbf{error}}$

---

**Fig. 4.** Reduction Rules

initialised final field is made. (EVAL-NEW) determines firstly which class table to look-up to find the constructors. This depends upon the path for which the new class is being created—the new owner. A new object is created, with the owner field set appropriately, all non-final fields set to `null`, and all final fields set to `error` to indicate that they are not initialised. The result is an expression which will evaluate the constructor and return the new location. (EVAL-METH) finds the code for the method in the class table, and reduces to the body with the targets and arguments substituted. (EVAL-LET) is standard let statement. (EVAL-SEQ) is standard sequential composition, which discards the result of the first expression. (EVAL-CONTEXT) is also standard, stating that the evaluation of any expression proceeds by evaluating one of its redexes. (EVAL-ERROR) detects an error condition and immediately reduces to an error.

## 5 Type System

The type system is defined in terms of the following judgements:

$$\begin{array}{llll} \Gamma \vdash \diamond & \Gamma \text{ is a good typing environment} & \Gamma \vdash T & T \text{ is a good type} \\ \Gamma \vdash T \uparrow \mathbf{A} & \text{Class } \mathbf{A} \text{ corresponds to type } T & \Gamma \vdash T \leq T' & T \text{ is a subtype of } T' \\ \Gamma \vdash e : T & \text{expression } e \text{ has type } T & & \end{array}$$

Typing environments have the following syntax:

$$\Gamma ::= \emptyset \mid \mathbf{A} \mid \Gamma, x : T \mid \Gamma, \iota : T \mid \Gamma, \iota.\overline{ff} = v$$

---


$$\begin{array}{c}
\text{(ENV-EMPTY)} \qquad \text{(ENV-ABSCCLASS)} \qquad \text{(ENV-DECL)} \\
\frac{}{\emptyset \vdash \diamond} \qquad \frac{\vdash A \text{ cls}}{A \vdash \diamond} \qquad \frac{\Gamma \vdash T \quad x \notin \text{dom}(\Gamma)}{\Gamma, x : T \vdash \diamond} \\
\hline
\text{(ENV-EQ)} \\
\frac{\iota.\text{ff} = v' \notin \Gamma \quad \Gamma \vdash \iota \uparrow A \quad \text{ff} : T \in \text{finals}(A) \quad \Gamma \vdash v : T[\iota/\text{this}]}{\Gamma, \iota.\text{ff} = v \vdash \diamond}
\end{array}$$


---

**Fig. 5.** Good Environments—(ENV-DECL) applies to both variables and addresses

A typing environment maps variables, including the receiver `this`, and addresses to types. The absolute class name  $A$  in a typing environment indicates the context in which type checking is performed. Equations of the form  $\iota.\text{ff} = v$  in a typing environment keep track of paths; they are used to track paths when proving soundness, and are *not* used in the static semantics.

Figure 5 defines well-formed environments,  $\Gamma \vdash \diamond$ . The rules are mostly straightforward. The non-obvious rule, (ENV-EQ), is applicable for run-time environments, and enables an equation to be added to the typing environment, if no equation is already present for the given object field. It simply requires that  $v$  is a value which can be stored in final field  $\text{ff}$  of object  $\iota$ .

### 5.1 Well-formed Types

A well-formed type  $T$  corresponds to an absolute class,  $A$ , as expressed through the judgement  $\Gamma \vdash T \uparrow A$  in Figure 6. Only types which correspond to absolute classes are well formed, from rule (TYPE).

---


$$\begin{array}{c}
\text{(TYPE-START)} \qquad \text{(TYPE-DECL)} \qquad \text{(TYPE-THIS)} \\
\frac{\Gamma \vdash \diamond}{\Gamma \vdash \circ \uparrow \circ} \quad \frac{x : T \in \Gamma}{\Gamma \vdash x \uparrow A} \quad \frac{x \neq \text{this} \quad \Gamma \vdash T \uparrow A}{\Gamma \vdash x \uparrow A} \quad \frac{\Gamma \vdash \diamond \quad A, \text{this} : \text{this} \in \text{dom}(\Gamma)}{\Gamma \vdash \text{this} \uparrow A} \\
\hline
\text{(TYPE-FINAL-FIELD)} \\
\frac{\Gamma \vdash T \uparrow A \quad \text{ff} : T' \in \text{finals}(A) \quad A, \text{this} : \text{this} \vdash T' \uparrow A'}{\Gamma \vdash T.\text{ff} \uparrow A'} \\
\hline
\text{(TYPE-CLASS)} \qquad \text{(TYPE-OWNER)} \qquad \text{(TYPE)} \\
\frac{\Gamma \vdash T \uparrow A \quad \vdash A.C \text{ cls}}{\Gamma \vdash T.C \uparrow A.C} \quad \frac{\Gamma \vdash T \uparrow A.C}{\Gamma \vdash T.\text{owner} \uparrow A} \quad \frac{\Gamma \vdash T \uparrow A}{\Gamma \vdash T}
\end{array}$$


---

**Fig. 6.** Good Types

Rule (TYPE-START) allows  $\circ$  as the starting point for absolute types. Rule (TYPE-DECL) states that variables and addresses can be used as singleton types.

---

$\frac{\text{(SUB-REFL)}}{\Gamma \vdash T \leq T}$	$\frac{\text{(SUB-TRANS)}}{\Gamma \vdash T \leq T' \quad \Gamma \vdash T' \leq T'' \quad \Gamma \vdash T \leq T''}$	$\frac{\text{(SUB-SUBCLASS)}}{\Gamma \vdash T \uparrow A \quad \vdash A.C \sqsubseteq_s A.C' \quad \Gamma \vdash T.C \leq T'.C'}$
$\frac{\text{(SUB-DECL)}}{x : T \in \Gamma \quad \Gamma \vdash x \leq T}$	$\frac{\text{(SUB-FINAL-FIELD)}}{\Gamma \vdash T \uparrow A \quad ff : T' \in \text{finals}(A) \quad \Gamma \vdash T.ff \leq T'[T/\text{this}]}$	$\frac{\text{(SUB-ABS)}}{\Gamma \vdash T \uparrow A \quad \Gamma \vdash T \leq A}$
$\frac{\text{(SUB-OWNER-1)}}{\Gamma \vdash T.C.\text{owner} \quad \Gamma \vdash T.C.\text{owner} = T}$	$\frac{\text{(SUB-OWNER-2)}}{\Gamma \vdash T \uparrow A.C \quad \Gamma \vdash T \leq T.\text{owner}.C}$	$\frac{\text{(SUB-OWNER-3)}}{\Gamma \vdash T \leq T'.\text{owner}.C \quad \Gamma \vdash T.\text{owner} = T'.\text{owner}}$
$\frac{\text{(SUB-NEST-FINAL)}}{\Gamma \vdash T \leq T' \quad \Gamma \vdash T'.ff \quad \Gamma \vdash T.ff \leq T'.ff}$		$\frac{\text{(SUB-NEST-CLASS)}}{\Gamma \vdash T \leq T' \quad \Gamma \vdash T'.C \quad \Gamma \vdash T.C \leq T'.C}$

---

**Fig. 7.** Good Subtyping— $T = T'$  is interpreted as  $T \leq T'$  and  $T' \leq T$ .

The absolute class to which a variable or address corresponds is determined using its declared type. Rule (TYPE-THIS) declares `this` to be a singleton type, which corresponds to the absolute class given by the present context. By rule (TYPE-CLASS), any valid type for which it makes sense to have a nested class `C` can be extended with suffix `C`. Similarly, rule (TYPE-FINAL-FIELD) enables any type which contains a final field `ff` to be extended with suffix `ff`. (TYPE-OWNER) states that any type can be extended with suffix `owner`, as long as the type corresponds to an actual absolute class, *not*  $\circ$ .

Note that we not only have more types than *vc*, as described in the introduction. We even permit types to go outside of a hierarchy. For example, in a topmost class `A` the type `this.owner.B` refers to an instance of the topmost class `B`, even though the two classes are not enclosed in another class. This is useful when using adoption (Section 7).

We can prove that 1) a type corresponds to at most one absolute class; 2) if a type corresponds to an absolute class, then this absolute class is a class (possibly inferred through inheritance); and 3) types correspond to more specialised absolute classes in more specialised contexts.

**Lemma 3.** *For program  $P$ , environment  $\Gamma$ , type  $T$ , absolute classes  $A_0, A, A'$ :*

1. *If  $\Gamma \vdash T \uparrow A$  and  $\Gamma \vdash T \uparrow A'$ , then  $A = A'$ .*
2. *If  $\Gamma \vdash T \uparrow A$ , then  $\vdash A$  cls.*
3. *If  $\vdash A_1 \sqsubseteq_i^* A_2, A_1, \Gamma \vdash T \uparrow A_3$ , and  $A_2, \Gamma \vdash T \uparrow A_4$ , then,  $\vdash A_3 \sqsubseteq_i^* A_4$ .*

## 5.2 Subtyping

Subtyping in Tribe is very rich. The rules are presented in Figure 7.

As usual, the type system supports reflexivity (SUB-REFL) and transitivity (SUB-TRANS) of the subtype relation. (SUB-SUBCLASS) converts subclassing into

subtyping. (SUB-DECL) states that a singleton type which corresponds to either a variable or an address is a subtype of its declared type. (SUB-ABS) establishes a subtype relation between any type and its absolute class.

(SUB-FINAL-FIELD) establishes a relationship between a type with a final field suffix and the type of the field in the given context. This is achieved by substituting the prefix into the type of the final field. For example, if  $\Gamma \vdash T \uparrow \circ.A.B$  and class  $\circ.A.B$  has final fields  $ff_1 : \circ.C.D$  and  $ff_2 : \text{this}.E$ , then (SUB-FINAL-FIELD) gives that  $\Gamma \vdash T.ff_1 \leq \circ.C.D$  and  $\Gamma \vdash T.ff_2 \leq T.E$ .

Rule (SUB-OWNER-1) is two rules combined into one statement. The equality is to be read both left to right and right to left as a subtype rule. A type can have  $C.\text{owner}$  added or removed from the end (whenever it produces a sensible type). This works since  $C$  is a nested type, and  $\text{owner}$  goes back to the enclosing instance. Rule (SUB-OWNER-2) allows any type to be given in terms of a type in the surrounding class. This is essential for obtaining that  $\text{this} : \text{this}.\text{owner}.C$ , for the appropriate  $C$ . Rule (SUB-OWNER-3) is again really two rules. This enables one to determine whether two types reside in the same family, in which case, their owners are the same.

Both rules (SUB-NEST-FINAL) and (SUB-NEST-CLASS) enable a prefix of a type to be replaced by any subtype. Combined with transitivity, these two rules enable any internal part of a type path to be replaced by its subtype. Rule (SUB-NEST-CLASS) resembles the rule ( $\leq$ -NEST) from Jx [16], with an additional well-formedness check.

The following lemma states that 1) subtyping of absolute classes implies inheritance; 2) the absolute class corresponding to a type is its most specific absolute supertype; 3) the absolute classes corresponding to subtypes are in the inheritance relationship, meaning that everything expected of some type is available in all subtypes (crucial for proving progress); 4) subtype relations are preserved in more specialised contexts. The last property means that subtype relations remain true in all *inherited* code. A similar property applies to typing, as we will see in the next section. Such preservation properties are crucial for proving soundness.

**Lemma 4.**

1. If  $\Gamma \vdash A \leq A'$ , then  $\vdash A \sqsubseteq_i^* A'$ .
2. If  $\Gamma \vdash T \uparrow A$ , and  $\Gamma \vdash T \leq A'$ , then  $\Gamma \vdash A \leq A'$ .
3. If  $\Gamma \vdash T \leq T'$ ,  $\Gamma \vdash T \uparrow A$ , and  $\Gamma \vdash T' \uparrow A'$ , then  $\vdash A \sqsubseteq_i^* A'$ .
4. If  $\vdash A_1 \sqsubseteq_i^* A_0$  and  $A_0, \Gamma \vdash T \leq T'$ , then  $A_1, \Gamma \vdash T \leq T'$ .

### 5.3 Expression Typing

Expression typing is given in Figure 8.

(EXPR-BULLET) gives to  $\circ$  the type  $\circ$  so that it  $\circ$  be used in expressions which create absolute types. (EXPR-DECL) gives a singleton type to any variable or address (as they are final). Through a combination of (SUB-DECL) and (EXPR-SUBSUMPTION) we obtain, for example, that if  $x : T \in \Gamma$ , then  $\Gamma \vdash x : T$ ,

---

$\frac{\text{(EXPR-BULLET)}}{\Gamma \vdash \diamond}$	$\frac{\text{(EXPR-DECL)}}{\Gamma \vdash \diamond \quad x : T \in \Gamma}$	$\frac{\text{(EXPR-NULL)}}{\Gamma \vdash T}$	$\frac{\text{(EXPR-ERROR)}}{\Gamma \vdash T}$
$\frac{\Gamma \vdash \circ : \circ}{\Gamma \vdash \circ : \circ}$	$\frac{\Gamma \vdash x : x}{\Gamma \vdash x : x}$	$\frac{\Gamma \vdash \mathbf{null} : T}{\Gamma \vdash \mathbf{null} : T}$	$\frac{\Gamma \vdash \mathbf{error} : T}{\Gamma \vdash \mathbf{error} : T}$
$\frac{\text{(EXPR-FINAL-FIELD)}}{\Gamma \vdash e : T \quad \Gamma \vdash T \uparrow A \quad \mathbf{ff} : T' \in \mathbf{finals}(A)}{\Gamma \vdash e.\mathbf{ff} : T.\mathbf{ff}}$	$\frac{\text{(EXPR-FIELD)}}{\Gamma \vdash e : T \quad \Gamma \vdash T \uparrow A \quad f : T' \in \mathbf{nonfinals}(A)}{\Gamma \vdash e.f : T'[T/\mathbf{this}]}$	$\frac{\text{(EXPR-FIELDASGN)}}{\Gamma \vdash p : T \quad \Gamma \vdash T \uparrow A \quad [\mathbf{final}] f : T' \in \mathbf{fields}(A)}{\Gamma \vdash e : T'[p/\mathbf{this}]}$	
$\frac{\text{(EXPR-NEW)}}{\Gamma \vdash p_0 \uparrow A \quad C(\overline{T \ x})\{\dots\} \in \mathbf{constructors}(A) \quad \Gamma \vdash \overline{p} : \overline{T}[\overline{p}/\overline{x}]}{\Gamma \vdash \mathbf{new} \ p_0.C(\overline{p}) : p_0.C}$			
$\frac{\text{(EXPR-CALL)}}{\Gamma \vdash p_0 \uparrow A \quad T \ m(\overline{T \ x})\{\dots\} \in \mathbf{methods}(A) \quad \Gamma \vdash \overline{p} : \overline{T}[p_0/\mathbf{this}][\overline{p}/\overline{x}]}{\Gamma \vdash p_0.m(\overline{p}) : T[p_0/\mathbf{this}][\overline{p}/\overline{x}]}$			
$\frac{\text{(EXPR-LET)}}{\Gamma \vdash e : T \quad \Gamma, x : T \vdash e' : T' \quad \Gamma \vdash T'}{\Gamma \vdash \mathbf{final} \ T \ x := e; e' : T'}$	$\frac{\text{(EXPR-SEQ)}}{\Gamma \vdash e : T \quad \Gamma \vdash e' : T'}{\Gamma \vdash e; e' : T'}$	$\frac{\text{(EXPR-SUBSUMPTION)}}{\Gamma \vdash e : T \quad \Gamma \vdash T \leq T'}{\Gamma \vdash e : T'}$	

---

**Fig. 8.** Expression Typing

as expected. (EXPR-NULL) is standard, giving any type to **null**. Rule (EXPR-ERROR) is similar. It is used to enable the subject reduction result to go through in the presence of errors.

(EXPR-FIELD) is as expected, giving the type of the field relative to the type of the object whose field it is (since the field type may contain **this**). (EXPR-FINAL-FIELD) instead extends the type with a singleton type, though whether the resulting type is singleton depends upon the type  $T$ . For example, from  $\Gamma \vdash x : x$  we might obtain that  $\Gamma \vdash x.\mathbf{ff} : x.\mathbf{ff}$ , which is a singleton type, but from  $\Gamma \vdash x : y.D$  we obtain  $\Gamma \vdash x.\mathbf{ff} : y.D.\mathbf{ff}$ , which is not a singleton type.

(EXPR-FIELDASGN) handles field assignment. For precise typing, the target containing the field must be a path. This path is then used to give a precise type to the value assigned to the field. This rule permits both assignment of both final and non-final fields. This is for uniformity, since final fields need to be initialised in the constructor. Attempts at assigning an initialised final field or attempts at accessing non-initialised fields result in errors, as previously discussed.

(EXPR-NEW) describes the type of object creation. The arguments  $\overline{p}$  must have the types required by the constructor, where any appearance of the formal parameter has been updated by the actual paths. Note that rule (EXPR-CALL) is similar to (EXPR-NEW) in the treatment of the arguments. In fact,  $vc$  unifies the two (following **BETA** and **gBeta**), eliminating methods.

The rules (EXPR-NEW) and (EXPR-METHOD) may seem somewhat odd in that they don't give a type to the receiver. Observe the following lemma:

**Lemma 5.** *If  $\Gamma \vdash p \uparrow A$ , then  $\Gamma \vdash p : p$ .*

---


$$\begin{array}{c}
\text{(GOOD-PROGRAM)} \\
\frac{\forall \text{class} \in P(\circ). \circ \vdash \text{class} \sqsubseteq_i^+ \text{ is acyclic}}{\vdash P}
\end{array}
\qquad
\begin{array}{c}
\text{(GOOD-METHOD)} \\
\frac{\Gamma, \bar{x} : \bar{T} \vdash e : T}{\Gamma \vdash T \text{ md}(T \ x)\{ e \}}
\end{array}$$

$$\begin{array}{c}
\text{(GOOD-CLASS)} \\
\frac{\begin{array}{l}
\Gamma = \text{A.C, this : this} \quad \text{sup} = \{ \text{A.C} \mid \text{C} \in \bar{C}' \} \quad \bar{\text{A}} = \text{sup} \cup \{ \text{A}' \mid \vdash \text{A.C} \sqsubseteq_f \text{A}' \} \\
\vdash \text{A cls for each } \text{A} \in \text{sup} \quad \forall f : T \in \text{flds}. \Gamma \vdash T \quad \text{A.C} \vdash \text{classes} \quad \Gamma \vdash \text{mthds} \\
\forall m, m' \in (\text{mthds} \cup \text{methods}(\bar{\text{A}})). \text{name}(m) = \text{name}(m') \implies \text{type}(m) = \text{type}(m') \\
\forall f. f : T, f : T' \in (\text{flds} \cup \text{fields}(\bar{\text{A}})) \implies T = T' \quad \Gamma, x : \bar{T} \vdash e : T' \text{ for some } T' \\
\forall \text{A}' : \vdash \text{A} \sqsubseteq_i \text{A}' \text{ and } \text{C}(\bar{T}' \ x')\{ \dots \} \in \text{constructors}(\text{A}') \implies \bar{T}' = \bar{T}
\end{array}}{\text{A} \vdash \text{class C extends } \bar{C}' \{ \text{C}(\bar{T}' \ x)\{ e \} \text{ flds classes mthds } \}}
\end{array}$$


---

**Fig. 9.** Program, Class and Method Typing.

With the above lemma, the additional premise  $\Gamma \vdash p_0 : p_0$  could be redundantly added to the premises of (EXPR-NEW), thus giving the target a type.

(EXPR-LET) is for “let” expressions. It is standard, including the last premise to ensure that variable  $x$ , which could occur in the resulting type, does not escape its scope. This is standard in dependent typing [1]. (EXPR-SEQ) handles sequential composition in a straightforward fashion. Finally, (EXPR-SUBSUMPTION) is the standard subsumption rule.

#### 5.4 Well-formed Programs and Classes

Figure 9 defines well-formed programs and classes.

A program is well-formed if all its top level classes are well-formed, and there are no cyclic dependencies in the inheritance hierarchy.

A class is well-formed if its super classes are all valid classes *in the present context*, the fields have types that are well-formed in the current context, all methods and nested classes are well formed. In addition, all methods of the same name which are either declared in the present class or inherited *must* have the same type (where *name* gets a method’s name and *type* gets a method’s signature). Any inherited fields or local fields with the same name must have the same type; this last constraint simplifies matters.<sup>5</sup> Finally, the constructor must have the same type as all inherited constructors.

A method is well-defined in the usual sense, though the types of arguments may depend upon the path of other arguments and on the path of the receiver.

Finally, we will give some further results about well-formed programs. Define the relation  $(fs, ms, cn) \leq (fs', ms', cn')$  to express that the members  $(fs, ms, cn)$  enhance those from  $(fs', ms', cn')$ , while preserving types.

<sup>5</sup> Another approach would be to have a compiler phase which annotated field use with the class containing their definition, as is done by Java and C++ compilers. The description of such an approach is not difficult, but we consider this matter orthogonal to soundness.

**Definition 4.** Define  $(fs, ms, cn) \leq (fs', ms', cn')$  iff the following four conditions hold:

- $(fs, ms, cn) \sqsubseteq (fs', ms', cn')$ ;
- If  $f : T \in fs'$  and  $f : T \in fs$ , then  $T' = T$ ;
- If  $m \in \text{names}(ms) \cap \text{names}(ms')$ , then  $\text{type}(ms(m)) = \text{type}(ms'(m))$ ; and
- If  $\mathcal{C}(\overline{T}x)\{ \dots \} \in cn$  and  $\mathcal{C}(\overline{T}'x)\{ \dots \} \in cn'$ , then  $\overline{T} = \overline{T}'$ .

The following lemma guarantees that in well formed programs 1) expressions preserve their type in more specialised contexts; 2) class table entries of inheriting classes enhance those of the inherited class while preserving types, and 3) method defined in or inherited by an absolute class is well formed in that class context. The last guarantee is stronger than what is required by rule (GOOD-CLASS), because the rule only guarantees well-formedness for methods directly defined in the class.

**Lemma 6.** For any well-formed program and its associated class table CT:

- If  $\vdash A_1 \sqsubseteq_i^* A_0$  and  $A_0, \Gamma \vdash e : T$ , then  $A_1, \Gamma \vdash e : T$ .
- If  $\vdash A \sqsubseteq_i^* A'$ , then  $\text{CT}(A) \leq \text{CT}(A')$ .
- If  $T \text{ md}(\overline{T}x)\{ e \} \in \text{methods}(A)$ , then  $A, \text{this} : \overline{T}, x : \overline{T} \vdash e : T$ .

## 6 Soundness of the Type System

The soundness of the Tribe type system is established using the standard technique of providing subject reduction and progress theorems due to Wright and Felleisen [21]. We now present a number of auxiliary notions and the rules for well-defined heaps.

In order to type run-time expressions, we need to extract a typing environment from a heap. This information includes the types of addresses and equations involving the values of final fields. This is achieved using the following definition:

**Definition 5.** Given a heap  $H$ , we define operation  $H^\gamma$  as follows:

$$\begin{aligned} \emptyset^\gamma &= \emptyset \\ (H, \iota \mapsto [fds]_{A.C})^\gamma &= H^\gamma, \iota : \mathbf{v}.C, \iota.\text{ff}_1 = v_1, \dots, \iota.\text{ff}_n = v_n \\ &\text{where } H(\iota).\text{owner} = \mathbf{v} \text{ and} \\ &fds = \{ \overline{\text{ff}} \mapsto v, \overline{\text{ff}'} \mapsto \text{error}, \overline{f} \mapsto - \}, v_i \neq \text{error} \end{aligned}$$

Note that only equations on *defined* final fields get recorded in  $H^\gamma$ . Uninitialised final fields, those containing **error**, do not contribute to the typing environment.

$H^\gamma$  has the same shape as an environment  $\Gamma$ , so the judgements from the static semantics can be reused to type dynamic expressions. Rather than *explicitly* inserting the extraction operation, we use it *implicitly*. Thus, whenever we write, for example,  $H \vdash T \uparrow A$ , we really mean  $H^\gamma \vdash T \uparrow A$ . This applies to all rule shapes.

---


$$\begin{array}{c}
\text{(OBJECT)} \\
\frac{\text{fields}(A) = \bar{f} : \bar{T} \quad \text{fds} = \bar{f} \mapsto \bar{v} \quad H \vdash \bar{v} : T[\iota/\text{this}] \quad \text{fds}(\text{owner}) \neq \text{error}}{H \vdash \iota \mapsto [\text{fds}]_A} \\
\\
\begin{array}{ccc}
\text{(EXPR-EQ)} & \text{(HEAP)} & \text{(CONFIG)} \\
\frac{\Gamma \vdash \iota' : T \quad \iota'.f = v \in \Gamma}{\Gamma \vdash v : T.f} & \frac{\forall \iota \mapsto o \in H : H \vdash \iota \mapsto o}{\vdash H} & \frac{\vdash H \quad H \vdash e : T}{\vdash H, e : T}
\end{array}
\end{array}$$


---

**Fig. 10.** Heap Typing

## 6.1 Well-formed heaps

The type rules for heaps, objects and configurations are given Figure 10.

(OBJECT) requires that all the fields have the correct type with respect to some absolute class and the present object id. The fact that the object id appears in the types takes care of the family requirements. In particular, the owner is a final field, and therefore, the type rules require that it should have type  $\text{this.owner}[\iota/\text{this}]$ , *i.e.*, for  $H(\iota).\text{owner} = \mathbf{v}$ , we require  $H^\gamma \vdash \mathbf{v} : \iota.\text{owner}$ , which is satisfied from the definition of  $H^\gamma$  and rule (EXPR-EQ).

(EXPR-EQ) is an additional rule used to establish equivalences between a type involving a path and object to which the path refers. (HEAP) states that a heap is well-formed whenever all of its objects are well-formed. (CONFIG) states that a well-formed heap and expression typed against it form a configuration with the type of the expression.

We have the following properties of well-formed heaps: 1) a well-formed heap can act as a well-formed environment; 2) the class corresponding to an address (as a singleton type) is the class of that address; 3) the surrounding class of the object surrounding an object is the surrounding class of the object's class; and 4) the chains of owners forms a finite tree rooted at  $\circ$ .

**Lemma 7.** *Assume  $\vdash H$ . Then:*

1.  $H^\gamma \vdash \diamond$ .
2.  $H^\gamma \vdash \iota \uparrow A$  if and only if  $H(\iota) = [\dots]_A$ .
3. Assume  $H(\iota).\text{owner} = \mathbf{v}$ . If  $H^\gamma \vdash \iota \uparrow A.C$ , then  $H^\gamma \vdash \mathbf{v} \uparrow A$ .
4. The graph  $G = (\text{dom}(H) \cup \{\circ\}, \{(\iota, \mathbf{v}) \mid \iota \in \text{dom}(H) \wedge H(\iota).\text{owner} = \mathbf{v}\})$  is a tree with root  $\circ$ .

## 6.2 Soundness

We can prove subject reduction and progress properties. We first state our variants of the standard meta-theory required to establish these results.

**Lemma 8.** *If  $A, \text{this} : \text{this}, \Gamma \vdash e : T$  and  $H \vdash \iota \uparrow A$ , then  $H, \Gamma[\iota/\text{this}] \vdash e[\iota/\text{this}] : T[\iota/\text{this}]$ .*

**Lemma 9 (Substitution).** *If  $\Gamma, x : T, \Gamma' \vdash e : T'$ , and  $\Gamma \vdash v : T$ , then  $\Gamma, \Gamma'[v/x] \vdash e[v/x] : T'[v/x]$ .*

**Definition 6 (Environment Extension, Heap Extension).**

- Define  $\Gamma \subseteq \Gamma'$  iff  $\exists \Gamma''$  with  $\Gamma' = \Gamma, \Gamma''$ .
- Define  $H \subseteq H'$  iff  $H^\gamma \subseteq H'^\gamma$ .

Note, that  $H \subseteq H'$  implies that  $\text{dom}(H) \subseteq \text{dom}(H')$ ,  $\forall \iota \in \text{dom}(H). \text{type}(H(\iota)) = \text{type}(H'(\iota))$ , and that final fields are more defined in  $H'$  than in  $H$ . The last point means that final fields containing **error** in  $H$  may become defined in  $H'$ , but already defined final fields must preserve their contents.

**Lemma 10 (Extension).** *If  $\Gamma \vdash e : T$ ,  $\Gamma \subseteq \Gamma'$  and  $\Gamma' \vdash \diamond$ , then  $\Gamma' \vdash e : T$ .*

**Lemma 11 (Retraction).** *If  $\Gamma, x : T' \vdash e : T$  and  $x \notin \text{fv}(e) \cup \text{fv}(T)$ , then  $\Gamma \vdash e : T$ .*

**Lemma 12 (Subformula Property).** *If  $\Gamma \vdash E[e] : T$ , then there exist a  $\Gamma'$  and a  $T'$  such that  $\Gamma \subseteq \Gamma'$  and  $\Gamma' \vdash e : T'$ .*

**Lemma 13 (Replacement).** *If  $\Gamma' \vdash e : T'$  is a sub-derivation of  $\Gamma \vdash E[e] : T$ , and  $\Gamma' \vdash e' : T'$ , then  $\Gamma \vdash E[e'] : T$ .*

**Definition 7.** *A redex is an expression which has the form of one of the left hands sides of our reduction rules (apart from the two context rules):  $\mathbf{v}.f$ ,  $\mathbf{v}.f := v$ ,  $\mathbf{new} \ \mathbf{v}.\mathbf{C}(\bar{v})$ ,  $\mathbf{v}.m(\bar{v})$ ,  $\mathbf{final} \ T \ x := v$ ;  $e$ , or  $v$ ;  $e$ .*

**Lemma 14.** *Any expression  $e$  factored uniquely as an evaluation context  $E[-]$  and another expression  $e'$ , such that  $e = E[e']$ . Either  $e'$  is a redex or an  $N$ .*

And now for subject reduction and progress, which together give soundness<sup>6</sup>.

**Theorem 1 (Subject Reduction).** *If  $\vdash H, e : T$  and  $H, e \rightsquigarrow H', e'$ , then  $H \subseteq H'$  and  $\vdash H', e' : T$ .*

*Proof.* By induction on the structure of  $e$ .

**Theorem 2 (Progress).** *If  $H \vdash e : T$ , then either  $e$  is a value, or there exist an  $H'$  and an  $e'$  such that  $H, e \rightsquigarrow H', e$ .*

*Proof.* By case analysis on the structure of  $e$ . The result depends on Lemma 6 to do all the work.

## 7 Advanced Tribe

This section addresses a number of issues concerning the design of Tribe.

<sup>6</sup> Proof sketches are available at <http://dsv.su.se/~tobias/appendix.pdf>.

---


$$\begin{array}{c}
\text{(EXPR-CALL-OVERRIDDEN)} \\
\frac{\Gamma \vdash p \uparrow A' \quad \Gamma \vdash p.qual \uparrow A \quad \vdash A' \sqsubseteq_i^* A}{T \ m(\overline{T \ x})\{ \dots \} \in \text{methods}(A) \quad \Gamma \vdash \overline{p} : T[p/\text{this}][\overline{p}/\overline{x}]} \\
\Gamma \vdash p :: qual.m(\overline{p}) : T[p/\text{this}][\overline{p}/\overline{x}]
\end{array}$$

$$\begin{array}{c}
\text{(EVAL-CALL-OVERRIDDEN)} \\
\frac{H(\iota) = [\dots]_{A'} \quad T \ m(\overline{T \ z})\{ e \} \in \text{methods}(A)}{H, \iota :: A.m(\overline{v}) \rightsquigarrow H, e[\iota/\text{this}, \overline{v}/\overline{z}]}
\end{array}$$


---

**Fig. 11.** Type and Reduction Rules for Overridden Method Call

## 7.1 Calling Overridden Methods

Name clashes and ambiguous super calls are problems faced by every language with multiple inheritance. In Tribe, we force the caller to explicitly state which class' implementation a method should be bound to, though we do so using a relative path to the desired class. This is more stable under changes to the inheritance hierarchy, so long as the targeted class is inherited where it is expected. Tribe provides uniform access for both methods appearing in super- and further bound classes. The syntax is:

$$e ::= \dots \quad | \quad p :: qual.m(\overline{p}) \quad \quad qual ::= \text{owner}^*.C^* \quad | \quad \mathbf{A}$$

In an overridden method call, *qual* is a qualifier which is used to refer to some inherited class, relative to the present context. It uses **owner** to enter surrounding classes and class names to go select specific classes within the surrounding class. The relative path to a class is resolved to a static class name at compiler time.

The type rule for overridden method call is found in Figure 11. The rule checks that the path of the target, merged with the qualifier, corresponds to some class the present class inherits from that contains the named method. The rest is the same as for ordinary method call.

An overridden method call binds to the implementation in the class denoted statically by the qualifier. Thus our approach is the same as that taken in C++, Java and Smalltalk, except that our paths are relative rather than absolute. At compile time, all method calls that use qualifiers are rewritten as follows: Replace the call  $p :: qual.m(\overline{p})$  by  $p :: A.m(\overline{p})$ , where  $\Gamma \vdash p.qual \uparrow A$ .

The reduction rule, which uses the absolute class name, appears in Figure 11.

The following example illustrates how the qualifier can be used to resolve ambiguous method calls (comments from  $*.Y.A$ 's point-of-view).

```

class X { // owner.owner.X
  class A { void m() { ... } }
}
class Y extends X { // owner
  class C { void m() { ... } }
  class D { void m() { ... } }

```

```

class A extends C, D { // inherits classes X.A, Y.C, Y.D
  this::owner.C.m();      // uses Y.C's impl. of m()
  this::owner.D.m();      // uses Y.D's impl. of m()
  this::owner.owner.X.A.m(); // uses A.B's impl. of m()
}
}

```

As in C++, qualifiers can be used for paths not starting with `this`.

Finally, recall that class construction (§3) non-deterministically resolved method calls when no overriding was present. A simple super-call can break this non-determinism. For example, assume that a class `C` inherits `A` and `B`, which both contain method `m`. The following code in `C` chooses the `m` method from `A`.

```

void m() { this::owner.A.m(); }

```

## 7.2 Adoption and Over-the-top Types

In Tribe, every class contains an nested class `Object` as the root of the nested subclass hierarchy. Furthermore, subclassing occurs only between classes nested within the same class. While offering simplicity, this approach suffers from practical limitations. In this section, we describe the notions of *adoption* and of *over-the-top types* which overcome these limitations.

Consider the following top-level classes—*not* nested within another class:

```

class Observer {
  this.owner.Subject subject; // relative type
}
class Subject { }

```

The `subject` field of `Observer` refers to a `Subject` type using a relative type rather than the absolute type `o.Subject`. This is okay as the absolute class of `this.owner` is `o`, which contains class `Subject`. Such relative types appearing in a top-level class—*over-the-top types*—are useful in combination with adoption.

*Adoption* occurs when a top-level class hierarchy is grafted into some other class, enabling better reuse of code which may have been written without family polymorphism in mind. The following syntax incorporates adoption into Tribe:

$$\begin{aligned}
class &::= class C \text{ extends } \bar{C} \{ cnstr fld^* class^* mthd^* adpt^* \} \\
adpt &::= adopt \ o.C
\end{aligned}$$

Note that adopting classes which are not at the top-level has no clear semantics.

The following code uses adoption to graft classes `o.Observer` and `o.Subject` into `o.Spy` to produce the classes `o.Spy.Observer` and `o.Spy.Subject`.

```

class Spy {
  adopt *.Observer;
  adopt *.Subject;
  class Observer { ... } // further binds adopted Observer
  class MaleSubject extends Subject { } // extension of adopted class
}

```

Class `o.Spy.Observer` inherits the field `this.owner.Subject subject`, whose type now corresponds to `o.Spy.Subject`. Without over-the-top types, this field would still have type `o.Subject`, and losing the advantages of virtual classes.

The following minor additions and modifications are required to incorporate adoption into the type system. Firstly, the following additional rules are added to subclassing and further binding resolution:

$$\frac{\text{(ADOPT-CLASS)}}{\frac{\text{class } C \cdots \{ \dots \text{adopt } o.C \} \in P(A)}{\vdash A.C.C_i \text{ cls}}} \quad \frac{\text{(ADOPT-SUBCLASS)}}{\frac{\text{class } C \cdots \{ \dots \text{adopt } o.C \} \in P(A)}{\vdash A.C.C_i \sqsubseteq_s o.C_i}}$$

The only change that need be made to (CLASS-TABLE) is to add  $\vdash o.C \text{ cls}$  for each adopted class, and to check that the class `C` is not already inherited. The final change required is ensuring that every top-level class referred to in an adopted class through an over-the-top type is also adopted. If, for example, `o.Subject` above is not adopted, the type of field `subject` in `o.Spy.Observer` would refer to a non-existent class. Worse, a totally different class with the same name could be introduced into `o.Spy`, which would be unsound.

We have not proven the soundness of this extension, but believe it to be so.

Adoption sometimes requires the adoption of many classes, potentially *every* top-level class. We plan to investigate alternative designs as future work.

### 7.3 Dynamic Type Casts

Extending Tribe to support dynamic casts is straightforward, as every object carries around sufficient information, an absolute class and an owner. The type rule is obvious. The reduction rules for type cast are as follows:

$$\frac{\text{(EXPR-CAST)}}{\frac{\Gamma \vdash e : T' \quad \Gamma \vdash T}{\Gamma \vdash (T)e : T}} \quad \frac{\text{(EVAL-TYPECAST-SUCCESS)}}{\frac{H \vdash v : T}{H, (T)v \rightsquigarrow H, v}} \quad \frac{\text{(EVAL-TYPECAST-FAIL)}}{\frac{\neg(H \vdash v : T)}{H, (T)v \rightsquigarrow H, \mathbf{error}}}$$

### 7.4 Decidability of Type Checking

As our type system is not syntax directed, it is not clear whether type checking is decidable. At the time of writing, we have determined that  $\Gamma \vdash T \uparrow A$  is decidable, and that decidability depends on the decidability of subtyping. Note that *vc*'s type system is decidable, and we believe that a fragment of our system corresponding to *vc* can easily be carved out (and would be decidable).

## 8 Related Work

There has been a range of research on family polymorphism since Ernst's original proposal [6], built on `gBeta` [5]. Systems that support family polymorphism can be divided into two categories: those that support *class families* and those that support *object families*. In class families, classes are nested in other classes,

whereas in object families, classes are nested in objects and types depend on paths. In terms of our initial graph example, class families can prevent non-coloured nodes from appearing in a coloured graph, but cannot prevent attempts at connecting nodes belonging to different graph instances. Object families can do both.

In this sense Concord [12] and CAESAR [15] groups are class families as well as the nested class system of .FJ [11]. The class families of Concord and .FJ are also “shallow” as they do not allow more than one level of nesting. Nystrom et al.’s Jx [16] presents a notion of nested inheritance that works much like class families with support for nestings of arbitrary depth. To enable family-polymorphic methods, Jx introduces the notion of *prefix types* of the form  $C[C']$  that denote the innermost enclosing class of class  $C'$  that is a subtype of  $C$ . To solve a similar problem, .FJ uses a bounded type parameter that is passed in separately. We believe that neither of these solutions is as intuitive as using relative paths based on `owner` or `out`.

Scala [8] and *vObj* [17] have both object family and class family polymorphism, although *vObj* does not support virtual classes, only virtual types. They both lack an `out` or `outer` construct but can use the syntax `C.this` to refer to the innermost object enclosing this with type  $C$ . Note that Tribe allows the description of such a type using a path starting from any field or variable, not just `this`; and allows classes at any point in a part.

Our proposal is closest both in spirit and expressiveness to *vc* [7]. The main difference is that in *vc*, all paths are relative to the current `this`, whereas Tribe supports absolute paths; paths beginning with variables; and paths intermixing classes and objects in any order. One practical benefit is the way Tribe handles family parameters. Writing our example from page 3 in *vc* would not allow the use of `e.owner` to obtain the class enclosing the argument edge, requiring the “family” class to be passed in separately. Tribe also has richer types and subtypes. For example, *vc* only defines subtyping for class names in the same path, whereas Tribe defines subtyping for variables and absolute class names.

Table 1 summarises this brief comparison of related systems. Ernst et al.’s *vc* paper includes a more comprehensive account [7].

**Table 1.** Comparison. <sup>1)</sup> through `myGrp`, only one level and for class families only; <sup>2)</sup> through prefix types, for class families only; <sup>3)</sup> only from `this`; <sup>4)</sup> only from `this` through `C.this` that denotes the innermost class enclosing `this` with type  $C$ .

	Concord	Caesar	.FJ	Jx	<i>vc</i>	<i>vObj</i>	Scala	Tribe
Object families	no	no	no	no	yes	yes	yes	yes
Class family types ( $C.C'$ )	yes	yes	yes	yes	no	yes	yes	yes
Relative paths	yes <sup>1</sup>	no	no	yes <sup>2</sup>	yes <sup>3</sup>	yes <sup>4</sup>	yes <sup>4</sup>	yes
Over-the-top types	no	no	no	no	no	no	no	yes
Adoption	no	no	no	no	no	no	no	yes

## 9 Conclusions and Future Work

We have presented Tribe, a type system for generalised class and family polymorphism. Tribe is simpler and more powerful than existing systems, with more flexible path-based types and extended subtype relations, resulting in a more expressive of the calculus with little additional conceptual overhead.

We are currently implementing a prototype programming system based upon Tribe, and more generally, we hope Tribe will serve as a suitable basis for research in a range of areas:

**Ownership Types** The savvy reader may have cottoned onto the fact that we have used the brand-name `owner` to refer to the surrounding instance.

Our original motivation for doing this work was related to ownership types, but we found the formalism for `vc` to be too unwieldy to modify. Hence, we devised our own, incorporating a number of extensions that we found useful. The work on ownership types in this setting will be reported in future work.

**Dynamic Nesting Structure** The `owner` field must be final for soundness.

But this prevents object aggregations (or nesting structures) from changing or evolving. It would be interesting to determine what would be required to change this. This can probably be achieved by linearising the objects whose type/ownership/structure will change, leveraging on DeLine and Fähndrich’s work on type states [4], and Clarke and Wrigstad’s External Uniqueness [3, 22].

**Generics** Adding generics and/or virtual types to our language would be, we expect, relatively simple to do (following perhaps Scala [8]). Type checking Scala is, unfortunately, undecidable. Venturing into this territory thus will require careful steps in order to remain decidable (assuming we can establish decidability for our system).

## References

1. David Aspinall and Martin Hofmann. Dependent types. Chapter in [18].
2. Gilad Bracha. *The Programming Language Jigsaw: Mixins, Modularity and Multiple Inheritance*. PhD thesis, University of Utah, 1992.
3. David Clarke and Tobias Wrigstad. External uniqueness is unique enough. In *Proceedings of the 17th European Conference on Object-Oriented Programming (ECOOP)*, Darmstadt, Germany, 2003.
4. Robert DeLine and Manuel Fähndrich. The fugue protocol checker: Is your software baroque? Technical Report MSR-TR-2004-07, Microsoft Research, 2003.
5. Erik Ernst. *gBeta—A Language with Virtual Attributes, Block Structure, and Propagating, Dynamic Inheritance*. PhD thesis, University of Aarhus, Denmark, 1999.
6. Erik Ernst. Family polymorphism. In *Proceedings of the 15th European Conference on Object-Oriented Programming*, London, UK, 2001. Springer-Verlag.
7. Erik Ernst, Klaus Ostermann, and William R. Cook. A virtual class calculus. In *Proceedings of Principles of Programming Languages (POPL)*, Charleston, South Carolina, USA, January 2006.

8. Martin Odersky et al. An overview of the scala programming language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland, 2004.
9. Manuel Fähndrich and K. Rustan M. Leino. Declaring and checking non-null types in an object-oriented language. In *OOPSLA '03: Proceedings of the 18th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, 2003.
10. James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *Java Language Specification*. Addison-Wesley Professional, 3rd edition, 2005.
11. Atsushi Igarashi, Chieri Saito, and Mirko Viroli. Lightweight family polymorphism. In *Proceedings of the 3rd Asian Symposium on Programming Languages and Systems (APLAS'05)*, volume 3780 of *LNCS*, Tsukuba, Japan, 2005.
12. Paul Jolly, Sophia Drossopoulou, Christopher Anderson, and Klaus Ostermann. Simple dependent types: Concord. In *6th ECOOP Workshop on Formal Techniques for Java-like Languages*, June 2004.
13. Ole Lehrmann Madsen, Birger Møller-Pedersen, and Kristen Nygaard. *Object-Oriented Programming in the BETA Programming Language*. Addison-Wesley, Reading, MA, USA, 1993.
14. Jan-Willem Maessen and Xiaowei Shen. Improving the java memory model using crf. In *OOPSLA '00: Proceedings of the 15th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, 2000.
15. Mira Mezini and Klaus Ostermann. Conquering aspects with caesar. In Mehmet Aksit, editor, *Proceedings of the 2nd International Conference on Aspect-Oriented Software Development (AOSD)*, pages 90–100, Boston, USA, March 2003.
16. Nathaniel Nystrom, Sephen Chong, and Andrew C. Myers. Scalable extensibility via nested inheritance. In *Proceedings of Objects, Programming Languages, Systems and Applications (OOPSLA)*, Vancouver, Canada, October 2004.
17. Martin Odersky, Vincent Cremet, Christine Röckl, and Matthias Zenger. A nominal theory of objects with dependent types. In *Proc. ECOOP'03*, Springer LNCS, July 2003.
18. Benjamin Pierce, editor. *Advanced Topics in Types and Programming Languages*. MIT Press, 2004.
19. Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Andrew P. Black. Traits: Composable units of behaviour. In Luca Cardelli, editor, *ECOOP 2003 – Object-Oriented Programming: 17th European Conference*, volume 2473 of *Lecture Notes In Computer Science*, pages 248–274. Springer-Verlag, July 2003.
20. Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley, 3rd edition, 1997.
21. Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.
22. Tobias Wrigstad. *Ownership-Based Alias Management*. PhD thesis, Stockholm University, 2006. Forthcoming.