

A Core Calculus of Metaclasses

Sam Tobin-Hochstadt Eric Allen

Sun Microsystems Laboratories
Burlington MA 01803
<first>.<last>@sun.com

Abstract

Metaclasses provide a useful method of abstraction for programmers working with object-oriented languages, but they have not seen the formal exploration applied to more conventional object-oriented programming features. In order to elucidate the structure of metaclasses and their relationship with static typing, we present a core calculus for a nominally-typed object-oriented language with metaclasses and prove type soundness over this core. This calculus is presented as an adaptation of Featherweight GJ [13], and is powerful enough to capture metaclass relationships beyond what are expressible in common object-oriented languages, including arbitrary metaclass hierarchies and classes as values. We also explore the integration of object-oriented design patterns with metaclasses, and show how our calculus integrates several such patterns cleanly.

1 Introduction

One of the stated benefits of object-oriented languages is their ability to model aspects of the world in the object hierarchy. However, most such languages are unable to model many simple and common relationships. For example, consider the relationship between physical quantities, units of measurement, and physical dimensions. It is obvious that “length” is a dimension, and that “3 feet” is a length. However, modeling this relationship in a conventional language such as the Java™ Programming Language is impossible without attributing to some of these concepts properties which do not properly belong to them. It is natural to model `3 feet` as an instance of class `Length`. But if we were to define a class `Dimension` and define `Length` to be an instance

of class `Dimension`, then `Length` would not be a class and so `3 feet` could not be an instance of `Length`. Alternatively, we might define `Length` to be a subclass of `Dimension`. But then we still cannot define `3 feet` to be an instance of `Length` because anything that is an instance of `Length` would also be an instance of class `Dimension` and `3 feet` is obviously not a dimension.

The fundamental limitation of conventional object-oriented languages is that they do not allow a single concept to serve both as a class and as an instance. In [18], Welty and Ferrucci lay out a comprehensive case that this limitation prevents conventional modeling languages from capturing many important aspects of the world. In their primary example, they show the difficulties in modeling a simple ontology including the notions `Species` and `Eagle`, while also capturing the relationship between `Eagle` and `Harry`, a particular eagle. One solution they consider and reject is to separate the notion of “eagle” into two concepts: `the Eagle` and `Eagle`, with the former an instance of `Species` and the latter a class whose instances are specific birds. This approach is inadequate because it fails to express a relationship between the notions `the Eagle` and `Eagle`.

In order to overcome these limitations in expressiveness, several object-oriented languages that are not statically typed, such as Smalltalk, Python, and Self, allow for more flexible class relationships [11], [14], [17]. In the case of Smalltalk and Python, classes are instances of metaclasses. Static members of a class are modeled as ordinary members of the corresponding metaclass. But traditional metaclass systems have placed important limitations on expressiveness. For example, in Smalltalk, the metaclass hierarchy is only two levels deep. If we want to represent the relationships “A is an instance of B”, “B is an instance of C”, and “C is an instance of D”, we can only do this in Smalltalk if D is the special Smalltalk class `Metaclass`. This does not allow the multi-level hierarchies required to model the examples discussed above [11]. Self, on the other hand, is a prototype-based object-oriented language,

where there are no classes at all; object instantiation consists of cloning an existing instance. The members of the clone can be modified and added to at will, making static checking difficult.

An extremely expressive system of metaclasses, very similar to ours, was presented by Cointe in [6]. In his paper, he makes a number of persuasive arguments for the benefits of metaclasses. In response to complaints that the Smalltalk80 metaclass system was too complex, he suggested that a more powerful and flexible system, as he and we present, has the potential to be much simpler. He claims “The ObjVlisp model’s primary advantage is uniformity.” We agree, and believe that the treatment of metaclasses presented here unifies both disparate metaclass systems, and different ad-hoc mechanisms describing the properties of classes as objects in languages without metaclasses.

Cointe also quotes from [4]: “With respect to Simula, Smalltalk also abandons ... strong typing, allowing it to ... introduce the notion of meta-classes”. In this paper, we show that such a compromise is not necessary, and that static typing can coexist with the flexibility of metaclasses.

We present and formalize a statically typed calculus with nominal subtyping for metaclasses where every class can serve as both an instance of a class and as a class of instances, and there is no limitation on the levels of nesting of metaclasses. Our work is motivated by [2], which presents a system for integrating the static checking of dimensions of physical quantities with a conventional object-oriented type system. In that system, metaclasses are used extensively to model the type relationships of dimensions. The notion of metaclass in that system is more general than that of languages such as Smalltalk, where each class has a unique metaclass that defines its class members. In this paper, we formalize the core of that system related to metaclasses and prove it sound. In the process, we discover several interesting properties about the structure that a nominal type system for metaclasses must take.

In the interests of economy and clarity of presentation, in the following discussion we restrict ourselves to simple examples, and do not delve in the complexities of object-oriented analysis, or of dimension types. This is certainly not because we believe that metaclasses are most useful for the development of toy object-oriented programs. On the contrary, metaclasses are likely to show their value in large systems, both in terms of modeling and in the treatment of design patterns in section 3.

The remainder of this paper is organized as follows. In section 2 we present a core calculus for metaclasses we dub MCJ. We introduce the language and describe its key features, as well as some of the important design

choices. The motivating examples from the introduction are presented in MCJ in section 2.3, demonstrating the benefits of metaclasses. We discuss in section 3 how the need for many object-oriented “design patterns” is obviated through the use of metaclasses. We then describe its formal properties in section 4. In section 4.2 we provide a formal semantics for the language, and in section 4.3 we prove a type soundness result. In section 5 we consider related work and in section 6 we describe conclusions and future directions.

2 MCJ

2.1 Overview

MCJ is a generically-typed object-oriented calculus based on Featherweight GJ that includes metaclasses. An MCJ program consists of a sequence of class definitions followed by a trailing expression. This trailing expression is evaluated in the context of the class table induced by the class definitions. Before discussing the formal specification of MCJ, we first give a high-level overview of the nature of MCJ class definitions.

Each class is an instance of a *metaclass*. A metaclass is either a user-defined class or class `Object`. We say that if a class `C` is an instance of a class `D`, then `C` is an *instance class* of `D`, and that `D` is the *immediate containing class*, or *kind*, of `C`. If class `E` is a superclass of `D` then we also say that `C` is an instance class of `E` and that `E` is a containing class of `C`.

There are two key distinguishing features of MCJ. First, a class (or an instantiation of a generic class) can be used as an expression. Second, all classes have both a superclass and a kind. Instances of a class inherit behavior from the superclass. The class itself, when used in an expression context, inherits behavior from its kind.

A class definition consists of a header plus a collection of fields and methods associated with the class. The header names the class and specifies the type parameters, as well as the superclass and the kind. For example, the following header declares a class `C` with superclass `Object` and kind `D` and no type parameters:

```
class C kind D extends Object {...}
```

Each class may have both class members and instance members. Class members define the behavior of the class when used as an instance. Instance members define the behavior of instances of a class. In the abstract syntax presented in this paper, class members are distinguished from instance members by order, and with the keyword `class`¹. Members of a class definition

¹This syntax is not ambiguous since MCJ has no inner classes.

are laid out as follows: first, class fields are defined, followed by class methods, then instance fields and finally instance methods. For example, the following class definition includes a class method `m` and an instance method `n`:

```
class C kind Object extends Object {
  class Object m() {...}
  Object n() {...}
}
```

Class fields and methods bear a resemblance to static fields and methods, but there are important differences. For a given class `C`, instance members of `C`'s kind are inherited as class members of `C`. For example, a class method `m` in `C` with the same name as an instance method of `C`'s kind overrides the definition of `m` in the kind (and it must have the same signature as the overridden method). Note that if an instance class `C` of type `T` is assigned to a variable `v` of type `T`, references to the instance members of `v` refer to the class members of `C`. As in [2], a class is allowed to define both a static method and an instance method of the same name. We prevent ambiguity in member references by requiring that all references to class members in MCJ must explicitly denote the receiver.

Superclasses behave as in Featherweight GJ, providing implementation inheritance of instance behavior and subtyping. A class, when used as an expression, is an instance of its kind. If a method is invoked on a class `C`, first the class methods of `C` are examined. If the method is found there, it is invoked. Otherwise, the *instance* methods of the kind `D` of `C` are searched, and if it is found there, the method is invoked. Otherwise, the superclass hierarchy of `D` is examined as it is for an ordinary instance of `D`.

For example, consider the following class definitions:

```
class A kind Object extends Object {...}
class B kind Object extends A {...}
class C kind B extends Object {...}
```

To resolve the method call `C.m()` first the class methods of `C` are examined, and if a match is found, it is invoked. Then the *instance* methods of `B` and then `A` are examined.

If the method call were instead `b.m()`, where `b` denotes an instance of `B`, then the instance methods of first `B` then `A` would be examined.

The generic type system of MCJ is similar to the generic type system in Featherweight GJ. However, as in [5, 2], we allow for type variables to occur in type-dependent contexts such as casts, preventing the use of type erasure as an implementation technique. In addition, the receiver of a reference to a class member

may be a type variable. Unlike the system presented in [1], MCJ does not support first-class genericity because a naked type variable must not appear in the `extends` clause or `kind` clause of a class definition and `new` expressions on naked type variables are not supported. Also, polymorphic methods (which are orthogonal to the features we explore) are not supported.

Because classes can be used as expressions, we need a bound on the behavior of classes when used in expression contexts. Therefore, we place two bounds on every type variable `T` in the header of the class definition. The first bound on `T` is a bound on the kind of an instantiation of `T`; the second bound is a bound on the superclass of `T`. For example, the following class header declares a class `C` with one type parameter `T` that must be instantiated with a type of kind `D` that is a subtype of `E`:

```
class C<T extends (D,E)>
  kind Object
  extends Object {...}
```

2.2 Key Design Points

Private Constructors One benefit of a metaclass system is that constructors no longer need to play such a central role in the language. Class instances exist at the beginning of a program execution; they need not be constructed. Non-class instances are naturally constructed with class (factory) methods [10]; in this way a class instance can be passed as a type parameter and factory methods can be called on the class instance to construct new instances without stipulating the class of the constructed instance. This flexibility is similar to that provided with first-class genericity [5, 1], but obviates the need for `with` clauses because the bound on the kind of a type variable stipulates what factory methods can be called on it. A class can define class methods with arbitrary signatures that return new instances of the class, but each such method must ultimately include a `new` expression, which can occur only within the scope of the class whose instance it returns. A `new` expression looks like a call to method named `new`, that takes a single argument for each type parameter and each instance field of the syntactically enclosing class. The result of a `new` expression is a new instance of the enclosing class whose type parameters are instantiated with, and whose fields are initialized with, the given arguments. In our formal semantics, `new` expressions are annotated with the name of the enclosing class. These annotations need not be added by a programmer; they could be added easily by a straightforward syntactic preprocessing over the program, since they can be determined merely by the lexical scopes of the `new` expressions.

All uses of constructors in Featherweight GJ [13] are macro-expressible [9] in MCJ, that is, they can be expressed via local transformations. Of course, not all Featherweight GJ programs can be expressed in MCJ, because of the differences in the type system. Namely, MCJ does not include polymorphic methods.

Metaclasses allow other kinds of flexibility in object creation, as discussed in section 3.

Fields and Initialization To ensure type safety, we must have an initial value for every field, or prevent fields from being used before they are initialized. For instance fields of objects created with a constructor, this is achieved by requiring that the constructor have an argument for every instance field. Combined with the solution for flexibility in object creation outlined above, this allows us the simplicity in semantics of FGJ, and the flexibility of more general constructors.

However, classes which are themselves instances of other classes have fields as well, and these fields must also be initialized before they are used. These fields include those instance fields of the kind which become class fields of the new class. Our solution is to require that *all class fields, including those obtained by inheritance, must be redeclared* and provided with values.

typeOf Given that class references can be used as expressions, it is natural to ask: what is the type of a class reference? In the world of the λ -calculus, such types are known as kinds. In MCJ, however, the kind of a class does not capture all of the properties of the class as a value. For example, a class may add new class fields or class methods which are not present in the kind. Therefore, each class freely generates a new type, which is the type of the class considered as a value. We represent this type with the `typeOf` operator, which produces a type that is not also a class. Since this type is not a class, it cannot be the superclass or kind of another class, and it cannot serve as the instantiation of a generic type parameter.

These new types can create complex relationships between classes and types. For example, the following class headers:

```
class A() kind Object extends Object {...}
class B() kind A extends Object {...}
```

induce the following type relationships:

```
B : typeOf [B]
typeOf [B] <: A
B instanceof A
```

The name `typeOf` has been used in other contexts to refer to a runtime operation that determines the dy-

namic type of an object. However, despite the name similarity, these functions bear no relationship to the operator described here. In MCJ, an application of `typeOf` is a static type reference.

Non-transitivity The standard subclassing relationship, like subtyping, is transitive. That is, if `B` is subclass of `C` and `C` is a subclass of `D`, then `B` is a subclass of `D`. However, this relationship does not hold for `instanceof` relationships. Therefore, if `B` is an instance of `C` and `C` is an instance of `D`, then no judgment about the relationship of `B` to `D` can be inferred.

Cyclic class hierarchies One property of metaclasses that we have not yet discussed is that cycles can occur among instance class relationships. The simplest such cycle is:

```
class C<> : C<> <Object {}
```

Although this pathology appears to be dangerously close to Russell's paradox, it does not lead to inconsistencies. Class `C` is an instance of itself. It can be used in any context where one of its instances can be used. `C` can be thought of as a self-replicating value. This class hierarchy causes no problems for method lookup because method resolution never proceeds through more than a single containing class to an instance class. Therefore, we see no reason to disallow it.

2.3 Examples, Reprised

Having seen the outlines of MCJ, we return to the motivational examples discussed earlier. First, the relationships between dimensions is easy to capture.² We define

```
class Dimension kind Object extends Object { ... }
class Length kind Dimension extends Object { ... }
class Meter kind Length extends Object { ... }
```

Here `Meter` is a singleton class (there is only one `Meter`, in the Platonic sense). We have easily expressed the desired type relationships, and can statically check program invariants that rely on dimensional relationships.

The example from [18] is also easily expressed:

```
class Species kind Object extends Object { ... }
class Eagle kind Species extends Object {
    class Eagle make(String name) { ... }
}
Eagle harry = Eagle.make("Harry")
```

²Of course, dimensions have many subtleties, not captured here. All we present is the essence of the type relationships.

Here we simply create an object to represent our concrete instance of an `Eagle`. The naive implementation of the above code in a conventional OO language would look something like this:

```
class Species extends Object { ... }
class TheEagle extends Species { ... }
class Eagle extends Object { ... }
Eagle harry = new Eagle("Harry")
```

There are at least two substantial problems with this code. First, `TheEagle` and `Eagle` have no relationship to each other in the type system. One is a singleton class, representing a particular species, and one is a name for a set of objects, being the members of that species. The fundamental nature of species has a two-level containment relationship, which the type system is failing to express. This means that our type checker cannot determine what we want, and cannot help us avoid mistakes.

Another problem, which is in some ways just a symptom of the first problem, but which bears special attention, is the use of generic types. Imagine the following method, where `T` is a type parameter in scope:³

```
foo(T x) { ... }
```

If we want to perform the call `foo(harry)`, then the bound on `T` must be either `Object`, in which case `foo` must have no knowledge of the object, or `Eagle`, in which case the function cannot handle multiple `Species`. The “solution” in conventional languages is to use a bound of `Object` and insert a downcast, which fails at runtime if the wrong argument is passed. In MCJ `T` can be bound by kind `Species`, and instantiated with `Eagle`, allowing the original call to be type-checked. This gives the programmer both safety and expressiveness.

3 Design Patterns as Language Features

Design Patterns, as exemplified in [10], have had a very substantial impact on the world of object-oriented programming and beyond. Design patterns allow programmers to increase the flexibility and abstraction of their software designs. However, few design patterns have been integrated into programming languages.

One of the advantages of our metaclass framework is that it allows clean expression of several patterns in the language, rather than adding additional abstraction techniques on top of the language. For example, a number of the classic design patterns deal with object creation, as discussed above in section 2.2. It is clear

³Generic methods would be straightforward to add to MCJ, but we excluded them to simplify the presentation of the type system, and since they do not add complexity to the proofs.

that the concept of a Factory method is easily (and necessarily) expressed in MCJ. However, a number of other patterns also find easy expression. The Abstract Factory pattern is a simple application of inheritance in the metaclass hierarchy. The Prototype pattern is, as mentioned in the Design Patterns book, supported in languages with metaclasses. Finally, there is no need for the Singleton pattern, since the class itself can serve as a singleton. While this is possible in the Java Programming Language and in C++, it is undesirable since the static behavior that would need to be used cannot inherit from another class. In fact, the inflexibility of class or static operations is one of the rationales cited for the Singleton pattern.

Seeing that a number of patterns can be expressed quite simply in MCJ, the question arises: is the simple expression of these patterns a benefit? We argue that it is. The existence of design patterns is a sign of shortcomings in language design. Fundamentally, a design pattern is an abstraction that cannot be expressed in the language. For example, the Singleton and Visitor patterns are abstractions that have obvious invariants, but they cannot be expressed directly and they require complex cooperation from many parts of the system.

Sophisticated macro systems, such as those found in Common Lisp [16] or Scheme [8] allow for expression of many forms of abstraction, as do extremely flexible object systems such as CLOS. However, these solutions add significantly to what a programmer must understand in order to use the language productively.

An alternative solution explored in this paper is to provide additional flexibility in the language, which does not allow for arbitrary additional abstraction, but instead makes the abstractions commonly needed easy to express. One does not need the full power of macros to encode the Factory pattern. In MCJ, we have presented a language where several kinds of patterns are easy to express, without adding an additional language on top of the original.

4 Formal Specification of MCJ

Having outlined the motivation for metaclasses, and the basics of their use, we now turn to a formal exposition of the syntax and semantics of MCJ, followed by an outline of the proof of soundness for the calculus.

4.1 Syntax

The syntax of MCJ is given in Figure 1. When describing the formal semantics of MCJ, we use the following metavariables:

- Expressions or mappings: `e`, `d`, `r`

CL	::=	class C< \bar{X} < (\bar{T} , \bar{T}) > : A < A {class \bar{T} \bar{f} \bar{e} ; \bar{M} ; \bar{T} \bar{f} ; \bar{N} ;} }	class declaration
N	::=	T m (\bar{T} \bar{x}) {return e; }	method declaration
M	::=	class T m (\bar{T} \bar{x}) {return e; }	class method declaration
e	::=	x	variable reference
		e.f	field access
		e.m(\bar{e})	method invocation
		new _C < \bar{R} >(e)	instance creation
		(T)e	cast
		R	type reference
		Φ_G	mapping
T	::=	R	non-typeOf type
		typeOf [R]	typeOf application
R	::=	X	type variable
		A	
A	::=	C< \bar{R} >	type application
		Object	

Figure 1: MCJ Syntax

- Field names: f, g
- Variables: x
- Method names: m
- Values: v
- Method declarations: M, N
- Type names: C, D
- Types: I, J, K, T, U, V, W
- Non-typeOf Types: O, P, Q, R, S
- Types that are either Object or a type application: A, B
- Ground types (contain no type variables): G
- Type variables: X, Y, Z
- Mappings (field name \mapsto value): Φ

As in Featherweight Java and many other works, \bar{x} stands for a possibly-empty sequence of x. $[\bar{X} \mapsto \bar{S}]$ denotes a substitution of the \bar{S} for the \bar{X} , which can be applied to either an expression or a type, and which can substitute either type or expression variables.

A number of symbols are used to abbreviate keywords: < stands for **extends** and : stands for **kind**. Also @ represents concatenation of sequences of syntactic constructs. Finally, $CT(T)$ is a lookup in the class table for the definition of T.

A number of restrictions on MCJ programs are implicit in the formal rules. First, we assume that all

sequences of methods and fields are free of duplicates. Second, there is an implicit well-formedness constraint on programs that no class be a superclass of itself, either directly or indirectly (however, as discussed below, cycles in the kind hierarchy are allowed). Third, we assume that **this** is never used as the name of a variable, method or field. We also take **Object** to be a distinguished member of the hierarchy, with no specific definition, which does not have a superclass or a kind.

For example, the simplest MCJ class is:

```
class C<> : Object <Object{}
```

This is a class named C, with no type arguments, with kind **Object** and superclass **Object**, which has no fields or methods whatsoever. In examples that follow, we omit empty <>. Note that both the kind and the superclass are **Object**. MCJ does not have a distinguished class **Class**. Interestingly, we determined that a class **Class** would need no consideration from the type system. Since both **Class** and **Object** would sit atop the hierarchy with no contents, we determined there was no need to include both of them. In a more substantial language, where **Object** might have methods or fields, there might be a use for a distinguished **Class** class. However, it would not need any special treatment by the type system.

A more complicated MCJ class is:⁴

```
class Pair<A extends (Object, Object),
        B extends (Object, Object)>
  kind Object extends Object
```

⁴Here we use **extends** for < and **kind** for ..

```

{
  class Pair<A,B> make(A a, B b){ new(A,B) (a,b) }
  A fst;
  B snd;
  Pair<B,B> setfst(B b){
    return new <B,B> (b, this.snd);
  };
}

```

This class specifies a polymorphic pair data structure. It also demonstrates a number of important MCJ features. First, we see two different kinds of methods. The `make` method creates a new instance of `Pair`. This method contains a `new` expression, which initializes the instance fields of the class positionally, so that the first argument to `make` becomes `fst`, and the second becomes `snd`. Also, note that we must mention the type parameters of the newly created instance.

Second, we have an instance method, `setfst`, which creates a new pair with the old second element, and a new first element which has the same type as the second element.

Finally, adding the following to the above definition of `Pair` gives a full MCJ program:

```

class 0 kind Object extends Object
{
  class 0 make(){return new () ()};
}

Pair<0,0>.make(0.make(), 0.make()).fst

```

This program, which creates two instances of `0`, then creates a `Pair` to hold them both, and finally selects one of the two to become the value of the program, demonstrates three of the expression forms in MCJ. `Pair<0,0>` is a type, which, by itself, can be a value. `0.make()` is a method invocation (here with a class, not an instance, as the receiver). The entire expression is a field access, of the `fst` instance field.

Other kinds of expressions are seen in the body of the `Pair` class. `new<A,B> (a,b)` is a new expression, which creates an instance of the enclosing class (here `Pair`). In the body of `make` is a reference to the variable `b`, which has the obvious value. The final directly expressible expression is the cast, which works as follows: `(Object)0.make()` is an expression of type `Object` that evaluates to an instance of `0`.

With an understanding of expressions, we can examine the rest of the `Pair` class. There are two methods: a class method (`make`) and an instance method (`setfst`). There are also two fields, both of which are instance fields and initialized the `new` expressions. In MCJ, a `new` expression is different from those in Featherweight GJ;

each `new` expression evaluates to an instance of the syntactically enclosing class.⁵ Had there been additional fields in the superclass of `Pair`, it would have been necessary to initialize them in the `new` expression as well. `new` expressions are explained in detail in section 2.2.

4.2 Semantics

There are two forms that a value can take in MCJ: that of an instance class, and that of a conventional (non-class) value. If we were to distinguish these two forms of value, we would significantly increase the number of rules necessary to describe our semantics because every rule referring to a value would have to be written twice. To avoid this complexity, we introduce a special *mapping construct* (similar to a record value) to denote the results of computations.⁶ A mapping takes field names to expressions, and is annotated with a ground type. Mappings from a sequence of fields \bar{f} to a sequence of expressions \bar{e} with type G are written $\{\bar{f} \mapsto \bar{e}\}_G$. A mapping denoting an instance class C consists of a sequence of class fields mapped to a sequence of expressions and is annotated with the type `typeOf(C)`. Note that the right hand sides of such maps need not always be values, and thus computation can take place inside of a mappings. Mappings are not available to the programmer, and thus can only be created by operation of the reduction rules. Therefore, unlike Featherweight Java, our reductions do not operate entirely in the expressible syntax of the language. We use the metavariable Φ to range over mappings. When we need to refer to the type annotation of a mapping explicitly, the metavariable is written Φ_G .

The semantics are given in figures 2, 3, 4, and 5 with auxiliary functions given in figure 6.

4.2.1 Typing

Rules governing the typing of MCJ programs are given in figures 2, 3 and 4. The $bound_\Delta$ function is defined as follows:

$$\begin{aligned}
bound_\Delta(X) &= \Delta(X) \\
bound_\Delta(\text{typeOf}[X]) &= \Delta(\text{typeOf}[X]) \\
bound_\Delta(S) &= S
\end{aligned}$$

This function maps type variables and `typeOf` applied to type variables to their bounds, and leaves others unchanged.

The metavariables Δ and Γ range over bounds environments, written $\bar{X} \triangleleft \bar{S}$ and type environments, written

⁵In the grammar of figure 1, `new` expressions are annotated with this enclosing class. In the examples, this redundant information is elided.

⁶This simplification was suggested by Jan-Willem Maessen in response to an earlier draft of the MCJ semantics.

$\bar{x} : \bar{T}$ respectively. A bounds environment contains two bounds for each type variable, so that if $\Delta = X <: (A, B)$, then $\Delta(X) = A$ and $\Delta(\text{typeOf}[X]) = B$.

The notation $A <: B$ means A is a subtype of B . Subtyping judgments are made in the context of a bounds environment, which relates a type to the declared bound of that type from the class header.

Type judgments are of the form $\Delta; \Gamma \vdash e : T$, which states that in bounds environment Δ and type environment Γ , expression e has type T . Type judgments are not transitive in the way that `instanceof` relationships are with respect to subtyping: if $\Delta; \Gamma \vdash e : T$, and $\Delta \vdash T <: S$, it is not necessarily the case that $\Delta; \Gamma \vdash e : S$. This is important, since our proofs depend on having unique derivations of a given typing judgment. We represent empty environments with \emptyset and we abbreviate judgments of the form $\emptyset; \emptyset \vdash e : T$ and $\emptyset \vdash T <: S$ as $e : T$ and $T <: S$ respectively.

We now discuss several non-obvious aspects of our type system resulting from the need to statically check uses of metaclasses.

Casts and Mappings We follow [1] in typing all casts as statically correct, so as to avoid the complications of “stupid casts”. Mappings, which are not expressible in the source language, are typed merely by a well-formedness constraint.

4.2.2 Well-Formedness

There are three kinds of well-formedness constraints on MCJ programs. First, type well-formedness, written “ $\Delta \vdash T$ ok”, states that type T refers to a defined class, and that if it is a type application, the arguments satisfy the bounds.

More important are the class and method well-formedness constraints, which together determine if a class table is well-formed, and thus part of a legal MCJ program. Method well-formedness is a judgment of the form “ M ok in G ”, where M is a method. We make use of the latter portion (G) of this judgment as a second argument, allowing us to use this rule to check both instance and class methods. Method well-formedness consists mostly in fitting the body to the return type, and checking for invalid overrides.

Class well-formedness involves all of the above checks. All methods must be well-formed, and all class fields must have correct initialization expressions. Further, class fields must contain the instance fields of the kind. Finally, all `new` expressions must have the correct class name annotation.

4.2.3 Evaluation

The evaluation rules for MCJ are given in figure 5. The evaluation relation, written $e \rightarrow e'$, states that e transitions to e' in one step. The reflexive transitive closure of this relation is written $e \rightarrow^* e'$.

Bad Casts There is one way that evaluation in MCJ can get stuck: the bad cast. This problem occurs for all languages that allow static downcasts. A expression e is a bad cast iff $e = (S)\Phi_T$ where S is not a supertype of T .

The evaluation relation given here is non-deterministic. For example, no order is prescribed for evaluating the arguments to a `new` expression or method call, or for evaluating the initializers for class fields. There are also a number of places where either congruence or reduction rules can be applied. Therefore, in the presence of non-termination or bad casts, the results may differ depending on evaluation order. However, confluence can be regained simply by requiring that [C-MAPPING] is applied whenever possible. This restriction, combined with the requirement in the premise of [R-NEW] that the arguments be values, ensures that every program with an error or non-termination will either cause an error, or fail to terminate.

4.3 Type Soundness

With the above definitions, we are now able to turn to a proof of type soundness. Given the simplicity we are able to achieve in the definitions, the proof is not significantly more complicated than the soundness proof given for Featherweight GJ [13]. However, there are a number of significant lemmas, of which we list the most important:

Lemma 1 (Fields are Preserved by Subtypes)

If $fields(U) = \bar{F}$ and $\emptyset \vdash V <: U$ then $\exists \bar{G}$ where $fields(V) = \bar{F} @ \bar{G}$.

Proof We prove this by induction over the derivation for $fields(V)$.

Case $V = \text{Object}$: Then $U = \text{Object}$ and $\emptyset = \emptyset @ \emptyset$.

Case $V = C\langle\bar{R}\rangle$: Then
$$CT(C) = \text{class } C\langle\bar{X}\langle(\bar{I}, \bar{J})\rangle : B\langle A\{...\bar{H}...\}\rangle$$

Continue by induction on the derivation of $\emptyset \vdash V <: U$. The only interesting cases are [S-SUPER] and [S-TRANS].

Subcase [S-Super]: $[\bar{X} \mapsto \bar{R}]A = U$. Thus $\bar{G} = \bar{H}$.

$\Delta \vdash T <: T$ [S-REFLEX]	$\frac{X \triangleleft T \in \Delta}{\Delta \vdash X <: T}$ [S-BOUND]	$\frac{\Delta \vdash V <: T \quad \Delta \vdash T <: U}{\Delta \vdash V <: U}$ [S-TRANS]
$\frac{CT(C) = \text{class } C \triangleleft \bar{X} \triangleleft (\bar{I}, \bar{J}) > : B \triangleleft A \{ \dots \}}{\Delta \vdash \text{typeOf}[C \triangleleft \bar{S}] <: [\bar{X} \mapsto \bar{S}]B}$ [S-KIND]	$\frac{CT(C) = \text{class } C \triangleleft \bar{X} \triangleleft (\bar{I}, \bar{J}) > : B \triangleleft A \{ \dots \}}{\Delta \vdash C \triangleleft \bar{S} <: [\bar{X} \mapsto \bar{S}]A}$ [S-SUPER]	
$\Delta \vdash \text{Object ok}$ [WF-OBJECT]	$\frac{\Delta \vdash X <: T}{\Delta \vdash X \text{ ok}}$ [WF-VAR]	$\frac{\Delta \vdash T \text{ ok}}{\Delta \vdash \text{typeOf}[T] \text{ ok}}$ [WF-TYPEOF]
$\frac{CT(C) = \text{class } C \triangleleft \bar{X} \triangleleft (\bar{I}, \bar{J}) > : B \triangleleft A \{ \dots \} \quad \Delta \vdash \bar{S} <: [\bar{X} \mapsto \bar{S}]\bar{I} \quad \Delta \vdash \text{typeOf}[\bar{S}] <: [\bar{X} \mapsto \bar{S}]\bar{J} \quad \Delta \vdash \bar{S} \text{ ok}}{\Delta \vdash C \triangleleft \bar{S} \text{ ok}}$ [WF-CLASS]		

Figure 2: Subtyping and Well-Formed Types

$\Delta; \Gamma \vdash x : \Gamma(x)$ [T-VAR]	$\frac{\Delta \vdash T \text{ ok} \quad \Delta; \Gamma \vdash e : U}{\Delta; \Gamma \vdash (T)e : T}$ [T-CAST]	$\frac{\Delta \vdash S \text{ ok}}{\Delta; \Gamma \vdash S : \text{typeOf}[S]}$ [T-CLASS]
$\frac{\Delta; \Gamma \vdash e : U \quad \text{fields}(\text{bound}_\Delta(U)) = \bar{T} \bar{f}}{\Delta; \Gamma \vdash e.f_i : T_i}$ [T-FIELD]	$\frac{\Delta \vdash C \triangleleft \bar{S} \text{ ok} \quad \text{fields}(C \triangleleft \bar{S}) = \bar{T} \bar{f}}{\Delta \vdash \bar{U} <: \bar{T} \quad \Delta; \Gamma \vdash \bar{e} : \bar{U}}$ [T-NEW]	
$\frac{\text{fields}(G) = \bar{T} \bar{f} \quad \Delta; \Gamma \vdash \bar{e} : \bar{S} \quad \Delta \vdash \bar{S} <: \bar{T}}{\Delta; \Gamma \vdash \{\bar{f} \mapsto \bar{e}\}_G : G}$ [T-MAPPING]	$\frac{\text{mtype}(m, \text{bound}_\Delta(W)) = \bar{U} \rightarrow T \quad \Delta; \Gamma \vdash r : W \quad \Delta; \Gamma \vdash \bar{e} : \bar{V} \quad \Delta \vdash \bar{V} <: \bar{U}}{\Delta; \Gamma \vdash r.m(\bar{e}) : T}$ [T-INVK]	

Figure 3: Expression Typing

$\begin{aligned} & \text{params}(G) = \bar{X} \sqsubset (\bar{I}, \bar{J}) \\ & \Delta = \{ \bar{X} \triangleleft \bar{I}, \text{typeOf}[\bar{X}] \triangleleft \bar{J} \} \quad \Gamma = \{ \bar{x} : \bar{T}, \text{this} : G \} \\ & \Delta \vdash \bar{T} \text{ ok} \quad \Delta \vdash U \text{ ok} \quad \text{override}(m, \text{super}(G), \bar{T} \rightarrow U) \\ & \Delta; \Gamma \vdash e : W \quad \Delta \vdash W <: U \end{aligned}$	[WF-METHOD]
$\frac{\begin{aligned} & U m (\bar{T} \bar{x}) \{ \text{return } e; \} \text{ ok in } G \\ & \bar{M} \text{ ok in } \text{typeOf}[C \triangleleft \bar{X}] \quad \bar{N} \text{ ok in } C \triangleleft \bar{X} \quad \Delta = \{ \bar{X} \triangleleft \bar{I}, \text{typeOf}[\bar{X}] \triangleleft \bar{J} \} \\ & \Delta \vdash B \text{ ok} \quad \Delta \vdash A \text{ ok} \quad \Delta \vdash I \text{ ok} \quad \Delta \vdash \bar{J} \text{ ok} \quad \Delta \vdash \bar{T} \text{ ok} \quad \Delta \vdash \bar{W} \text{ ok} \\ & \text{fields}(B) \subseteq \bar{T} \bar{f} \quad \Delta; \emptyset \vdash \bar{e} : \bar{T} \quad \emptyset \vdash \bar{T}' <: \bar{T} \\ & \forall \text{new}_D \triangleleft \bar{S} > (\bar{d}) \in \bar{e}, \bar{M}, \bar{N}. D = C \end{aligned}}{\text{class } C \triangleleft \bar{X} \triangleleft (\bar{I}, \bar{J}) > : B \triangleleft A \{ \bar{T} \bar{f} \bar{e}; \bar{M}; \bar{W} \bar{g}; \bar{N}; \} \text{ ok}}$	
[WF-CLASSDEF]	

Figure 4: Well-Formed Constructs

Subcase [S-Trans]: Let T be the intermediate type. Then $\text{fields}(V) = \text{fields}(T) @ \bar{H}'$ and $\text{fields}(V) = \text{fields}(T) @ \bar{H}''$ by the induction hypothesis. Thus $\bar{G} = \bar{H}' @ \bar{H}''$.

Case $V = \text{typeOf}[C \triangleleft \bar{S}]$: Then $CT(C) = \text{class } C \triangleleft \bar{X} \triangleleft (\bar{I}, \bar{J}) > : B \triangleleft A \{ \bar{w} \bar{h} \bar{e} \dots \}$. We continue by induction over the derivation of $\emptyset \vdash V <: U$. The only complex cases are [S-KIND] and [S-TRANS].

Subcase [S-Trans]: As above.

Subcase [S-Kind]: $[\bar{X} \mapsto \bar{S}]A = U$ Then $\text{fields}(V) = [\bar{X} \mapsto \bar{S}]\bar{w} \bar{h}$. By well-formedness of C , we know that

$\text{fields}(B) \subseteq \bar{w} \bar{h}$. Then $[\bar{X} \mapsto \bar{S}]\text{fields}(B) \subseteq [\bar{X} \mapsto \bar{S}]\bar{w} \bar{h}$, and $\text{fields}([\bar{X} \mapsto \bar{S}]B) \subseteq [\bar{X} \mapsto \bar{S}]\bar{w} \bar{h}$ by lemma Substitution Distributes over fields .

Lemma 2 (Subtyping Preserves Method Typing)

If $\text{mtype}(m, U) = \bar{T} \rightarrow S$ and $\emptyset \vdash V <: U$ then $\text{mtype}(m, V) = \bar{T} \rightarrow S$.

Proof By induction over the derivation of $\emptyset \vdash V <: U$.

Case [S-Reflex]: Trivial.

Case [S-Bound]: Not applicable.

$\text{Object} \rightarrow \{\}\text{Object}$ [R-OBJECT]	$\frac{\text{fields}(\mathbf{C}\langle\bar{\mathbf{S}}\rangle) = \bar{\mathbf{T}} \bar{\mathbf{f}}}{\text{new}_{\mathbf{C}\langle\bar{\mathbf{S}}\rangle}(\bar{\mathbf{v}}) \rightarrow \{\bar{\mathbf{f}} \mapsto \bar{\mathbf{v}}\}_{\mathbf{C}\langle\bar{\mathbf{S}}\rangle}}$ [R-NEW]
$\frac{\emptyset \vdash \mathbf{G} <: \mathbf{T}}{(\mathbf{T})\Phi_{\mathbf{G}} \rightarrow \Phi_{\mathbf{G}}}$ [R-CAST]	$\frac{\text{field-vals}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{S}}\rangle]) = \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}}{\mathbf{C}\langle\bar{\mathbf{S}}\rangle \rightarrow \{\bar{\mathbf{f}} \mapsto \bar{\mathbf{e}}\}_{\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{S}}\rangle]}}$ [R-CLASS]
$\Phi_{\mathbf{G}}.f \rightarrow \Phi_{\mathbf{G}}(f)$ [R-FIELD]	$\frac{\text{mbody}(\mathbf{m}, \mathbf{G}) = (\bar{\mathbf{x}}, \bar{\mathbf{e}}_0)}{\Phi_{\mathbf{G}}.m(\bar{\mathbf{d}}) \rightarrow [\bar{\mathbf{x}} \mapsto \bar{\mathbf{d}}, \text{this} \mapsto \Phi_{\mathbf{G}}]\bar{\mathbf{e}}_0}$ [R-INVK]
$\frac{e \rightarrow e'}{e.f \rightarrow e'.f}$ [C-FIELD]	$\frac{e \rightarrow e'}{e.m(\bar{\mathbf{d}}) \rightarrow e'.m(\bar{\mathbf{d}})}$ [C-RCVR]
$\frac{\bar{e} \rightarrow \bar{e}'}{d.m(\bar{e}) \rightarrow d.m(\bar{e}')}$ [C-ARG]	$\frac{e \rightarrow e'}{(\mathbf{T})e \rightarrow (\mathbf{T})e'}$ [C-CAST]
	$\frac{\bar{e} \rightarrow \bar{e}'}{\text{new}_{\mathbf{C}\langle\bar{\mathbf{S}}\rangle}(\bar{e}) \rightarrow \text{new}_{\mathbf{C}\langle\bar{\mathbf{S}}\rangle}(\bar{e}')}$ [C-NEW]
	$\frac{\bar{e} \rightarrow \bar{e}'}{\{\bar{\mathbf{f}} \mapsto \bar{\mathbf{e}}\}_{\mathbf{G}} \rightarrow \{\bar{\mathbf{f}} \mapsto \bar{\mathbf{e}}'\}_{\mathbf{G}}}$ [C-MAPPING]

Figure 5: Computation Rules

$\text{fields}(\text{Object}) = \emptyset$ [F-OBJECT]	$\text{field-vals}(\text{Object}) = \emptyset$ [FV-OBJECT]	$\text{field-vals}(\mathbf{C}\langle\bar{\mathbf{S}}\rangle) = \emptyset$ [FV-CLASS]
$CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}; \bar{\mathbf{M}}; \bar{\mathbf{W}} \bar{\mathbf{g}}; \bar{\mathbf{N}}; \}$ $\text{fields}([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{A}) = \bar{\mathbf{U}} \bar{\mathbf{h}}$		
[F-CLASS]		
$\text{fields}(\mathbf{C}\langle\bar{\mathbf{R}}\rangle) = \bar{\mathbf{U}} \bar{\mathbf{h}} \cup [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{W}} \bar{\mathbf{g}}$ $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}; \bar{\mathbf{M}}; \bar{\mathbf{W}} \bar{\mathbf{g}}; \bar{\mathbf{N}}; \}$		
[F-TYPEOF]		
$\text{fields}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{R}}\rangle]) = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{T}} \bar{\mathbf{f}}$ $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}; \bar{\mathbf{M}}; \bar{\mathbf{W}} \bar{\mathbf{g}}; \bar{\mathbf{N}}; \}$		
[FV-TYPEOF]		
$\text{field-vals}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{R}}\rangle]) = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}$		
$\frac{\text{U m } (\bar{\mathbf{T}} \bar{\mathbf{x}}) \{ \text{return } e; \} \in \text{methods}(\mathbf{G})}{\text{mtype}(\mathbf{m}, \mathbf{G}) = \bar{\mathbf{T}} \rightarrow \bar{\mathbf{U}}}$ [MTYPE]	$\frac{\text{U m } (\bar{\mathbf{T}} \bar{\mathbf{x}}) \{ \text{return } e; \} \in \text{methods}(\mathbf{G})}{\text{mbody}(\mathbf{m}, \mathbf{G}) = (\bar{\mathbf{x}}, \bar{\mathbf{e}})}$ [MBODY]	
$\text{methods}(\text{Object}) = \emptyset$ [METHODSOBJECT]		
$CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}; \bar{\mathbf{M}}; \bar{\mathbf{W}} \bar{\mathbf{g}}; \bar{\mathbf{N}}; \}$		
[METHODSCONST]		
$\text{methods}(\mathbf{C}\langle\bar{\mathbf{R}}\rangle) = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{N}} \cup \text{methods}([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{A})$ $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \bar{\mathbf{T}} \bar{\mathbf{f}} \bar{\mathbf{e}}; \bar{\mathbf{M}}; \bar{\mathbf{W}} \bar{\mathbf{g}}; \bar{\mathbf{N}}; \}$		
[METHODSTYPEOF]		
$\text{methods}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{R}}\rangle]) = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{M}} \cup \text{methods}([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{B})$		
[METHODSOF]		
$\text{mtype}(\mathbf{m}, \mathbf{G}) = \bar{\mathbf{U}} \rightarrow \bar{\mathbf{V}}$ implies $\bar{\mathbf{W}} = \bar{\mathbf{U}}$ and $\bar{\mathbf{T}} = \bar{\mathbf{V}}$ [OVERRIDE]		
$\text{override}(\mathbf{m}, \mathbf{G}, \bar{\mathbf{W}} \rightarrow \bar{\mathbf{T}})$		
$CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \dots \}$ [PARAMS]	$CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \dots \}$ [PARAMSTYPEOF]	
$\text{params}(\mathbf{C}\langle\bar{\mathbf{T}}\rangle) = \bar{\mathbf{X}} \sqsubset (\bar{\mathbf{I}}, \bar{\mathbf{J}})$ $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \dots \}$ [SUPER]	$\text{params}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{T}}\rangle]) = \bar{\mathbf{X}} \sqsubset (\bar{\mathbf{I}}, \bar{\mathbf{J}})$ $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \dots \}$ [SUPERTYPEOF]	
$\text{super}(\mathbf{C}\langle\bar{\mathbf{T}}\rangle) = \mathbf{A}$	$\text{super}(\text{typeOf}[\mathbf{C}\langle\bar{\mathbf{T}}\rangle]) = \mathbf{B}$	

Figure 6: Auxiliary Functions

Case [S-Super]: In this case, $\mathbf{V} = \mathbf{C}\langle\bar{\mathbf{R}}\rangle$, where $CT(\mathbf{C}) = \text{class } \mathbf{C}\langle\bar{\mathbf{X}}\rangle \triangleleft (\bar{\mathbf{I}}, \bar{\mathbf{J}}) > : \mathbf{B} \triangleleft \mathbf{A} \{ \dots \bar{\mathbf{M}} \}$ and $\mathbf{U} = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{A}$. We need to show that $\text{mtype}(\mathbf{m}, \mathbf{C}\langle\bar{\mathbf{R}}\rangle) = \bar{\mathbf{T}} \rightarrow \bar{\mathbf{S}}$. By lemma *methods* is well-defined, it suffices to show that $\text{U m } (\bar{\mathbf{T}} \bar{\mathbf{x}}) \{ \text{return } e; \} \in \text{methods}(\mathbf{C}\langle\bar{\mathbf{R}}\rangle)$.

We know that $\text{U m } (\bar{\mathbf{T}} \bar{\mathbf{x}}) \{ \text{return } e; \} \in \text{methods}(\mathbf{U})$. So, we proceed by induction on the derivation of $\text{methods}(\mathbf{U})$.

Subcase $\mathbf{U} = \text{Object}$: Impossible, since Object has no methods.

Subcase $\mathbf{U} = \mathbf{D}\langle\bar{\mathbf{W}}\rangle$: Then, $\mathbf{U} = [\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{A}$. By [METHODSCONST], $\text{methods}(\mathbf{V}) = \text{methods}(\mathbf{C}\langle\bar{\mathbf{R}}\rangle) = ([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{M}}) \cup \text{methods}([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \mathbf{A}) = ([\bar{\mathbf{X}} \mapsto \bar{\mathbf{R}}] \bar{\mathbf{M}}) \cup \text{methods}(\mathbf{U})$. Therefore, any method in $\text{methods}(\mathbf{U})$ must be in $\text{methods}(\mathbf{V})$.

Case [S-Typeof]: Analogous to [S-SUPER].

Case [S-Trans]: Let the intermediate type be \mathbf{T} . Then $\emptyset \vdash \mathbf{V} <: \mathbf{T}$ and $\emptyset \vdash \mathbf{T} <: \mathbf{U}$. Thus, by the induction hypothesis, $\text{mtype}(\mathbf{m}, \mathbf{V}) = \text{mtype}(\mathbf{m}, \mathbf{T}) = \text{mtype}(\mathbf{m}, \mathbf{U})$.

Lemma 3 (Substitution Preserves Typing)

If $\Delta; \Gamma \vdash e : T$ and $\Gamma = \bar{x} : \bar{S}$ and $\Delta; \emptyset \vdash \bar{d} : \bar{U}$ and $\Delta \vdash \bar{U} <: \bar{S}$ then $\Delta; \emptyset \vdash [\bar{x} \mapsto \bar{d}]e : T'$ where $\Delta \vdash T' <: T$.

Proof With the above two lemmas, this one follows by straightforward structural induction over the derivation of $\Delta; \Gamma \vdash e : T$.

Given the above lemmas, the following subject reduction proof is a simple structural induction with a case analysis on the typing rule used to derive $\emptyset; \emptyset \vdash e : S$.

Theorem 1 (Subject Reduction) *If $\emptyset; \emptyset \vdash e : S$ and $e \rightarrow e'$ then $\emptyset; \emptyset \vdash e' : T$ where $\emptyset \vdash T <: S$.*

Proof We prove this by structural induction on the derivation of $e \rightarrow e'$.

Case [R-Object]: Immediate.

Case [R-New]: Immediate from the premises of [T-NEW] and [R-NEW].

Case [R-Class]: Immediate from the premises of [WF-CLASSDEF] and [R-CLASS].

Case [R-Cast]: By [T-CAST], $e = (T)\Phi_G$ must have type T . By [T-MAPPING], $e' = \Phi_G$ must have type G . By hypothesis of the reduction rule, $\emptyset \vdash G <: T$.

Case [R-Field]: We know that $e = \{\bar{f} \mapsto \bar{d}\}_G.f_i$ and that $e' = d_i$. Since e must have been typed by [T-FIELD], we know that $fields(G) = \bar{T} \bar{f}$ and $\emptyset; \emptyset \vdash e : T_i$. Further, since $\{\bar{f} \mapsto \bar{d}\}_G$ must have been typed by [T-MAPPING], we know that $\emptyset; \emptyset \vdash d_i : S$ where $\emptyset \vdash S <: T$.

Case [R-Invk]: We know that $e = \Phi_G.m(\bar{d})$ and that $e' = [\bar{x} \mapsto \bar{d}, \text{this} \mapsto \Phi_G]e_0$. Further, e was typed by [T-INVK] to have type U where $mtype(m, G) = \bar{T} \rightarrow U$ and $\emptyset; \emptyset \vdash \bar{d} : \bar{T}'$ and $\emptyset \vdash \bar{T}' <: \bar{T}$. As a premise of [R-INVK], we know that $mbody(m, G) = (\bar{x}, e'_0)$.

By the lemma *mtype* and *mbody* agree, $\emptyset; \bar{x} : \bar{T}, \text{this} : G \vdash e'_0 : U'$ where $\emptyset \vdash U' <: U$. Then by the lemma Substitution Preserves Typing, $\emptyset; \emptyset \vdash e' = [\bar{x} \mapsto \bar{d}, \text{this} \mapsto \Phi_G]e_0 : U''$ where $\emptyset \vdash U'' <: U'$.

Case [C-Cast]: Trivial, since $\emptyset \vdash (T)e : T$ for any e .

Case [C-Map]: Immediate from the induction hypothesis and the transitivity of subtyping.

Case [C-New]: Immediate from the induction hypothesis and the transitivity of subtyping.

Case [C-Arg]: Immediate from the induction hypothesis and the transitivity of subtyping.

Case [C-Rcvr]: We know that $e = e_0.m(\bar{d})$ and $e' = e'_0.m(\bar{d})$. Further, e must have been typed by [T-INVK], which means that $\emptyset; \emptyset \vdash e_0 : W$ for some ground type W , and that $\emptyset; \emptyset \vdash \bar{d} : \bar{V}$ and $\emptyset \vdash \bar{V} <: \bar{U}$, and also $mtype(m, W) = \bar{U} \rightarrow T$. By the induction hypothesis, $\emptyset; \emptyset \vdash e'_0 : W'$ where $\emptyset \vdash W' <: W$. Therefore, by lemma Subtyping Preserves Method Typing, $mtype(m, W') = \bar{U} \rightarrow T$ and thus $\emptyset; \emptyset \vdash e'_0.m(\bar{d}) : T$ by [T-INVK].

Case [C-Field]: If $e.f_i \rightarrow e'.f_i$, then $e \rightarrow e'$. Further, $e.f_i$ must have been typed by [T-FIELD] to have type T_i . Therefore, by the induction hypothesis, $\emptyset; \emptyset \vdash e : S$ and $\emptyset; \emptyset \vdash e' : S'$ where $\emptyset \vdash S <: S'$. Then, by the lemma Fields are Preserved by Subtypes, $fields(S') = fields(S) @ \bar{F}$ for some \bar{F} , and by [T-FIELD], $\emptyset; \emptyset \vdash e'.f_i : T_i$.

The proof of progress presents no additional complications, and requires only one new lemma.

Lemma 4 (Agreement of fields and field-vals) *If $fields(G) = \bar{T} \bar{f}$ and $field-vals(G) = \bar{T}' \bar{f}' \bar{e}$ then $\bar{T} = \bar{T}'$, $\bar{f} = \bar{f}'$ and $\emptyset; \emptyset \vdash \bar{e} : \bar{S}$ where $\emptyset \vdash \bar{S} <: \bar{T}$.*

Theorem 2 (Progress) *If $\emptyset; \emptyset \vdash e : S$ then one of the following holds:*

- $e = \{\bar{f} \mapsto \bar{v}\}_G$
- $e \rightarrow e'$
- $e = (S)e'$ and $\emptyset; \emptyset \vdash e' : T$ and $\emptyset \vdash T \not<: S$.

Proof By induction over the derivation of $\emptyset; \emptyset \vdash e : S$.

Case [T-Var]: This is a contradiction, since e is ground.

Case [T-Class]: In this case $e = C\langle \bar{S}' \rangle$ where $C\langle \bar{S}' \rangle$ is ground. Then by the lemma Agreement of field-vals and fields, $params(C\langle \bar{S}' \rangle) = \bar{X} \sqsubset (\bar{I}, \bar{J})$ for some Further, $field-vals(\text{typeOf}[C\langle \bar{S}' \rangle]) = \bar{T} \bar{f} \bar{e}$ for some \bar{T}, \bar{f} , and \bar{e} . Therefore, [R-CLASS] applies and $S \rightarrow \{\bar{f} \mapsto [\bar{x} \mapsto \bar{S}']C\langle \bar{S}' \rangle\}_{\text{typeOf}[C\langle \bar{S}' \rangle]}$.

Case [T-Mapping]: Either e is already a value, or $e = \{\bar{f} \mapsto \bar{e}\}_G$ where not all of the \bar{e} are values. Then by the induction hypothesis, there is some i such that either $e_i \rightarrow e'_i$, in which case [C-MAPPING] applies, or e_i contains a bad cast, and the case is complete.

Case [T-Cast]: Here there are three cases:

- $e = (S)e'$ where e' is not a mapping. Then [C-CAST] applies.

- $e = (S)\Phi_T$ where $\emptyset \vdash T <: S$. Then [R-CAST] applies.
- $e = (S)\Phi_T$ where $\vdash T \not<: S$. Then e is a bad cast.

Case [T-New]: From the antecedent of [T-NEW] the premise of [R-NEW] applies.

Case [T-Invk]: Here there are two cases.

- $e = r.m(\bar{d})$ where r is not a mapping. Then by the induction hypothesis, either r contains a bad cast or $r \rightarrow r'$ and [C-RCVR] applies.
- $e = \Phi_T.m(\bar{d})$ We know from the antecedent that $mtype(m, boundT) = \bar{U} \rightarrow V$ and therefore $mtype(m, T) = \bar{U} \rightarrow V$ since T is ground. Therefore, since $mbody$ is defined everywhere $mtype$ is defined, $mbody(m, T) = (\bar{x}, e_0)$ for some \bar{x} and e_0 . Thus [R-INVK] applies.

Case [T-Field]: Here there are two cases, either the receiver is a mapping or not. In the first, by the antecedent of the typing rule, we can lookup the field successfully and apply [R-FIELD]. Otherwise, we can apply [C-FIELD].

From these, we can conclude the desired type safety result.

Theorem 3 (Soundness) *If $e : S$ then either*

- $e \rightarrow^* \{\bar{f} \mapsto \bar{v}\}_G$
- $e \rightarrow^* e'$ where e' is an invalid cast
- e reduces infinitely.

Proof Immediate from Subject Reduction and Progress.

5 Related Work

A number of object-oriented languages have included some form of metaclass system. Most notable among these is Smalltalk [11], but others include Common Lisp with CLOS [14].

All of these systems share a common architecture of the metaclass system in which each class has its own freely generated metaclass, defined by the class methods and fields of the class. In contrast, MCJ provides a hierarchy for structuring metaclass relationships, which provides significantly more modeling and abstraction flexibility.

The metaclass system present in Python [17] is quite similar to that provided here, with inheritance of instance methods as class methods from arbitrary other classes as metaclasses. However, Python is

dynamically-typed, and many of the uses to which Python metaclasses are put are not possible in a statically typed language. The work on static type systems for Python has not included metaclasses [15].

Several type systems have also been proposed for languages with metaclasses, including the Strongtalk language [3], which turns Smalltalk into a structurally-typed language with extensive static checking. However, because the Smalltalk metaclass system is so different from the one in MCJ, many of the interesting aspects of the type system do not carry over. Furthermore, the Strongtalk papers do not provide a formal semantics and analysis of the system. A formal analysis of inheritance in Smalltalk is provided in [7] but this again does not consider the hierarchy of metaclasses presented here.

Graver and Johnson [12] present another type system for Smalltalk, with a formalism and sketch of a safety proof. Again, the metaclass type system is substantially different, reflecting the underlying Smalltalk system. Additionally, the paper is concerned primarily with optimization as opposed to static checking.

Cointe [6] presents a model of metaclasses extremely similar to that presented here. In it, he provides several overlapping motivations to our own. One is to regularize the metaclass system of Smalltalk, and another is to enable additional programming flexibility. To this we add modeling freedom, and a relation to static methods in more recent OO languages. Cointe’s work, however, does not provide a formal model, so it is difficult to determine the exact relationship between the systems. Additionally, his work is in a untyped setting (Lisp and Smalltalk) and thus the safety theorems proved here are not possible.

When viewed as “instance generators”, our metaclasses are similar to prototypes in untyped languages such as Self. Prototypes generate new instances, which themselves can generate new instances. Our language is more restrictive than prototype-based languages in the sense that all metaclasses and instance classes must be declared statically (i.e., by writing down class definitions). But our language is more expressive in the sense that we include classes and subclassing relationships. Also, unlike typical prototype-based languages, our language is statically typed.

Formalized calculi for object-oriented languages are abundant in the programming languages community today, including Featherweight Java [13], upon which MCJ is based. However, none of them have considered static methods, which are the closest analogue of the metaclass functionality in MCJ.

Finally, the motivation for this work comes from the language MetaGen, introduced in [2]. MetaGen can be seen as an extension of MCJ, which provides numerous

other advanced type features. Here we restrict ourselves to metaclasses and analyze the properties of the system formally.

6 Conclusions and Future Work

With MCJ, we have devised a core calculus for metaclasses that is more flexible than that available in more traditional metaclass systems such as Smalltalk and that allows clean expression of many common design patterns. In doing so, we have demonstrated that metaclasses can be added to a nominally-typed, statically-checked language without either significant complication of the semantics or difficulty in the proof of soundness.

In addition to this contribution, we have elucidated several other important points about the integration of metaclasses into an object-oriented system. The `typeof` type operator is to our knowledge unique, and plays a key role in the soundness of the system. The discovery that the `Class` class played no special role, and that thus the class and metaclass hierarchies could both be rooted at `Object` is also novel. Finally, we have shown how an expressive framework of metaclasses has positive effects in other areas of the language, such as mechanisms for object construction.

A natural extension of the work in this paper is to expand MCJ to include more of the features presented in [2], so as to allow inclusion of the system for checking dimensions of physical quantities. Such an extension would allow for a proof of “dimensional soundness” in the resulting system. Further, any realistic system will have imperative features, and while such features do not seem likely to interact badly with metaclasses, no sound system is built on such assumptions.

Another interesting extension would be to expand our calculus to include either multiple inheritance (as does Python) or some alternative such as mixins or traits, as there may be interesting interactions between metaclasses and these features that have yet to be discovered.

Acknowledgments

We would like to thank Jan-Willem Maessen for his extremely helpful feedback on the formal rules presented in this paper. We also thank David Chase for not believing in us unless we were right, Victor Luchangco for many valuable discussions, and Guy Steele for his many helpful comments.

References

- [1] Eric Allen, Jonathan Bannet, and Robert Cartwright. A first-class approach to genericity. In *Proceedings of the 2003 ACM SIGPLAN conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 96–114. ACM Press, 2003.
- [2] Eric Allen, David Chase, Victor Luchangco, Jan-Willem Maessen, and Guy Steele. Object-Oriented Units of Measurement. In *Proceedings of the 2004 ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications*, 2004. Available at <http://research.sun.com/projects>, under Programming Languages.
- [3] Gilad Bracha and David Griswold. Strongtalk: Typechecking Smalltalk in a Production Environment. In *Proceedings of the OOPSLA '93 Conference on Object-Oriented Programming Systems, Languages and Applications*, pages 215–230, September 1993.
- [4] Luca Cardelli. A semantics of multiple inheritance. In *Proc. of the international symposium on Semantics of data types*, pages 51–67. Springer-Verlag New York, Inc., 1984.
- [5] Robert Cartwright and Guy L. Steele, Jr. Compatible genericity with run-time types for the java programming language. In *Proceedings of the 1998 ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 201–215. ACM Press, 1998.
- [6] Pierre Cointe. Metaclasses are first class: the ObjVlisp model. In *Proceedings of the OOPSLA '87 Conference on Object-Oriented Programming Systems, Languages and Applications*, pages 156–162. ACM Press, 1987.
- [7] William R. Cook. *A Denotational Semantics of Inheritance*. PhD thesis, Brown University, May 15 1989.
- [8] R. Kent Dybvig. Writing Hygienic Macros in Scheme with Syntax-Case. Technical report, Indiana University Computer Science Dept., July 03 1992.
- [9] Matthias Felleisen. On the expressive power of programming languages. In Neil D. Jones, editor, *ESOP'90, 3rd European Symposium on Programming*, volume 432 of *Lecture Notes in Computer Science*, pages 134–151, Copenhagen, Denmark, 15–18 May 1990. Springer.

- [10] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, Reading, Massachusetts, 1994.
- [11] Adele Goldberg and David Robson. *Smalltalk-80: The Language and Its Implementation*. Addison-Wesley, Reading, Massachusetts, 1989.
- [12] Justin O. Graver and Ralph E. Johnson. A type system for Smalltalk. In *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 136–150, San Francisco, California, January 1990.
- [13] Atshushi Igarashi, Benjamin Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 132–146. ACM Press, 1999.
- [14] Gregor Kiczales, Jim des Rivières, and Daniel G. Bobrow. *The Art of the Metaobject Protocol*. MIT Press, Cambridge, MA, 1991.
- [15] Michael Salib. Static Type Inference with Starkiller. In *PyCon DC, 2004*. <http://www.python.org/pycon/dc2004/papers/1/paper.pdf>.
- [16] Guy L. Steele Jr. *Common Lisp: The Language, Second Edition*. Digital Press, Bedford, Massachusetts, 1990.
- [17] Guido van Rossum. Unifying types and classes in Python 2.2. <http://www.python.org/2.2.2/descrintro.html>, 2002.
- [18] Christopher A. Welty and David A. Ferrucci. What’s in an instance? Technical report, Rochester Polytechnic Institute Computer Science Dept., 1994.