A Virtual Class Calculus

Erik Ernst

University of Aarhus, Denmark

eernst@daimi.au.dk

Klaus Ostermann

Darmstadt Univ. of Technology, Germany

ostermann@informatik.tu-darmstadt.de

William R. Cook

University of Texas at Austin, USA

cook@cs.utexas.edu

Abstract

Virtual classes are class-valued attributes of objects. Like virtual methods, virtual classes are defined in an object's class and may be redefined within subclasses. They resemble inner classes, which are also defined within a class, but virtual classes are accessed through object *instances*, not as static components of a class. When used as types, virtual classes depend upon object identity – each object instance introduces a new family of virtual class types. Virtual classes support large-scale program composition techniques, including higher-order hierarchies and family polymorphism. The original definition of virtual classes in BETA left open the question of static type safety, since some type errors were not caught until runtime. Later the languages Caesar and gbeta have used a more strict static analysis in order to ensure static type safety. However, the existence of a sound, statically typed model for virtual classes has been a long-standing open question. This paper presents a virtual class calculus, *vc*, that captures the essence of virtual classes in these full-fledged programming languages. The key contributions of the paper are a formalization of the dynamic and static semantics of *vc* and a proof of the soundness of *vc*.

Categories and Subject Descriptors D.3.3 [*Language Constructs and Features*]: Classes and objects, inheritance, polymorphism; F.3.3 [*Studies of Program Constructs*]: Object-oriented constructs, type structure; F.3.2 [*Semantics of Programming Languages*]: Operational semantics

General Terms Languages, theory

Keywords Virtual classes, soundness

1. Introduction

Virtual classes are class-valued attributes of objects. They are analogous to virtual *methods* in traditional object-oriented languages: they follow similar rules of definition, overriding and reference. In particular, virtual classes are defined within an object's class. They can be overridden and extended in subclasses, and they are accessed relative to an object instance, using late binding. This last characteristic is the key to virtual classes: it introduces a dependence between static types and dynamic instances, because dynamic instances contain classes that act as types. As a result, the actual, dynamic value of a virtual class is not known at compile time, but it is known to be a particular class which is accessible as a specific

attribute of a given object, and some of its features may be statically known, whereas others are not.

When an object is passed as an argument to a method, the virtual classes within this argument are also accessible to the method. Hence, the method can declare variables and create instances using the virtual classes of its arguments. This enables the definition and use of higher-order hierarchies [9, 28], or hierarchies of classes that can manipulated, extended and passed as a unit. The formal parameter used to access such a hierarchy must be immutable; in general a virtual class only specifies a well-defined type when accessed via an immutable expression, which rules out dynamic references and anonymous values.

Virtual classes from different instances are not compatible. This distinction enables family polymorphism [8], in which families of types are defined that interact together but are distinguished from the classes of other instances. Virtual classes support arbitrary nesting and a form of mixin-based inheritance [3]. The root of a (possibly deeply) nested hierarchy can be extended with a set of nested classes which automatically extend the corresponding classes in the original root at all levels.

Virtual classes were introduced in the late seventies in the programming language BETA, but documented only several years later [21]. Methods and classes are unified as *patterns* in BETA. Virtual patterns were introduced to allow redefinition of methods. Since patterns also represent classes, it was natural to allow redefinition of classes, i.e. virtual classes. Later languages, including Caesar [22, 23] and gbeta [7, 8, 9] have extended the concept of virtual classes while remaining essentially consistent with the informally specified model in BETA [20]. For example, they have lifted restrictions in BETA that prevented virtual patterns (classes) from inheriting other virtual patterns (classes). So in this sense the design of virtual classes has only recently been fully developed.

Unfortunately, the BETA language definition and implementation allows some unsafe programs and inserts runtime checks to ensure type safety. Caesar and gbeta have stronger type systems and more well-defined semantics. However, their type systems have never been proven sound. This raises the important question of whether there exists a sound, type-safe model of virtual classes.

This paper provides an answer to this question by presenting a formal semantics and type system for virtual classes and demonstrating the soundness of the system. This calculus is at the core of the semantics of Caesar and gbeta and would presumably be at the core of every language supporting family polymorphism [8] and incremental specification of class hierarchies [9].

The calculus does not allow inheritance from classes located in other objects than **this**, and we use some global conditions to prevent name clashes. The significance of these restrictions and the techniques used to overcome them in the full-fledged languages are described in Section 5 and 8. The approach to static analysis taken in this paper was pioneered in BETA, made strict and complete in gbeta, and adapted and clarified as an extension to Java in Caesar.

The claim that virtual classes are inherently not type-safe should now be laid to rest.

The primary contributions of this paper are:

- Development of *vc*—a statically typed virtual class calculus, specified by a big-step semantics with assignment. The formal semantics supports the addition of virtual classes to mainstream object-oriented languages.

- Proof of the soundness of the type system. This paper includes the theorems, and the proofs are available in an accompanying technical report [10]. We use a proof technique that was developed for big-step semantics of object-oriented languages [6]. The preservation theorem ensures that an expression reduces to a value of the correct type, or a null pointer error, but never a dynamic type error. No results are proven about computations that do not terminate.

- We strengthen the traditional approach to soundness in big-step semantics by proving a *coverage* lemma, which ensures that the rules cover all cases, including error situations. This lemma plays a role analogous to the progress lemma for a small-step semantics [29]: it ensures that evaluation does not get stuck as a result of a missing case in the dynamic semantics.

2. Overview of Virtual Classes

Virtual classes are illustrated by a set of examples using an informal syntax in the style of Featherweight Java [17] or ClassicJava [12]. The distinguishing characteristics of *vc* include the following:

- Class definitions can be nested to define virtual classes.

- An instance of a nested class can refer to its *enclosing object* by the keyword **out**.

- Objects contain mutable *variables* and immutable *fields*. Fields are distinguished from variables by the keyword **field**. Fields must all be initialized by constructor arguments.

- A type is described by a *path* to an object and the name of a class in that object.

- The types of arguments and the return type of a method can use virtual classes from other arguments.

These concepts are illustrated in the examples given below. A formal syntax for *vc* is defined in Section 3. The main difference between the informal and formal syntax is that the formal syntax unifies classes and methods into a single construct, thus highlighting the syntactic and semantic unification of these concepts.

```
class Base { // contains two virtual classes
  class Exp {}
  class Lit extends Exp {
    int value; // a mutable variable
  }
  Lit zero;    // a mutable variable
  out.Exp TestLit() {
        out.Lit l;
        l = new out.Lit();
        l.value = 3;
        l;
  }
}
```

Figure 1. Defining virtual classes for expressions.

```
class WithNeg extends Base {
  class Neg extends Exp {
    Neg(out.Exp e) { this.e = e; }
    field out.Exp e;
  }
  out.Exp TestNeg() {
        new out.Neg(TestLit());
  }
}
```

Figure 2. Adding a class for negation expressions.

```
class WithEval extends Base {
  class Exp {
    int eval() { 0; }
  }
  class Lit {
    int eval() { value; }
  }
  int TestEval() {
        out.TestLit().eval();
  }
}
```

Figure 3. Adding an evaluation method on expressions.

```
class NegAndEval extends WithNeg, WithEval {
  class Neg {
    Neg(out.Exp e) { this.e = e; }
    int eval() { −e.eval(); }
  }
  int TestNegAndEval() {
        out.TestNeg().eval();
  }
}
```

Figure 4. Combining the negation class and evaluation method.

2.1 Higher-Order Hierarchies

Virtual classes provide an elegant solution to the *extensibility problem* [5, 19]: how to easily extend a data abstraction with both new representations and new operations. This problem is also known as the *expression problem* because a canonical example is the representation of the abstract syntax of expressions [36, 34, 38]. We present a solution to a simplified version of a standardized problem definition [15].

In Figure 1, the class Base contains two virtual classes: a general class Exp representing numeric expressions and subclass Lit representing numeric literals. All classes in *vc* are virtual classes and can be arbitrarily nested. Top-level classes are virtual by means of an implicit root class containing all top-level declarations. The method TestLit is explained below.

A *family* is a collection of virtual classes that depend upon each other. For example, the classes Exp and Lit are a family that exists within class Base. A family can be extended by subclassing the class in which it is defined. For example, Figure 2 extends the family to include a class Neg representing negation expressions.

Every virtual class has an *enclosing object*, to which the class can refer explicitly via the keyword **out**. In Figure 2, class Neg contains a *field* of type **out**.Exp. The type **out**.Exp is a reference to the class Exp in the enclosing instance of Neg. In general the type

out.A in class B denotes the sibling A of B. Because of subclassing and late binding, the dynamic value of **out** in Neg may be an instance of WithNeg or a subclass thereof. The **out** keyword can be repeated to access enclosing objects of the enclosing object.

The test functions in Figures 1 and 2 create a test instance of each class. The objects are created by accessing a virtual class (Lit or Neg) in the enclosing object. The return type of the methods is out.Exp rather than Exp because activation records are treated as separate objects whose enclosing object is the object containing the method, hence a property of the object containing the method must be accessed via **out**, whereas method parameters are accessed via **this**. A test can be run by invoking new WithNeg().TestNeg().

Redefinition of a virtual class occurs when it is declared and it is already defined in a superclass. In Figure 3, Exp and Lit are redefined to include an eval method; it is a redefinition because the family WithEval extends Base and they both define Exp and Lit. All superclasses in *vc* are *virtual superclasses* because redefinition of a class that is used as superclass affects its subclasses as well, so that the entire family is redefined.

The *static path* of a class definition is the lexical address of a class definition defined by the list of names of lexically enclosing class definitions. The static paths of the class definitions in Figure 3 are WithEval, WithEval.Exp and WithEval.Lit. Static paths never appear in programs, because virtual classes are always accessed through an object instance, not a class. However, they are useful for referring to specific class definitions.

Note that references to classes are "late bound" just like methods: when Base.TestLit is called from WithEval.TestEval the references to Lit are interpreted as WithEval.Lit, not Base.Lit.

A virtual class can have multiple superclasses, as in the definition of NegAndEval in Figure 4, which composes WithNeg and WithEval and adds the missing implementation of evaluation for negation expressions.

Hierarchies are not only first-class values, they can also be composed as a consequence of composing the enclosing class. The semantics of this composition is that nested virtual classes are composed, continuing recursively into nested classes. This phenomenon was introduced as *propagating combination* in [7] and later referred to as *deep mixin composition* [38]. This is achieved by combining the superclasses of the virtual class using *linearization*. For example, the class NegAndEval.Neg implicitly extends class WithNeg.Neg. Its also extends both Base.Exp and WithEval.Exp.

This behavior is a form of mixin-based inheritance [3] in that new class bodies are inserted into an existing inheritance hierarchy. For example, although WithNeg.Neg in Figure 2 has Exp as a declared superclass, after linearization it has WithEval.Exp as its immediate superclass.

2.2 Path-based Types

The example in Figure 5 illustrates path-based types and family polymorphism. The argument types in the previous examples have had the form C or **out**.C, where **out** can be repeated multiple times. Types can also be named via fields, which are immutable object instances that may contain virtual classes. The variable n defined at the bottom of Figure 5 has type f1.Exp, meaning that only instances of Exp whose enclosing object is identical to the value of f1 may be assigned to n. In general, a type consists of a path that specifies how to access an object, together with a class name. To ensure that this is well-defined, the path must only contain **out** and/or immutable fields, but not mutable variables. Hence, type compatibility depends on object identity, but types do not depend on values in any other way. More specifically, the type system makes sure that two types are only compatible if they are known to have identical enclosing objects.

```
class Test {
  int Test(out.WithNeg f1, out.NegAndEval f2) {
    this.f1 = f1; this.f2 = f2;
    n = buildNeg(f1, n);  // OK
    // n.eval();  -- Static error
    f2.zero = new f2.Lit();  // OK
    // n2 = buildNeg(f2, f1.zero)  -- Static error
    n2 = buildNeg(f2, f2.zero);  // OK
    n2.eval();  // OK
  }
  ne.Neg buildNeg(out.out.WithNeg ne, ne.Exp ex) {
    new ne.Neg(ex);
  }
  field out.WithNeg f1
  field out.NegAndEval f2
  f1.Exp n
  f2.Exp n2
}
new Test(new NegAndEval(), new NegAndEval())
```

Figure 5. Example of family polymorphism

Although the resulting types may resemble Java package/class names, they are very different because objects play the role of packages, and the class that creates a package can be subclassed.

2.3 Family Polymorphism

A *family object* is an object that provides access to a class family. A family object may be the enclosing object for an expression, but it may also be a method argument or the value of a field. As a provider of classes, and hence types, it enables type parameterization of classes and methods. But virtual classes are different from parameterized types: while type parameters are bound statically at compile-time, virtual classes are bound dynamically at runtime. Thus virtual classes enable a new kind of subtype polymorphism known as family polymorphism [8].

Family objects can also be used to create new objects, even though the classes in the family object are not known at compile time. To achieve the same effect in a main-stream language like Java, a factory method [13] must be used. However, the typing relation between related classes is then lost, whereas a family object testifies to the interrelatedness of its nested family classes.

In Figure 5, f1 and f2 inside Test are used as family objects. The constructor call in the last line of the example shows how f1 is polymorphically initialized with a subtype of its static types. The field f1 of class Test is declared to be an out.WithNeg, but the constructor is called with an argument of type NegAndEval, which illustrates that entire class hierarchies are first class values, subject to subtype polymorphism via their family objects, and the nested family classes are usable for both typing and object creation.

The assignments and calls in the body of the Test constructor illustrate the expressiveness of the type system. For example, although the buildNeg method is not aware of the eval method introduced by WithEval, it is possible to assign the result to n2 and call eval on the returned value. This is an important special case of family polymorphism where the types of arguments or the return type of a method depend on other arguments. The example also shows a few cases that are rejected by the type checker because they would potentially lead to a type error at runtime.

3. Syntax

The formal syntax of *vc* has been designed to make the presentation of the semantics as simple as possible, hence the formal syntax

Grammar of *vc*

CL	::=	**class** C **extends** $\overline{\text{C}}$ { K $\overline{\text{CL}}$; $\overline{\text{T}}$ f; $\overline{\text{T}}$ v }
K	::=	T C($\overline{\text{T f}}$) { e; }
T	::=	path.C
path	::=	spine.$\overline{\text{f}}$
spine	::=	**this**.$\overline{\text{out}}$
e	::=	**null** | e ; e | path | path.v |
		path.v = e | **new** path.C($\overline{\text{e}}$)

Identifiers

class names	C	
field names	f	
variable names	v	
members	m	= f $\cup$ v
(C, f, and v are pairwise disjoint)		

Figure 6. Syntax of virtual class calculus *vc*

deviates from the informal syntax used in the examples in a few points that will be described in this section.

3.1 Notational Conventions

Our formal definitions use a number of syntactic conventions. A bar above a metavariable denotes a list: $\overline{\text{p}}$ stands for $p_1, ..., p_k$ for some natural number $k \geq 0$. If $k = 0$ then the list is empty. The length of $\overline{\text{p}}$ is $|\overline{\text{p}}|$. The same notation is used for lists whose elements are separated by dots or commas, e.g., $f_1.f_2.\cdots.f_k = \overline{\text{f}}$. A list may also be represented by a combination of barred and unbarred variables: $\overline{\text{f}}$.f stands for $f_1.\cdots.f_k.f$, where f denotes the last item of the list. Following common convention, $\overline{\text{T}}$ f represents a list of pairs $T_1 \ f_1 \cdots T_k \ f_k$ rather than a pair of lists. An empty list is written nil_x, where x identifies the kind of items that the list should contain. The subscript x may be omitted if it is clear from context. The notation [f] represents a list with a single element f. Finally, in function definitions with overlapping branches the first matching case is used.

3.2 Formal Syntax of *vc*

The formal syntax of *vc* is defined in Figure 6. A class definition CL consist of a name, the superclass names $\overline{\text{C}}$, a constructor K, a list of nested class definitions $\overline{\text{CL}}$, declarations $\overline{\text{T}}$ f of immutable fields, and declarations $\overline{\text{T}}$ v of mutable variables. A constructor K consists of a return type T, the class name, the formal parameters $\overline{\text{T}}$ f, and an expression e. The constructor has a return type because it can return other things than the new object, which enables the encoding of methods as classes.

The keyword **field** from the informal syntax is not needed, because field and variable names are separate in the formal syntax and use different metavariables—f for fields and v for variables. Field and variable names must be unique within the program in order to simplify the handling of name clashes in connection with class composition. Class names are unique in that two definitions of the same class name must have a common superclass. We will later discuss the implications and possible relaxations of these restrictions. Note, however, that any program in which the names are reused can always be rewritten to a program with unique names.

Expressions include standard forms for the current object or any of the enclosing objects via spine, access to fields of the current or an enclosing object via path, access and assignment of variables, path.v, and path .v = e, and the null value, **null**. Method calls and object construction are unified in the expression **new** path.C($\overline{\text{e}}$).

Types in the syntax of *vc* have the form path.C. A path has the form **this**.$\overline{\text{out}}$.$\overline{\text{f}}$. Thus a type allows a class C to be identified by navigating to any enclosing object and then traversing fields to find the object which contains C.

Metavariable

$$\text{static paths} \quad \text{p} \quad ::= \quad \overline{\text{C}}$$

Class table

$$CT(\text{p}) = CT2(\text{p}, \overline{\text{CL}}_{root})$$

$$\frac{\text{CL}_i = \textbf{class } C \textbf{ extends } \overline{\text{C}} \{ \ ... \ \}}{CT2(\text{C}, \overline{\text{CL}}) = \text{CL}_i}$$

$$\frac{\text{CL}_i = \textbf{class } C \textbf{ extends } \overline{\text{C}} \{ \ K \ \overline{\text{CL}}'; \ ... \ \}}{CT2(\text{C.p}, \overline{\text{CL}}) = CT(\text{p}, \overline{\text{CL}}')}$$

All members

$$Members(\text{nil}_\text{p}) = \text{nil}_{\text{T f}}, \text{nil}_{\text{T v}}$$

$$\frac{Members(\overline{\text{p}}) = \overline{\text{T}} \ \overline{\text{f}}, \overline{\text{T}}' \ \overline{\text{v}} \qquad CT(\text{p}) = \textbf{class } C \textbf{ extends } \overline{\text{C}} \{ \ K \ \overline{\text{CL}}; \ \overline{\text{T}}'' \ \overline{\text{f}}'; \overline{\text{T}}''' \ \overline{\text{v}}'\}}{Members(\text{p} \ \overline{\text{p}}) = \overline{\text{T}}'' \ \overline{\text{f}}' \ \overline{\text{T}} \ \overline{\text{f}}, \overline{\text{T}}''' \ \overline{\text{v}}' \ \overline{\text{T}}' \ \overline{\text{v}}}$$

Constructor

$$\frac{CT(\text{p}) = \textbf{class } C \textbf{ extends } \overline{\text{C}} \{ \ K \ \overline{\text{CL}}; \ \overline{\text{T}}'' \ \overline{\text{f}}'; \overline{\text{T}}''' \ \overline{\text{v}}'\}}{Constr(\text{p}) = K}$$

Figure 7. Auxiliary definitions

Primitive types like bool and int are omitted; they just add complexity to the formalism without adding value. A member m is either a field or a variable.

3.3 Translating Informal Notation to *vc*

The translation of the informal language to the formal syntax of *vc* is straightforward. The most significant difference is that *vc* unifies methods and classes into a single definition construct. This technique originated in Simula, where classes were simply functions that returned the current activation record. In *vc* activation records are first-class values that are accessed by **this**. Thus a class is simply a definition that returns **this**, while a method is a definition that returns any other value.

Hence, method definitions in the informal language correspond to class declarations in *vc*, where the constructor represents the method body. More formally, the translation is as follows:
T C($\overline{\text{T}}$ f) { $\overline{\text{T}}$ v; e; } $\Rightarrow$ **class** C **extends** { K nil_{CL}; $\overline{\text{T}}$ f; $\overline{\text{T}}$ v }
where K = T C($\overline{\text{T}}$ f) { e; }. Method calls are translated by prefixing them with the keyword **new**.

As in Java, constructors in the informal syntax do not specify a return type or return value, but these must be specified in *vc*. For a class definition C in the informal syntax, the constructor return type is always **out**.C and the returned value is always **this**.

In the informal syntax a class definition with no superclasses may omit the **extends** clause. In the formal syntax it must be present, but the list of superclasses can be empty. The assignments of the constructor arguments is omitted in the formal syntax; instead, the name of the constructor arguments are matched against the field names. Constructors are required in *vc*, while the informal syntax assumes a default constructor if none is given.

The informal notation omits **this** when followed by **out** or a field. *vc* has no implicit scoping rules, and all access to fields, variables, and classes must be disambiguated by a spine.

$\iota_{root} \mapsto \llbracket \ \bot \ \| \ \mathsf{C_{root}} \ \| \qquad \rrbracket$
$\iota_1 \mapsto \llbracket \ \iota_{root} \ \| \ \mathsf{NegAndEval} \ \| \ \mathsf{zero} : \mathbf{null} \ \rrbracket$
$\iota_2 \mapsto \llbracket \ \iota_{root} \ \| \ \mathsf{NegAndEval} \ \| \ \mathsf{zero} : \iota_5 \ \rrbracket$
$\iota_3 \mapsto \llbracket \ \iota_{root} \ \| \ \mathsf{Test} \ \| \ \mathsf{f1} : \iota_1 \ \mathsf{f2} : \iota_2 \ \mathsf{n} : \iota_4 \ \mathsf{n2} : \iota_6 \ \rrbracket$
$\iota_4 \mapsto \llbracket \ \iota_1 \ \| \ \mathsf{Neg} \ \| \ \mathsf{e} : \mathbf{null} \ \rrbracket$
$\iota_5 \mapsto \llbracket \ \iota_2 \ \| \ \mathsf{Lit} \ \| \ \mathsf{value} : 0 \ \rrbracket$
$\iota_6 \mapsto \llbracket \ \iota_2 \ \| \ \mathsf{Neg} \ \| \ \mathsf{e} : \iota_5 \ \rrbracket$

Figure 9. Dynamic Heap after executing the example in Figure 5

The informal language allows more general expressions where the calculus only allows paths: e.m, **new** e.C($\overline{e}$), and e.v $=$ e'. The general forms are translated into the calculus by rewriting e.m as **new this**.C'(e) where C' is a new local class with a field T f where T is the type of e, and whose constructor returns **this**.f.m. The translation is legal because the member is accessed through the new field. The other two constructs (**new** e.C($\overline{e}$), and e.v $=$ e') are handled similarly. The consequence of this is that the formal treatment need not take types inside temporary objects into account. This is a significant simplification, and handling types in temporaries does not produce useful extra insight.

3.4 Auxiliary Definitions

Figure 7 gives some auxiliary definitions. A *static path* p is a list of class names $\overline{\mathsf{C}}$. The function *CT* looks up a class definition. We assume the existence of a globally available program in the form of a list of top-level class declarations $\overline{\mathsf{CL}}_{root}$, which would otherwise embellish many relations and functions. *CT* is a partial function from static paths to class definitions. It uses the helper function *CT2*, which recursively enters each class definition named in the path starting from root. For example, the static path Base.Lit denotes the definition of Lit inside Base in Figure 1.

A static path that identifies a valid class is called a *mixin*. The set of mixins in a program is equivalent to the static paths p for which $CT(\mathsf{p}) \neq \bot$. Since there is a one-to-one correspondence between a mixin (a static path) and its class definition, we also use the term mixin to refer to the body of the corresponding class, i.e., the part of a class declaration between the curly brackets { ... }.

The function *Members* collects all field and variable declarations found in a list of mixins $\overline{\mathsf{p}}$. The function *Constr*(p) returns the constructor of $CT(\mathsf{p})$ given a static path p.

4. Operational Semantics

The operational semantics is defined in big-step style. The semantic domains, evaluation relation, and helper functions are given in Figure 8. Both the operational semantics and the type system have also been implemented in Haskell.

4.1 Objects and the Heap

As in most object-oriented languages, an object in *vc* combines state and behavior. An **Object** is a tuple containing a pointer to its enclosing object ι, a class name C, and a list of fields and variables with their values.

The fields and variables are the state of the object; fields are immutable while variables can be updated. The heap is standard: a map H from addresses ι to objects. The top-level root object has the special address ι_{root}. An example heap is given in Figure 9.

The features of the object are determined by the enclosing object ι and the class C. The enclosing object specifies the environment containing the class from which the object ι' was created: an object ι' with enclosing object ι and class C must have been created by evaluating an expression equivalent to **new** ι.C(...).

An object's features are defined by a list of mixins, or class bodies; these class bodies contain the declarations of members and nested classes. In *vc* there are no methods, but classes may be used as methods. The list of mixins of an object is computed from the class name and the mixins of the enclosing object.

Note that the definition of **Object** is optimized for a situation where all path expressions associated with an object should be understood relative to the same environment—the same enclosing object. It would be a relevant extension of *vc* to allow inheritance from classes inside other objects than **this** (i.e., to allow superclasses on the form path.C), but it would then be necessary to maintain an environment for each mixin or for each feature. It is possible to do this, and for instance the static analysis and run-time support for gbeta maintains a separate enclosing object for each mixin. This causes a non-trivial amount of extra complexity, even though the basic ideas are unchanged. It is part of future work to extend *vc* correspondingly.

4.2 Mixin Computation

The *Mix* function computes the behavior, or mixin list, of an object ι in the heap H. It does so by first computing the mixins of the enclosing object. All definitions of C and its superclasses are assembled into this mixin list. The mixin list of the root object has only a single element, namely the empty static path.

The *Assemble* function[1] computes the mixin list for a class C relative to an enclosing mixin list $\overline{\mathsf{p}}$. It calls *Defs* to collect all the definitions of C located in any of the class bodies specified by $\overline{\mathsf{p}}$. If the resulting list of mixins is empty then the class is not defined and *Assemble* returns $\bot$. Otherwise, the result is a list of static paths that identifies all definitions of C contained in the list of enclosing mixins.

As an example, let us consider the computation of $Mix(\mathsf{H}, \iota_4)$ in the program in Figure 1-4 and the sample heap in Figure 9. Assume that the mixin list $\overline{\mathsf{p}}$ of the enclosing object ι_1 has been computed to yield [Base, WithNeg, WithEval, NegAndEval]. Then $Defs(\overline{\mathsf{p}}, \mathsf{Neg}) = $ [WithNeg.Neg, NegAndEval.Neg].

The complete mixin list must also include the mixins of all the superclasses. To do so, *Assemble* maps *Expand* over the list of static paths that was computed with *Defs*, and linearizes the result. *Expand* assembles each of the superclasses of C, linearizes the result, and appends the class itself to the resulting list. In our example $[Expand(\overline{\mathsf{p}}, \mathsf{p}) \mid \mathsf{p} \leftarrow \mathsf{WithNeg.Neg} \ \mathsf{NegAndEval.Neg}] = \overline{\mathsf{p}}'\overline{\mathsf{p}}''$, where $\overline{\mathsf{p}}' = $ [Base.Exp, WithEval.Exp, WithNeg.Neg] and $\overline{\mathsf{p}}'' = $ [NegAndEval.Neg].

Linearization sorts an inheritance graph topologically, such that method calls are dispatched along the sort order. The function *Linearize* linearizes a list of mixin lists, i.e., it produces a single mixin list which contains the same mixins as those in the operands; the order of items in each of the input lists is preserved in the final result, to the degree possible. *Linearize* is defined in terms of a binary linearization function, *Lin2*. This function is an extension of the C3 linearization algorithm [1, 7] which has been used in gbeta and Caesar for several years. The linearization algorithm allows a programmer of a subclass to control the ordering of the class's mixins by choosing the order in which the superclasses appear in the **extends** clause.

Lin2 produces the same results as C3 linearization in every case where C3 linearization succeeds—this result follows trivially from the fact that the definition of C3 is just the four topmost cases in the definition of *Lin2*. The cases where C3 linearization fails are

---

[1] The [… | …] notation used in the definition of *Defs*, *Assemble*, and *Expand* means list comprehension as for example in Haskell. Note that we append an element to a list by just writing the element to append after the list. For example, $[2n \mid n \leftarrow 1...5, n > 3] 42$ is the list [8, 10, 42].

Objects and the Heap:

$$\text{Address} = \text{natural numbers} \qquad \iota$$
$$\text{Object} = \{\ [\![\ \iota\ \|\ C\ \|\ \overline{f:\text{val}}\ \ \overline{v:\text{val}}'\ ]\!]\ \} \qquad [\![\ ...\ ]\!]$$
$$\text{Heap} = \text{Address} \xrightarrow{\text{fin}} \text{Object} \qquad H$$
$$\text{Value} = \text{Address} \cup \{\textbf{null}\} \qquad \text{val}$$

Evaluation rules:

$$\leadsto:\ e \times \textbf{Heap} \times \textbf{Address} \to \textbf{Value} \cup \{\text{TypeErr, NullErr}\} \times \textbf{Heap}$$

$$\textbf{null}, H, \iota \leadsto \textbf{null}, H \quad \text{(R1)}$$

$$\frac{\mathcal{W}alk(H, \iota, \text{path}) = \text{val}}{\text{path}, H, \iota \leadsto \text{val}, H} \ \text{(R3)}$$

$$\frac{\begin{array}{c} e, H, \iota \leadsto \text{val}, H' \\ e', H', \iota \leadsto \text{val}', H'' \end{array}}{e\,\textbf{;}\,e', H, \iota \leadsto \text{val}', H''} \ \text{(R2)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \\ H(\iota')(v) = \text{val} \end{array}}{\text{path.v}, H, \iota \leadsto \text{val}, H} \ \text{(R4)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \qquad e, H, \iota \leadsto \text{val}, H' \\ H'(\iota')(v) \neq \bot \qquad H'' = H'[\iota' \mapsto H'(\iota')[v \mapsto \text{val}]] \end{array}}{\text{path.v} = e, H, \iota \leadsto \text{val}, H''} \ \text{(R5)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \qquad H = H_1 \\ e_i, H_i, \iota \leadsto \text{val}_i, H_{i+1} \ \text{for } i \in \{1...|\overline{e}|\} \\ H' = H_{|\overline{e}|+1} \qquad \overline{p} = \mathcal{A}ssemble(\mathcal{M}ix(H', \iota'), C) \\ \mathcal{M}embers(\overline{p}) = \overline{T}\,\overline{f}, \overline{T}'\,\overline{v} \qquad |\overline{f}| = |\overline{\text{val}}| \\ \iota'' \text{ is new in } H' \qquad \mathcal{C}onstr(p_{|\overline{p}|}) = T\,C(\_)\{e';\} \\ H'' = H'[\,\iota'' \mapsto\ [\![\ \iota'\ \|\ C\ \|\ \overline{f:\text{val}}\ \ \overline{v:\textbf{null}}\ ]\!]\ ] \\ e', H'', \iota'' \leadsto \text{val}, H''' \end{array}}{\textbf{new } \text{path.}C(\overline{e}), H, \iota \leadsto \text{val}, H'''} \ \text{(R6)}$$

Enclosing object:

$$\mathcal{E}ncl(\ [\![\ \iota\ \|\ \_\ \|\ ...\ ]\!]\ ) = \iota$$

Evaluation functions:

$$\mathcal{W}alk(H, \iota, \textbf{this}) = \iota$$
$$\mathcal{W}alk(H, \iota, \text{spine}.\textbf{out}) = \mathcal{E}ncl(H(\iota')) \quad \text{if } \mathcal{W}alk(H, \iota, \text{spine}) = \iota' \neq \iota_{\text{root}}$$
$$\mathcal{W}alk(H, \iota, \text{path.f}) = \text{val} \quad \text{if } H(\mathcal{W}alk(H, \iota, \text{path}))(f) = \text{val}$$

$$\mathcal{W}alk(H, \iota, \text{path.f}) = \text{NullErr} \quad \text{if } \mathcal{W}alk(H, \iota, \text{path}) = \textbf{null}$$
$$\mathcal{W}alk(H, \iota, \text{path.f}) = \text{TypeErr} \quad \text{if } H(\mathcal{W}alk(H, \iota, \text{path}))(f) = \bot$$
$$\mathcal{W}alk(H, \iota, \text{spine}.\textbf{out}) = \text{TypeErr} \quad \text{if } \mathcal{W}alk(H, \iota, \text{spine}) = \iota_{\text{root}}$$

Error handling:

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \textbf{null}, H \\ \hline \text{path.v}, H, \iota \leadsto \text{NullErr}, H \\ \text{path.v} = e, H, \iota \leadsto \text{NullErr}, H \\ \textbf{new } \text{path.}C(\overline{e}), H, \iota \leadsto \text{NullErr}, H \end{array}}{} \ \text{(ER1)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \qquad H(\iota')(v) = \bot \\ \hline \text{path.v}, H, \iota \leadsto \text{TypeErr}, H \\ \text{path.v} = e, H, \iota \leadsto \text{TypeErr}, H \end{array}}{} \ \text{(ER2)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \\ \mathcal{A}ssemble(\mathcal{M}ix(H, \iota'), C) = \bot \\ \hline \textbf{new } \text{path.}C(\overline{e}), H, \iota \leadsto \text{TypeErr}, H \end{array}}{} \ \text{(ER3)}$$

$$\frac{\begin{array}{c} \text{path}, H, \iota \leadsto \iota', H \\ \mathcal{A}ssemble(\mathcal{M}ix(H, \iota'), C) = \overline{p} \\ \mathcal{M}embers(\overline{p}) = \overline{T}\,\overline{f}, \_ \qquad |\overline{e}| \neq |\overline{f}| \\ \hline \textbf{new } \text{path.}C(\overline{e}), H, \iota \leadsto \text{TypeErr}, H \end{array}}{} \ \text{(ER4)}$$

Mixin Computation:

$$\mathcal{M}ix(H, \iota_{\text{root}}) = [\text{nil}_c]$$
$$\mathcal{M}ix(H, \iota) = \mathcal{A}ssemble(\mathcal{M}ix(H, \iota'), C)$$
$$\quad \text{where } H(\iota) = [\![\ \iota'\ \|\ C\ \|\ ...\ ]\!]$$

$$\mathcal{A}ssemble(\overline{p}, C) = \mathcal{L}inearize[\ \mathcal{E}xpand(\overline{p}, p)\ |\ p \leftarrow \mathcal{D}efs(\overline{p}, C)\ ]$$

$$\mathcal{D}efs(\overline{p}, C) = check[\ p.C\ |\ p \leftarrow \overline{p}, CT(p.C) \neq \bot\ ]$$
$$\quad \text{where} \quad check(\overline{p}) = \begin{cases} \bot & |\overline{p}| = 0 \\ \overline{p} & otherwise \end{cases}$$

$$\mathcal{E}xpand(\overline{p}, p) = \mathcal{L}inearize([\ \mathcal{A}ssemble(\overline{p}, C)\ |\ C \leftarrow \overline{C}\ ]\,p)$$
$$\quad \text{where } CT(p) = \textbf{class } C'\ \textbf{extends}\ \overline{C}\ \{\ ...\ \}$$

$$\mathcal{L}inearize(\text{nil}_{\overline{p}}) = \text{nil}_p$$
$$\mathcal{L}inearize(\overline{\overline{p}}\ \overline{p}) = \mathcal{L}in2(\mathcal{L}inearize(\overline{\overline{p}}), \overline{p})$$
$$\mathcal{L}in2(\text{nil}_p, \text{nil}_p) \quad = \quad \text{nil}_p$$
$$\mathcal{L}in2(\overline{p}\,p, \overline{p}'\,p) \quad = \quad \mathcal{L}in2(\overline{p}, \overline{p}')\,p$$
$$\mathcal{L}in2(\overline{p}, \overline{p}'\,p') \quad = \quad \mathcal{L}in2(\overline{p}, \overline{p}')\,p', \text{ if } p' \notin \overline{p}$$
$$\mathcal{L}in2(\overline{p}\,p, \overline{p}') \quad = \quad \mathcal{L}in2(\overline{p}, \overline{p}')\,p, \text{ if } p \notin \overline{p}'$$
$$\mathcal{L}in2(\overline{p}\,p\,\overline{p}''p, \overline{p}'p') \quad = \quad \mathcal{L}in2(\overline{p}\,\overline{p}''p, \overline{p}')\,p'$$

Figure 8. Operational semantics of *vc*

exactly the cases covered by the bottommost clause in the definition of $\mathcal{L}in2$, i.e., the cases where the two operands contradict each other with respect to the ordering of shared mixins (intuitively this means that they disagree about which mixin should be the more specific one); in these cases, $\mathcal{L}in2$ resolves the conflict by letting the rightmost operand decide the outcome.

The final result of computing $\mathcal{M}ix(H, \iota_4)$ is the mixin list Base.Exp WithEval.Exp WithNeg.Neg NegAndEval.Neg .

$\mathcal{L}in2$ is a total function on lists of mixins, and the set of mixins in the result is equal to the union of the sets of mixins in the operands. For soundness the set of mixins is relevant but the ordering makes no difference, so this generalization of C3 enhances the expressive power without affecting type safety.

4.3 Evaluation Rules and Error Handling

The evaluation relation $e, H, \iota \leadsto r, H'$ reduces an expression, a heap, and a current object to a value or an error and a new heap. The current object plays the role of the environment.

The expression **null** evaluates to the null value (R1). An expression sequence $e\,\textbf{;}\,e'$ evaluates to the result of evaluating e' in the heap that results from evaluating e (R2).

Evaluation of a path path does not affect the heap (R3). The value of the path is computed by the function $\Downarrow$, which "walks" a path from an address ι in the heap H to return the value specified by the path. As a base case, $\Downarrow$ returns ι when applied to the trivial path, **this**; spine.**out**^{n} locates the nth enclosing object of ι; finally a path path.f finds the object ι' for path and then returns the value of the field f in the object ι'.

Variable lookup path.v evaluates path to get ι', which is then looked up in the heap to get the variable's value (R4). An assignment $\text{path.v} = e$ evaluates path and e to ι' and val (R5). It then checks that the variable is defined on the object and updates the heap to set variable v of ι' to val. The notation $H(\iota)(m)$ means lookup of the value of a field or variable m in the object ι. The notation $[v \mapsto \text{val}]$ appended to an object denotes (functional) update of the variable v of that object, and $H[\iota \mapsto ...]$ denotes heap update.

In (R6) a new object **new** path.C($\overline{\text{e}}$) is constructed by instantiating the virtual class C defined in the enclosing object ι' identified by path. The behavior $\overline{\text{p}}$ of the new object is assembled from the mixins of the enclosing object as described in Section 4.2. If the enclosing object does not contain a definition of C, then *Assemble* returns $\bot$ and rule (R6) does not apply. The mixin list $\overline{\text{p}}$ also specifies the members and the most specific constructor of the new object. To construct the object, the heap is extended to define a new address ι'' bound to a new object with enclosing object ι', class C, fields initialized to the evaluated constructor arguments, and variables initialized to null. The constructor body is then evaluated in the context of this new object. The result of the constructor is the result of the entire expression. If the constructor body is **this** (i.e., the class is used as a class in the conventional sense), then the result of the constructor call is ι''.

Two different kinds of error can occur during evaluation: Type errors (TypeErr) and null pointer errors (NullErr). The rule (ER1) handles access to a property of an object, where the object is **null**. (ER2) to (ER4) define the situations in which a type error occurs, namely if a member to be read or written is not available (ER2), or when creating an instance of a class C, but the enclosing object has no definition of C, i.e., its mixin list is empty (ER3), or the number of parameters does not match (ER4).

The rules for propagating errors are standard and straightforward, so they are omitted; the sequel assumes that NullErr or TypeErr errors are propagated. The complete list of error rules is given in the technical report accompanying this paper [10].

5. Type System

The *vc* type system uses nominal typing based on *paths* to objects containing virtual classes. Typing domains, type checking rules, and functions for abstract interpretation are given in Figure 10.

5.1 Types

The type of an expression describes an object ι obtained by evaluation of it in one of two ways. In the first case a path which leads to the object ι itself is computed statically, and in the second case a path to the *enclosing* object of ι is computed, as well as a class name characterizing the class of ι itself. The former is an *object type*, u, and the latter is a *class type*, s. An object type contains more information than a class type, because every object type can be converted into a class type, but not vice versa. Since a path only makes sense as seen from a lexical point in the program, typing judgements have the form $\text{p}' \vdash \text{e} : \text{t}$, where t is a type and p' represents the current **this** object.

An object type u has the form $\langle \text{p} \rangle . \overline{\text{f}}$. If an expression has the object type $\langle \text{p} \rangle . \overline{\text{f}}$ as seen from p', then p is a prefix of p', and the object denoted by the expression can be reached by going **out** ($|\text{p}'| - |\text{p}|$) steps and then following $\overline{\text{f}}$ in the heap. More formally, if the program and heap H are well-formed, the expression e is typable by $\text{p}' \vdash \text{e} : \langle \text{p} \rangle . \overline{\text{f}}$ in this program, the object ι_0 is appropriate as **this** for p', and e evaluates by e, H, $\iota_0 \rightsquigarrow \iota$, H', then $\mathcal{W}alk(\text{H}', \iota_0, \textbf{this}.\textbf{out}^j.\overline{\text{f}}) = \iota$, where $j = \mathcal{D}epth(\text{H}', \iota_0) - |\text{p}|$.

A class type s is on the form $\langle \text{p} \rangle . \overline{\text{f}}.\text{C}$. If an expression e has type $\langle \text{p} \rangle . \overline{\text{f}}.\text{C}$ and e, H, $\iota_0 \rightsquigarrow \iota$, H' as above then $\langle \text{p} \rangle . \overline{\text{f}}$ is an object type describing the enclosing object $\mathcal{E}ncl(\text{H}'(\iota))$, and ι is an instance of the class C which is nested in $\mathcal{E}ncl(\text{H}'(\iota))$, or a subclass thereof.

The type checker computes object types for paths or path-like expressions (like a sequence containing a path as last element). For an expression like path.v or **new** path.C, an object type cannot be computed because, in general, there is no path to that object. However, there is always a path to its enclosing object in these cases, hence such expressions can be assigned a class type.

5.2 Abstract interpretation of the heap

The operational semantics defines functions to navigate a heap and compute mixin lists of objects. In particular, $\mathcal{E}ncl$ navigates to an enclosing object, $\mathcal{W}alk$H follows a path starting from some object, and $\mathcal{M}ix$ computes the mixin list of an object. An abstract interpretation of these functions is at the core of the type system: $\mathcal{E}$, $\mathcal{W}$, and $\mathcal{M}$ are the static versions of $\mathcal{E}ncl$, $\mathcal{W}alk$, and $\mathcal{M}ix$, respectively. They serve the same purpose as their dynamic counterparts, but they receive and produce types instead of objects. Before going into the details of their definition, we will at first state some properties of $\mathcal{E}$, $\mathcal{W}$, and $\mathcal{M}$ and discuss the connection with $\mathcal{E}ncl$, $\mathcal{W}alk$H, and $\mathcal{M}ix$ (the formal statements and proofs of these properties are in [10]).

The most important connections between the static and dynamic semantics are (a) if a navigation along a path is ok in the abstract interpretation of the heap then the corresponding navigation is also ok in the dynamic heap, and (b) navigation preserves agreement. Agreement, which is formally defined later in this section, states that an object ι has type t as seen from an object ι_0 in a heap H, written H, $\iota_0 \vdash \iota \triangleright \text{t}$. Given a well-formed program and a well-formed heap and H, $\iota_0 \vdash \iota \triangleright \text{t}$, then the following holds:

1. Enclosing types agree with enclosing objects: if t is not the type of the root object, then $\mathcal{E}ncl(\text{H}(\iota))$ exists and H, $\iota_0 \vdash \mathcal{E}ncl(\text{H}(\iota)) \triangleright \mathcal{E}(\text{t})$.

2. The statically known set of mixins is a subset of the dynamic set of mixins, $\mathcal{M}ix(\text{H}, \iota) \supseteq \mathcal{M}(\text{t})$.

3. If a field or variable exists according to the abstract interpretation then it exists in the heap: $\mathcal{E}xists(\text{t}, \text{m}) \Rightarrow \text{H}(\iota)(\text{m}) \neq \bot$.

4. If t is an object type u and a path is valid in both the heap and its abstract interpretation, then the results will agree: given $\mathcal{W}alk(\text{H}, \iota, \text{path}) = \text{val}$ and $\mathcal{W}(\text{u}, \text{path}) = \text{t}'$ then H, $\iota_0 \vdash \text{val} \triangleright \text{t}'$.

Both the heap and its abstract interpretation are also *enclosing-correct*, which informally means that for any declared field path.C f, the enclosing object of the value of the field must be equal to the object specified by the path, relative to the object containing the field. More formally, a well-formed dynamic heap ensures $\mathcal{W}alk(\text{H}, \iota, \text{path}) = \mathcal{E}ncl(\text{H}(\mathcal{W}alk(\text{H}, \iota, \text{f})))$, where path.C f $\in \mathcal{M}embers(\mathcal{M}ix(\text{H}, \iota))$ and $\text{H}(\iota)(\text{f}) \neq \textbf{null}$. Similarly, the static semantics ensures $\mathcal{W}(\text{u}, \text{path}) = \mathcal{E}(\mathcal{W}(\text{u}, \text{f}))$, where $\mathcal{D}clType(\text{u}, \text{f}) = \text{path}.\text{C}$.

Let us now consider the definition of these functions in detail. The $\mathcal{W}$ function takes an object type u and a path path or a syntactic type T and produces an object type or a class type, if it succeeds. If the second argument is a path path, the intuition is that $\mathcal{W}$ computes a type for the object that is reached from the object described by u by traversing path in the heap. A naive approach would be to concatenate path to the path in u, but it would be hard to tell whether such a concatenated path leads to the same object as another concatenated path. The ability to decide whether two paths lead to the same object, however, is crucial for determining the subtyping relation, since only objects with identical enclosing object are compatible. For this reason, $\mathcal{W}$ returns a *canonical* representation of the combined path, namely a type. It is canonical in that the path inside the type has the form spine.$\overline{\text{f}}$. Object types can hence be compared by simple equality tests in order to determine whether they refer to the same object.

For the empty path **this**, $\mathcal{W}$ simply returns u (first case). For paths ending in **out**, the function $\mathcal{E}$ is used to find the enclosing type (second case). Paths ending in a field or a class are checked for validity: an appropriate field or class must exist. The last case in $\mathcal{W}$ extends the domain of the second argument to T; this is the only case where $\mathcal{W}$ returns a class type. As an exam-

Typing domains:

$$u ::= \langle p \rangle.\overline{f} \qquad q ::= \textbf{this} \mid \textbf{out} \mid f$$
$$s ::= \langle p \rangle.\overline{f}.C \qquad Q ::= \overline{q} \mid \overline{q}.C$$
$$t ::= u \mid s$$

Expression Typing:

$$\frac{\mathcal{M}(t) \neq \bot}{p \vdash \textbf{null} : t} \quad (T1) \qquad \frac{\mathcal{W}(\langle p \rangle, path) = u}{p \vdash path : u} \quad (T3)$$

$$\frac{p \vdash e : t \quad p \vdash e' : t'}{p \vdash e\,;e' : t'} \quad (T2) \qquad \frac{p \vdash path : u \quad \mathcal{W}(u, \mathcal{D}clType(u, v)) = s}{p \vdash path.v : s} \quad (T4)$$

$$\frac{p \vdash path.v : s \quad p \vdash e : t \quad \mathcal{C}(t) <: s}{p \vdash path.v = e : t} \quad (T5)$$

$$\begin{array}{c} p \vdash path : u \quad p' \in \mathcal{M}(u.C) \quad p \vdash \overline{e} : \overline{t} \\ Constr(p') = T_0\, C(\overline{T}\,\overline{f})\, \dots \quad |\overline{T}| = |\overline{t}| \\ s_i = \begin{cases} \mathcal{W}(u, \textbf{this}.Q) & \text{if } T_i = \textbf{this}.\textbf{out}.Q \\ \mathcal{W}(u_j, \textbf{this}.Q) & \text{if } T_i = \textbf{this}.f_j.Q \wedge t_j = u_j \\ \quad \text{for } i = 0 \dots |\overline{t}| \end{cases} \\ \mathcal{C}(t_i) <: s_i \text{ for } i = 1 \dots |\overline{t}| \\ \hline p \vdash \textbf{new } path.C(\overline{e}) : s_0 \end{array} \quad (T6)$$

Conversion to class types:

$$\begin{array}{rcl} \mathcal{C} & :: & t \to s \\ \mathcal{C}(\langle p.C \rangle) & = & \langle p \rangle.C \\ \mathcal{C}(u.f) & = & \mathcal{W}(u, \mathcal{D}clType(u, f)) \\ \mathcal{C}(s) & = & s \end{array}$$

Mixins:

$$\begin{array}{rcl} \mathcal{M} & :: & t \to \overline{p} \\ \mathcal{M}(\langle \rangle) & = & [\textsf{nil}_c] \\ \mathcal{M}(u.C) & = & \mathcal{A}ssemble(\mathcal{M}(u), C) \\ \mathcal{M}(u) & = & \mathcal{M}(\mathcal{C}(u)) \end{array}$$

Enclosing object type:

$$\begin{array}{rcl} \mathcal{E} & :: & t \to u \\ \mathcal{E}(u.C) & = & u \\ \mathcal{E}(u) & = & \mathcal{E}(\mathcal{C}(u)) \end{array}$$

Static lookup:

$$\begin{array}{rcl} \mathcal{W} & :: & u \times (path \cup T) \to t \\ \mathcal{W}(u, \textbf{this}) & = & u \\ \mathcal{W}(u, spine.\textbf{out}) & = & \mathcal{E}(\mathcal{W}(u, spine)) \\ \mathcal{W}(u, path.f) & = & \mathcal{W}(u, path).f \text{ if } \mathcal{E}xists(\mathcal{W}(u, path), f) \\ \mathcal{W}(u, path.C) & = & \mathcal{W}(u, path).C \text{ if } \mathcal{E}xists(\mathcal{W}(u, path), C) \end{array}$$

Program Typing:

$$\frac{\mathcal{M}(\langle p \rangle.C) \neq \bot}{p \vdash C\ OK} \quad (WF1) \qquad \frac{\mathcal{W}(\langle p \rangle, T) \neq \bot}{p \vdash T\ OK} \quad (WF2)$$

$$\frac{C = C' \Rightarrow T = T', \overline{T}\,\overline{f} = \overline{T}'\,\overline{f}'}{T\, C(\overline{T}\,\overline{f})\ \{e;\}\ \text{overrides}\ T'\, C'(\overline{T}'\,\overline{f}')\ \{e';\}\ OK} \quad (WF3)$$

$$\begin{array}{c} K = T\, C(\overline{T}''\,\overline{f}')\ \{\, e;\} \quad \mathcal{M}(\langle p \rangle.C) = \overline{p} \\ Members(\overline{p}) = \overline{T}''\,\overline{f}',\ \_ \\ p \vdash \overline{C}\ OK \quad p.C \vdash \overline{T}\ OK \quad p.C \vdash \overline{T}'\ OK \quad p.C \vdash T\ OK \\ p.C \vdash e : t \quad \mathcal{C}(t) <: \mathcal{W}(\langle p.C \rangle, T) \\ K' = Constr(p_j) \Rightarrow K\ \text{overrides}\ K'\ OK \\ \hline p \vdash\ \textbf{class } C \textbf{ extends } \overline{C}\ \{K\ \overline{CL};\ \overline{T}\,\overline{f};\, \overline{T}'\,\overline{v}\}\ OK \end{array} \quad (WF4)$$

There is a strict partial order $\sqsubset_f$ on field names such that
$$\forall p, f. \quad spine.\overline{f}.C\, f \in Members(p) \Rightarrow \forall i.\ f_i \sqsubset_f f$$

There is a strict partial order $\sqsubset_c$ on class names such that

$$\frac{\forall p.\ CT(p) = \textbf{class } C \textbf{ extends } \overline{C}\ \{\dots\} \Rightarrow \forall i.\ C_i \sqsubset_c C}{CT \text{ is acyclic}} \quad (WF5)$$

$$\begin{array}{c} CT \text{ is acyclic} \\ \forall p, p', C : CT(p.C) \neq \bot, CT(p'.C) \neq \bot \Rightarrow \\ p''.C \in \mathcal{M}(\langle p \rangle.C) \cap \mathcal{M}(\langle p' \rangle.C) \\ \forall p \neq p' : CT(p) = \textbf{class } C \dots \{K\ \overline{CL};\ \overline{T}\,\overline{f}; \overline{T}'\,\overline{v}\} \\ CT(p') = \textbf{class } C' \dots \{K'\ \overline{CL}';\ \overline{T}''\,\overline{f}'; \overline{T}'''\,\overline{v}'\} \\ \Rightarrow \overline{f} \cap \overline{f}' = \emptyset, \overline{v} \cap \overline{v}' = \emptyset \\ \forall p, C : CT(p.C) \neq \bot \Rightarrow p \vdash CT(p.C)\ OK \\ \hline CT\ OK \end{array} \quad (WF6)$$

Subtyping:

$$s <: s \quad (\text{S-Refl}) \qquad \frac{s <: s' \quad s' <: s''}{s <: s''} \quad (\text{S-Trans})$$

$$\frac{\mathcal{M}(u) = \overline{p} \quad CT(p_j.C) = \textbf{class } C \textbf{ extends } \dots C' \dots}{u.C <: u.C'} \quad (\text{S-Decl})$$

Declared type of member:

$$\begin{array}{rcl} \mathcal{D}clType(t, m) & = & T \quad \text{where } T\, m \in Members(\mathcal{M}(t)) \\ \mathcal{E}xists(t, m) & = & (\mathcal{D}clType(t, m) \neq \bot) \\ \mathcal{E}xists(u, C) & = & (\mathcal{M}(u.C) \neq \bot) \end{array}$$

Figure 10. Typing rules

ple based on the definitions in Figures 1 and 2, we would have $\mathcal{W}(\langle \textsf{WithNeg}.\textsf{Neg} \rangle.e, \textbf{this}.\textbf{out}.\textsf{Lit}) = \langle \textsf{WithNeg} \rangle.\textsf{Lit}$.

Object types can be converted into class types by means of the $\mathcal{C}$ function as follows: If the object type is just a static path and no field accesses, then the enclosing object is described by the same static path with the last element removed, and the class is that last element (first case). If the object type ends with a field, the field is replaced by its declared type ($\mathcal{D}clType$ is explained below) and the $\mathcal{W}$ is called to normalize the resulting path (second case). If the type is already a class type, there is nothing to do (third case).

The $\mathcal{M}$ function computes the statically known mixin structure of an object described by a type. The type $\langle \rangle$ describes the root object which has only one mixin, namely the empty class path: $[\textsf{nil}_c]$ (first case). For an object type $u.C$, u is a type that describes the enclosing object, hence its mixin list can be recur-

sively computed from the enclosing object. This mixin list and the class name C are sufficient to compute the mixin list for this type by calling the $\mathcal{A}ssemble$ function (second case). Finally, to compute the mixin list of an object type it is first converted to a class type (third case). For example, with the code in Figure 1-5 the mixin lists are $\mathcal{M}(\langle \textsf{Test} \rangle.\textsf{f1}.\textsf{Neg}) = [\textsf{Base}.\textsf{Exp}, \textsf{WithNeg}.\textsf{Neg}]$ and $\mathcal{M}(\langle \textsf{Test} \rangle.\textsf{f2}.\textsf{Exp}) = [\textsf{Base}.\textsf{Exp}, \textsf{WithEval}.\textsf{Exp}]$.

The $\mathcal{D}clType$ function uses $\mathcal{M}$ to look up a field or variable declaration in the mixin list of a given type.

$\mathcal{C}, \mathcal{E}, \mathcal{W}, \mathcal{M}$ and $\mathcal{D}clType$ depend on each other in non-trivial ways, so it is not obvious that evaluation of these functions will terminate. A proof is given in [10]. Informally, the functions terminate because the arguments to recursive calls of $\mathcal{W}$ inside $\mathcal{W}$ and $\mathcal{D}clType$ are smaller, and the recursive call inside $\mathcal{C}$ replaces a field by its declared type. The latter case is also guaranteed to termi-

nate because programs are well-formed only if there are no cyclic dependencies on field types, as explained later in this section.

5.3 Subtyping

Subtyping determines the compatibility of values for assignment or parameter binding. It is defined only on class types but object types can always be converted to class types via $\mathcal{C}$. The main rule for the subtyping relation, (S-DECL), defines type compatibility through a combination of path equality and examination of declared subclass relationships. The latter is standard in object-oriented type systems: a class B is a subtype of A if B is derived by subclassing from A. This traditional definition is modified in vc to take into account virtual classes: two classes can only be in a subtype relation *if they are contained in the same object*; this is a concrete manifestation of the fact that types depend on the enclosing object. Rule (S-DECL) ensures that subtypes are always based on the same object type u. Since an object type describes a path to an object, the enclosing objects must be identical. This comparison for identical enclosing object types works because object types are paths in a normalized form.

5.4 Expression Typing

Expressions are given a type in the context of a static path p which describes the current object **this**. As in the operational semantics, an environment is not needed because method parameters are encoded as fields.

The **null** value (T1) has any meaningful type, whereby "meaningful" is checked by ensuring that the type has mixins. The type of a sequence is the type of the last expression in the sequence (T2). Paths (T3) are given a type using the static lookup function $\mathcal{W}$ explained in Section 5.2. As is obvious from the definition, paths have an object type. Variable lookup (T4) also uses $\mathcal{W}$, but in this case the type of the variable is passed instead of the variable name. This is a manifestation of the fact that variables cannot be used in types. This also means, however, that the type of a variable access is always a class type, not an object type.

An assignment (T5) is checked by computing a type for the left hand side, which is known to be a class type by (T4), computing a type for the right hand side and then checking whether the left side is a subtype of the right side. If the left hand type is an object type, it is converted to a class type first.

The rule for object creation (T6) is the most complex, which is not surprising given that it also handles method calls. First, the type of the enclosing object u is computed. The statically known mixin structure of the new object, $\mathcal{M}(\mathsf{u.C})$, is computed, and a mixin is selected via the choice of $\mathsf{p'}$, which is then used to find the constructor signature. Note that all mixins will provide the same signature due to program well-formedness. The types of the arguments are computed; their number must be equal to the number of constructor arguments. The actual set of mixins at runtime may be larger than the statically known set, but program well-formedness ensures that the signature of the most specific constructor at runtime is identical to the one in the statically selected constructor.

To compare the syntactic types specified in the constructor with the types of the actual arguments, class types s_i are computed for every syntactic type in the constructor, including the return type. Intuitively, the syntactic types T_i must be adapted to the *viewpoint* p. To do that, the static lookup function $\mathcal{W}$ is used again. The types T_i are either of the form **this**.**out**.... or **this**.f_j...., depending on whether the argument type comes from the environment or another argument. (Syntactically, T_i could also have the form **this**.$\mathsf{C'}$ for some class name $\mathsf{C'}$, but this type would not be useful because it would refer to a virtual class of an object that does not yet exist.)

The first case applies to the traditional situation where the type of the argument is taken from the environment; TestLit in Figure 1

$$
\begin{array}{c}
\mathsf{Test} \vdash \mathbf{this} : \langle\mathsf{Test}\rangle. \quad \mathcal{M}(\langle\mathsf{Test}\rangle.\mathsf{buildNeg}) = \mathsf{Test.buildNeg} \\
\mathsf{Test} \vdash \mathsf{f2} : \langle\mathsf{Test}\rangle.\mathsf{f2} \qquad \mathsf{Test} \vdash \mathsf{f2.zero} : \langle\mathsf{Test}\rangle.\mathsf{f2.Lit} \\
\mathcal{C}onstr(\mathsf{Test.buildNeg}) = \\
\mathsf{ne.Neg}\ \mathsf{buildNeg}(\mathsf{out.out.WithNeg}\ \mathsf{ne},\ \mathsf{ne.Exp}\ \mathsf{ex}) \\
\mathsf{s}_0 = \mathcal{W}(\langle\mathsf{Test}\rangle.\mathsf{f2}, \mathbf{this}.\mathsf{Neg}) = \langle\mathsf{Test}\rangle.\mathsf{f2.Neg} \\
\mathsf{s}_1 = \mathcal{W}(\langle\mathsf{Test}\rangle., \mathbf{this}.\mathbf{out}.\mathsf{WithNeg}) = \langle\mathsf{WithNeg}\rangle. \\
\mathsf{s}_2 = \mathcal{W}(\langle\mathsf{Test}\rangle.\mathsf{f2}, \mathbf{this}.\mathsf{Exp}) = \langle\mathsf{Test}\rangle.\mathsf{f2.Exp} \\
\mathcal{C}(\langle\mathsf{Test}\rangle.\mathsf{f2}) = \langle\mathsf{NegAndEval}\rangle. <: \mathsf{s}_1 \\
\mathcal{C}(\langle\mathsf{Test}\rangle.\mathsf{f2.Lit}) = \langle\mathsf{Test}\rangle.\mathsf{f2.Lit} <: \mathsf{s}_2 \\
\hline
\mathsf{Test} \vdash \mathbf{new}\ \mathbf{this}.\mathsf{buildNeg}(\mathsf{f2},\mathsf{f2.zero}) : \langle\mathsf{Test}\rangle.\mathsf{f2.Neg}
\end{array}
$$

Figure 11. Type derivation for buildNeg(f2,f2.zero) in Figure 5

is an example. In this case, **this**.**out** refers to the enclosing object of the class. The type of this enclosing object is the type of path, or the object type u. The actual argument type s_i is then found by navigating from u into the tail of T_i using $\mathcal{W}$.

The latter case applies if an argument type depends on the virtual class of another argument, as for example buildNeg in Figure 5. In this case, f_j is initialized with the value of e_j at runtime. The actual argument type s_i is then found by navigating from t_j into the tail of T_i using $\mathcal{W}$. If an argument is used as type provider for another argument, then the expression for the argument needs to have an object type. This restriction is enforced by the condition $\mathsf{u}_j = \mathsf{t}_j$ in (T6).

The complete list of argument types s_i is then checked to be subtypes of the formal argument types. Finally, the viewpoint-adapted constructor return type s_0 is returned.

Figure 11 shows an example of a non-trivial usage of (T6) in the example from Figure 5. It has been slightly adjusted to fit to the formal syntax, see Section 3.3. The example illustrates only the last step in the typing derivation, the result of sub-derivations has been inlined. Notice in particular that the type of the expression contains the information that the result has the family f2.

5.5 Program Typing

In order to separate out the problem of cyclic inheritance relations and cyclic field type dependencies (the type of a field may depend on the value of other fields), declared names are partially ordered such that each of the two kinds of dependencies are known to be acyclic (WF5). Consequently, cyclic inheritance relations and cyclic relations via dependent types (which are expressed using fields) cannot occur. We could relax this restriction without affecting soundness, but with the current strict ruleset it is easy to see that the type analysis always terminates, without adding special checks for infinite loops in type computations.

The overall program well-formedness rule, (WF6), requires that the program is acyclic, that two class declarations of the same class name have a shared mixin, that field and variable declarations are unique, and that each class declaration is well-formed.

A class is OK (WF4) if the list of constructor arguments matches the list of fields in the statically known mixin structure of the class, if all superclasses are valid, if the type of the constructor expression is compatible to the declared return type, and if all other mixins that have the same class name have the same constructor signature, see also (WF3). The validity of superclass and type declarations ((WF1) and (WF2)) is checked using the $\mathcal{M}$ and $\mathcal{W}$ functions, which return $\bot$ if the class or type, respectively, is not known to exist in the context p.

Note that (WF4) implies that fields can only be declared in new class declarations (i.e., if there is no inherited class declaration with the same name); this restriction is not essential and we could easily add initialized fields (declared as $\mathsf{T}\ \mathsf{f} = \mathsf{e}$) to the calculus which could be declared in all classes. (In fact, we developed the whole

calculus with initialized and redefinable fields before we decided to add constructors and let fields be initialized via constructor arguments.) We have chosen to leave out initialized fields because they do add a number of details to rules, but do not provide much additional insight. We could also have allowed field declarations everywhere and accepted the possibility for additional run-time NullErr errors due to uninitialized fields, but we felt that the current strict approach is useful because it illustrates how to statically ensure that all fields are initialized. Also note that the restriction on fields does not affect the ability to declare variables and classes (possibly used as methods) in all class declarations, so there are no restrictions on ordinary width subtyping in the calculus.

As mentioned, (WF6) requires globally unique member names; that is, field and variable names must be unique throughout the program. This may seem like a serious restriction that could interfere with separate compilation, but it is in fact just a simple way to emulate an approach which is usable in a full-fledged language and which does not interfere with separate compilation. In particular, the gbeta compilation process extends all declared names with a unique identification of the enclosing class body (i.e., something that corresponds to the static path to the scope of the declaration). It is then resolved statically which name declaration each name usage refers to, and the name usage is then extended correspondingly. As a result, if a given object contains multiple members named m, they will at run-time be distinct members with extended names p_1_m, p_2_m, etc., and name usages will use these extended names for lookups. Hence, field and variable lookup uses early binding, which is also the desired semantics. In Caesar, such name clashes are detected and rejected at compile time, so the programmer has to rename one of the features in case of a clash.

For class or method lookup the desired semantics is late binding, so in this case the technique is slightly different. (WF6) requires any two declarations of a class with the same name to have a shared declaration of that class in their statically known sets of mixins. This global restriction may seem to interfere with separate compilation. However, it can be removed in a way which is similar to the one used for members. First, note that in *vc* it is easy to show that for a given class name C there must be a unique declaration of C which is in this sense shared among all declarations of C. In gbeta it is required that an "introductory" class declaration—i.e., one where no other declarations of the same class are known statically—is marked syntactically, not unlike the distinction between `virtual` and `override` methods in $C\#$. Each introductory declaration for a class is renamed with an identification of its enclosing class body, just like a member declaration. Each non-introductory class declaration is renamed like a member name *usage* to have the same extended name as its introduction. This implies that every class declaration has one particular introduction, which is resolved statically. Finally, class name usages are renamed to be like their extended statically known declarations. As a result, there is no need for global restrictions, and it is possible for multiple classes with the same name to coexist in the same object. With respect to binding time, there is early binding of the choice of class introduction (class *identity*), but late binding of the actual value (the dynamic set of mixins). Our formalization is thus much simpler, but it models the approach taken in full-fledged languages in a faithful albeit not always direct manner.

6. Wellformed Heaps and Agreement

The soundness of the operational semantics with respect to the type system depends upon having a well-formed heap, and agreement between a value and a type relative to a heap. The rules for heap well-formedness and agreement are given in Figure 12. Since the details of these definitions are not required to understand the *vc* calculus as such, the remainder of this section can be skipped

Well-formedness:

$$\frac{H(\iota)(m) = \textbf{null}}{\iota.m : T \ OK \ in \ H} \qquad \text{(WF-NULL)}$$

$$\frac{H(\iota)(m) = \iota' \quad}{\mathcal{W}alk(H, \iota', \textbf{out}) = \mathcal{W}alk(H, \iota, \text{path}) \quad p.C \in \mathcal{M}ix(H, \iota')}{\iota.m : \text{path}.C \ OK \ in \ H}$$
$$\text{(WF-MEM)}$$

$$\frac{T \ m \in \mathcal{M}embers(\mathcal{M}ix(H, \iota)) \Rightarrow \iota.m : T \ OK \ in \ H}{\iota \ OK \ in \ H} \quad \text{(WF-OBJ)}$$

$$\frac{H(\iota_{\text{root}}) = [\![\ \bot \ \| \ C_{\text{root}} \ \| \ ]\!]}{\iota_{\text{root}} \ OK \ in \ H} \qquad \text{(WF-ROOT)}$$

$$\frac{\forall \iota. \ \iota \ OK \ in \ H}{H \ OK} \qquad \text{(WF-HEAP)}$$

Agreement:

$$H, \iota_0 \vdash \textbf{null} \triangleright \text{val}_s \qquad \text{(A-NULL)}$$

$$H, \iota_0 \vdash \iota_{\text{root}} \triangleright \langle \rangle \qquad \text{(A-ROOT)}$$

$$\frac{j = \mathcal{D}epth(H, \iota_0) - |p|}{\mathcal{W}alk(H, \iota_0, \textbf{this}.\textbf{out}^j.\overline{f}) = \iota \quad H, \iota_0 \vdash \iota \triangleright \mathcal{C}(\langle p \rangle.\overline{f})}{H, \iota_0 \vdash \iota \triangleright \langle p \rangle.\overline{f}} \quad \text{(A-OTYPE)}$$

$$\frac{p'.C \in \mathcal{M}ix(H, \iota) \quad H, \iota_0 \vdash \mathcal{E}ncl(H(\iota)) \triangleright \mathcal{E}(u.C)}{H, \iota_0 \vdash \iota \triangleright u.C} \quad \text{(A-CTYPE)}$$

Auxiliary definitions:

$$\mathcal{D}epth(H, \iota) = \begin{cases} 0, & \text{if } \iota = \iota_{\text{root}} \\ 1 + \mathcal{D}epth(H, \mathcal{E}ncl(H(\iota))) \end{cases}$$

Figure 12. Dynamic well-formedness and agreement

by readers who are less interested in how the soundness result is reached.

A heap is well-formed if all its objects are well-formed (WF-HEAP). An object is well-formed if all its members are well-formed (WF-OBJ). An object member is well-formed if its value in the heap is null (WF-NULL). Otherwise a member m of object ι is well-formed if the member value $\iota' = H(\iota)(m)$ satisfies two conditions: (1) the enclosing object of the value, $\mathcal{W}alk(H, \iota', \textbf{out})$, is equal to the object $\mathcal{W}alk(H, \iota, \text{path})$ specified by the path in the declared type path.C; and (2) the mixins of the value, $\mathcal{M}ix(H, \iota')$, include a path ending with the class C. There is a special rule for well-formedness of the root object because it does not have an enclosing object.

Type agreement is specified as the agreement of an object at ι with a type T, relative to a dynamic heap H and a starting point ι_0. The starting point specifies an address in the dynamic heap that is related to the base of the type. **null** agrees with all types (A-NULL), and the root object agrees with the empty object type (A-ROOT).

Rule (A-OTYPE) handles object types, $\langle p \rangle.\overline{f}$. The rules ensure that the class path p is a prefix of the spine of ι_0, so the value j represents the number of enclosing objects that must be traversed from ι_0 to read an object with the same depth as p. The path **this**.**out**^{j}.$\overline{f}$ traverses to this object, and then traverses the field list

$\overline{f}$. The object ι must be located at the end of this path. In addition, ι must agree with the corresponding class type.

Rule (A-CTYPE) handles class types, $\langle p \rangle.\overline{f}.C$. It requires that the mixins of the value, $\mathcal{M}ix(H, \iota)$, include a path ending with the type's class C. It also requires that the actual enclosing object agrees with the enclosing type.

7. Soundness

The type system of *vc* is sound in the sense that a well-typed expression either returns a value that agrees with its type, terminates with a NullErr, or diverges, but never terminates with a TypeErr. The soundness result is composed of two formal results: *preservation* and *coverage*. Preservation is the standard theorem which characterizes the result of expressions that are well-typed and evaluate to a result. Coverage is a new technique for ensuring that errors do not prevent expressions from evaluating to a result.

Preservation assumes a valid program and heap. Given the static path p of a class in which an expression e has type t, and the address ι of an object that agrees with p; if the expression evaluates to a result r then either the result is NullErr or it is a value that agrees with t. Preservation also guarantees that the heap is still well-formed after the execution, and that the current object still agrees with its type.

THEOREM 1 (Preservation).

$$\begin{bmatrix} CT \ \mathsf{OK} \\ \mathsf{H} \ \mathsf{OK} \\ \mathsf{p} \vdash \mathsf{e} : \mathsf{t} \\ \mathsf{H}, \iota \vdash \iota \rhd \langle \mathsf{p} \rangle \\ \mathsf{e}, \mathsf{H}, \iota \leadsto \mathsf{r}, \mathsf{H}' \end{bmatrix} \Rightarrow \begin{bmatrix} \mathsf{H}' \ \mathsf{OK} \\ \mathsf{H}', \iota \vdash \iota \rhd \langle \mathsf{p} \rangle \\ \mathsf{r} = \mathsf{val} \ \wedge \ \mathsf{H}', \iota \vdash \mathsf{val} \rhd \mathsf{t} \\ \vee \\ \mathsf{r} = \mathsf{NullErr} \end{bmatrix}$$

This theorem only characterizes evaluations that terminate, which is a natural consequence of using a big-step semantics. Hence it is slightly weaker than the usual "progress and preservation" theorems in a small-step semantics, where it can be expressed that execution of a type correct program will never get stuck even if the execution continues forever.

Preservation alone does not ensure soundness however, because an expression may fail to evaluate due to a missing case in the evaluation rules. We have followed standard practice by including rules (ER1-4) to cover a variety of error cases in evaluation [14]. The complete list of error rules is given along with the soundness proof. The second half of our soundness proof ensures that *all* error cases have been handled. As a result, the only way an evaluation can fail to produce a value is if the computation diverges. This Lemma plays a role similar to the 'progress' theorem when using a small-step semantics.

The purpose of the coverage lemma is to show that the evaluation rules always produce a value unless the computation diverges. First we define a notion of finite evaluation. If the evaluation exceeds the bound for finite evaluation, it produces a special termination value. The evaluation rules for error propagation propagate this special value.

DEFINITION 1 (Finite Evaluation). *Define an evaluation relation $\leadsto_k$ as a copy of the rules for $\leadsto$. Replace each occurrence of $\leadsto$ in a premise by $\leadsto_{k-1}$. Replace $\leadsto$ in the conclusion of each rule and axiom with $\leadsto_k$. Note that the copied axioms are defined for all k. Add the following axiom:*

$$\mathsf{e}, \mathsf{H}, \iota \leadsto_0 \mathsf{KillErr}, \mathsf{H} \qquad (\text{KILL})$$

The finite evaluation relation $\leadsto_n$ returns KillErr if the derivation is more than n derivations deep. It is thus a finite approximation of the normal evaluation of an expression. The coverage lemma states that finite evaluation always produces a value.

LEMMA 1 (Coverage). *For all natural numbers n and e, H, ι, there exists r, H' such that*

$$\mathsf{e}, \mathsf{H}, \iota \leadsto_n \mathsf{r}, \mathsf{H}'$$

The coverage lemma ensures that the operational semantics produces a value even in the face of runtime errors, such as access to non-existing members, see (ER2) and (ER3) in Figure 8.

A terminating expression is one for which there is an n such that finite evaluation $\leadsto_n$ does not return KillErr. If the expression does not return KillErr, then it cannot use the KILL axiom. As a result, the derivation in $\leadsto_n$ can be translated to a derivation in $\leadsto$. Thus every terminating expression has a corresponding derivation in $\leadsto$.

Theorem 1 and Lemma 1 ensure the soundness of *vc*: execution of well-typed expressions will either produce a value of the correct type, return NullErr, or else diverge. But evaluation will never access non-existing fields, variables, or classes, and is never stuck.

Note that all proofs are provided in [10].

8. Related and Future Work

The idea of virtual classes and their different kinds of bindings stems from BETA [21]. The concept of virtual superclasses was explored but never fully realized in BETA and has not been supported in the BETA compiler since the early eighties. Virtual classes in their general form as defined in this paper have been presented informally in the works on family polymorphism and higher-order hierarchies in gbeta [8, 9], delegation layers [28], and Caesar [23]. *vc* represents the core of these languages.

In gbeta, classes can have superclasses of the form path.C, which enables a new kind of dynamic composition that is not expressible in *vc*. However, we have analyzed the required extensions to *vc* in order to support this kind of inheritance, and based on the experience from gbeta it does not seem very hard, although it does introduce many new details in the rules and proofs (essentially, mixins must be on the form $\langle u \rangle.p$ rather than simply p, and each member in an object must have its own enclosing object). We expect to explore this extension in some future work. Delegation layers are more dynamic than *vc* in that they use object-based delegation instead of class-based inheritance, which enables polymorphic composition of types at runtime. It is also a natural part of our future work to create a version of *vc* building on delegation, but in this case it is not obvious how hard it is. In Caesar, virtual classes are combined with some aspect-oriented mechanisms which make the language very suitable for integrating independently-developed software components. As in *vc*, both Caesar and gbeta distinguish mutable variables from immutable fields and use this distinction during type checking.

Odersky et al have presented a calculus with path-dependent types called νObj [26]. The most important difference to νObj is that *vc* allows virtual classes whereas νObj focuses on virtual types only. This means that no objects can be created as an instance of a virtual type (abstract type member) and no implementation can be specified before the virtual type is final-bound to a concrete type. Although it is possible to create a class that has a virtual super-class in νObj, this mechanism cannot express hierarchy specialization because the virtual superclass can only be replaced by a class that has *exactly* the same signature (e.g., does not add methods) [37]. Another difference is that *vc* has assignments, whereas νObj is purely functional. On the other hand, νObj is more powerful than *vc* w.r.t. the encoding of parametric polymorphism, which is not in the focus of this work. Finally, since our type-checker is completely syntax-directed (in particular, we have no subsumption rule), type-checking in *vc* is decidable, which is not the case for νObj.

In [25], a language with nested inheritance is described, which has a number of similarities with virtual classes. An important difference to their approach is that they use classes in classes rather

than classes in objects. The classes-in-classes model can trivially be simulated in a classes-in-objects model by using only one instance of each class containing virtual classes, but the converse does not hold—e.g., nested classes in [25] cannot access shared state of instances of enclosing classes. For example, in *vc* every nested class in Base and its subclasses can access the zero field declared in Figure 1. The expressive power of having access to the enclosing object is also illustrated by our straightforward encoding of methods by means of classes – accessing an instance variable foo of an object in a method bar is encoded as an access to the enclosing object **out**.foo in the corresponding class bar. Using nested inheritance, it would be possible to manually declare an instance variable, say 'enclosing', in each nested class and thus emulate the enclosing object, but it would require significantly more work to create and administrate such simulated enclosing objects, and it is not obvious that they could be given all the desired typing properties.

Another consequence is that a given program using nested inheritance has a fixed number of class families, whereas a given program in *vc* can have an unlimited number of distinct class families because every new family object contains a new class family. This enables a more fine-grained typing discipline in *vc*, because the type system will ensure that all these families are not mixed up. For example, this could be used to ensure that instances of Student nested in a given University are used only with the university from which they were obtained. With nested inheritance a simple instanceof test could reveal that all students were in fact members of the same class family, and hence the connection between a specific university and the associated students could not be expressed or enforced.

Family polymorphism by means of passing an instance of the enclosing class cannot be done directly in a classes-in-classes model. Instead, the authors of [25] propose the notion of *prefix types* to achieve a similar kind of polymorphism. Prefix types are a mechanism to refer to the (statically unknown) enclosing class of the class of an object. For example, A[b.class] denotes the enclosing class of the class of the object b.

The nested inheritance language itself is much bigger (and hence more complex) than our language. For example, there are seven different syntactic forms of type declarations and type schemas in [25], whereas the only form of type declaration is path.C in *vc*. Yet another difference is that different extensions to a class hierarchy cannot be combined in the nested inheritance language, as illustrated by our example in Figure 4. This is a consequence of the requirement in nested inheritance that the declared superclass of a class C must be a subtype of the inherited version of C, i.e., declared superclasses in redefinitions cannot be used to mix in additional features.

One feature which is well-known from related languages and calculi (including BETA and *νObj*) is that of final-bindings. A virtual class/type may be final-bound, which means that the current value must remain unchanged (e.g., no additional mixins can be included). This feature is useful because it provides a lower bound on the value of a class, which opens more opportunities for assignments to variables having a given virtual class/type as their declared type. It would hence make sense to add final bindings to *vc* as well, but this extension is orthogonal to our work because our focus is on extensibility and not on genericity. Moreover, many years of experience with BETA seems to indicate that final bounds are not that important when initialized immutable fields are available, because such fields can be used to obtain a lower bound on all virtual classes in a given object. It is likely that the trade-off is different in languages like *νObj* and Scala [27], because many details in the language design are different and closer to the functional paradigm.

There are a couple of other approaches that widen the expressibility of the static type system with respect to collaborating classes

and parametric polymorphism but do not support incremental hierarchy specification [4, 33, 16].

Thorup proposes a virtual type system for Java [32]. It supports instantiation of a virtual class and hence late bound classes, but it does not support virtual superclasses. Furthermore, the type system relies on dynamic type checks.

There have been a couple of approaches for hierarchy refinement in the context of product lines (e.g., [2, 30]) but polymorphic usage of a hierarchy variant is not in the focus of these works. It will be interesting to explore how virtual classes improve the expressibility of languages with respect to product lines. Virtual classes are interesting from a software architecture point of view because they enable both incremental specification of class hierarchies and composition of different extensions to a class hierarchy, a problem that is hard to solve in conventional object-oriented languages. Hence, the language constructs in *vc* are well-suited to implement layered software architectures like mixin layers [30] or GenVoca [2].

Family classes used as argument types give rise to covariant typing, which is known to be non-trivial to handle in a type-safe manner. Other examples of a strict and safe treatment of covariance are the formalization of variant parametric types in [18], and the inclusion of wildcards into the J2SE 5 version of the Java platform [35]. Note, however, that virtual classes are different from variant parametric types or parametric types with wildcards, because those mechanisms do not support family polymorphism, but they provide a different kind of flexibility through structural equivalence among type applications.

The notion of having a first-class representation of a hierarchy is also highly relevant to the domain of aspect-oriented programming, which can be seen as an approach to have multiple cross-cutting decompositions (that is, hierarchies) of a system [31, 24].

The only prior work related to our coverage lemma that we know of is a paper by Fisher and Reppy [11]. They also improve on the traditional approach to proving type soundness for big-step semantics by differentiating diverging expressions from errors. They use an 'evaluation height function', whose definition is similar in structure to a small-step operational semantics, to count the number of steps during evaluation. Their soundness proof involves showing that a well-typed term with an evaluation height of n will always evaluate to a value of the correct type. They *define* diverging programs as those for which the evaluation height function is undefined, but there is no proof that the evaluation height function correctly characterizes divergence of the operational semantics. In our technique, the correspondence between $\rightsquigarrow$ and $\rightsquigarrow_k$ is obvious by construction, all non-diverging programs have an evaluation tree because of the error rules, and missing rules are prevented due to the coverage lemma. Since Fisher and Reppy do not give full proofs, it is difficult to compare our techniques in detail.

9. Conclusions

We have presented the calculus *vc* of virtual classes with path-dependent types, described its dynamic and static semantics, and proved soundness. The approach to static analysis which was pioneered in BETA, made strict and complete in gbeta, and adapted for Java-like languages in Caesar has thereby been documented, clarified, and characterized as fundamentally sound. Our calculus has certain uniqueness requirements on declared names, but we have explained how these restrictions have been lifted in a full-fledged language at the cost of some extra complexity. All in all, we have hereby provided a foundation which shows that the widespread image of virtual classes as being inherently unsafe is too pessimistic.

and to Gary T. Leavens who helped us explaining several issues more clearly.

References

[1] K. Barrett, B. Cassels, P. Haahr, D. Moon, K. Playford, and P. T. Withington. A monotonic superclass linearization for Dylan. In *Proceedings OOPSLA '96*, pages 69–82. ACM Press, 1996.

[2] D. Batory, V. Singhal, J. Thomas, S. Dasari, B. Geraci, and M. Sirkin. The genvoca model of software-system generators. *IEEE Software*, 11(5), 1994.

[3] G. Bracha and W. Cook. Mixin-based inheritance. In *Proceedings OOPSLA/ECOOP'90. ACM SIGPLAN Notices 25(10)*, pages 303–311. ACM, 1990.

[4] K. B. Bruce, M. Odersky, and P. Wadler. A statically safe alternative to virtual types. In *Proceedings ECOOP '98. LNCS 1445*, pages 523–549. Springer, 1998.

[5] W. Cook. Object-oriented programming versus abstract data types. In *Proc. of the REX Workshop/School on the Foundations of Object-Oriented Languages*, LNCS 173. Springer-Verlag, 1990.

[6] S. Drossopoulou, F. Damiani, M. Dezani-Ciancaglini, and P. Giannini. More dynamic object re-classification: FickleII. *ACM Transactions On Programming Languages and Systems*, 24(2):153–191, 2002.

[7] E. Ernst. Propagating class and method combination. In *Proceedings ECOOP'99*, LNCS 1628, pages 67–91. Springer-Verlag, 1999.

[8] E. Ernst. Family polymorphism. In *Proceedings ECOOP '01*, LNCS 2072, pages 303–326. Springer, 2001.

[9] E. Ernst. Higher-order hierarchies. In *Proceedings ECOOP '03*, LNCS. Springer, 2003.

[10] E. Ernst, K. Ostermann, and W. Cook. A virtual class calculus. Technical report, University of Aarhus, Aarhus, Denmark, 2005.

[11] K. Fisher and J. Reppy. Statically typed traits. Technical Report TR-2003-13, University of Chicago, Chicago, USA, 2003.

[12] M. Flatt, S. Krishnamurthi, and M. Felleisen. A programmer's reduction semantics for classes and mixins. In *Formal Syntax and Semantics of Java*, pages 241–269, London, UK, 1999. Springer-Verlag.

[13] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns*. Addison Wesley, 1995.

[14] C. A. Gunter. *Semantics of Programming Languages: Structures and Techniques*. Foundations of Computing. MIT Press, 1992.

[15] R. Harrejon, D. Batory, and W. R. Cook. Evaluating support for features in advanced modularization technologies. In *Proceedings ECOOP '05*. Springer, 2005.

[16] A. Igarashi and B. Pierce. Foundations for virtual types. *Information and Computation*, 175(1):34–49, 2002.

[17] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.*, 23(3):396–450, 2001.

[18] A. Igarashi and M. Viroli. On variance-based subtyping for parametric types. In *Proceedings of ECOOP '02*. Springer LNCS 2374, 2002.

[19] S. Krishnamurthi, M. Felleisen, and D. P. Friedman. Synthesizing object-oriented and functional design to promote re-use. In *Proceedings of ECOOP '98*, LNCS 1445, 1998.

[20] O. L. Madsen, B. Møller-Pedersen, and K. Nygaard. *Object Oriented Programming in the Beta Programming Language*. Addison-Wesley, 1993.

[21] O. L. Madsen and B. MÃ¸ller-Pedersen. Virtual classes: A powerful mechanism in object-oriented programming. In *Proceedings of OOPSLA '89. ACM SIGPLAN Notices 24(10)*, pages 397–406, 1989.

[22] M. Mezini and K. Ostermann. Integrating independent components with on-demand remodularization. In *Proceedings OOPSLA '02, ACM SIGPLAN Notices 37(11)*, pages 52–67, 2002.

[23] M. Mezini and K. Ostermann. Conquering aspects with Caesar. In *Proceedings AOSD '03*, pages 90–99. ACM, 2003.

[24] M. Mezini and K. Ostermann. Modules for crosscutting models. In *International Conference on Reliable Software Technologies*. Springer LNCS 2655, 2003.

[25] N. Nystrom, S. Chong, and A. C. Myers. Scalable extensibility via nested inheritance. In *Proceedings OOPSLA '04*, pages 99–115. ACM Press, 2004.

[26] M. Odersky, V. Cremet, C. RÃ¶ckl, and M. Zenger. A nominal theory of objects with dependent types. In *Proceedings ECOOP '03*. Springer LNCS, 2003.

[27] M. Odersky and M. Zenger. Scalable component abstractions. In *OOPSLA '05: Proceedings ACM SIGPLAN Conference on Object oriented programming systems languages and applications*, pages 41–57. ACM Press, 2005.

[28] K. Ostermann. Dynamically composable collaborations with delegation layers. In *Proceedings of ECOOP '02. LNCS 2374*, pages 89–110. Springer, 2002.

[29] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.

[30] Y. Smaragdakis and D. Batory. Implementing layered designs with mixin-layers. In *Proceedings of ECOOP '98, LNCS 1445*, pages 550–570, 1998.

[31] P. Tarr, H. Ossher, W. Harrison, and S. M. Sutton. N degrees of separation: Multi-dimensional separation of concerns. In *Proceedings International Conference on Software Engineering (ICSE) '99*, pages 107–119. ACM Press, 1999.

[32] K. K. Thorup. Genericity in Java with virtual types. In *Proceedings ECOOP '97. LNCS 1241*, pages 444–471, 1997.

[33] K. K. Thorup and M. Torgersen. Unifying genericity - combining the benefits of virtual types and parameterized classes. In *Proceedings ECOOP '99*, 1999.

[34] M. Torgersen. The expression problem revisited. In *European Conference on Object-Oriented Programming*, 2004.

[35] M. Torgersen, E. Ernst, C. P. Hansen, P. von der Ahé, G. Bracha, and N. Gafter. Adding wildcards to the Java programming language. *Journal of Object Technology*, 3(11):97–116, Dec. 2004. http://www.jot.fm/issues/issue_2004_12/article5.

[36] P. Wadler. The expression problem. Message to java-genericity electronic mailing list, November 1998.

[37] M. Zenger. Personal communication, 2003.

[38] M. Zenger and M. Odersky. Independently extensible solutions to the expression problem. Technical Report IC/2004/33, École Polytechnique Fédérale de Lausanne, 2004.