Integrating Nominal and Structural Subtyping

Donna Malayeri and Jonathan Aldrich

Carnegie Mellon University, Pittsburgh, PA 15213, USA,
{donna, aldrich}@cs.cmu.edu

Abstract. Nominal and structural subtyping each have their own strengths and weaknesses. Nominal subtyping allows programmers to explicitly express design intent, and, when types are associated with run time tags, enables run-time type tests and external method dispatch. On the other hand, structural subtyping is flexible and compositional, allowing unanticipated reuse. To date, nearly all object-oriented languages fully support one subtyping paradigm or the other.

In this paper, we describe a core calculus for a language that integrates the key aspects of nominal and structural subtyping in a unified framework. We have also merged the flexibility of structural subtyping with statically typechecked external methods, a novel combination. We prove type safety for this language and illustrate its practical utility through examples that are not easily expressed in other languages. Our work provides a clean foundation for the design of future languages that enjoy the benefits of both nominal and structural subtyping.

1 Introduction

In a language with structural subtyping, a type U is a subtype of T if its methods and fields are a superset of T's methods and fields. The interface of a class is simply its public fields and methods; there is no need to declare a separate interface type. In a language with nominal subtyping, on the other hand, U is a subtype of T if and only if it is *declared* to be. Accordingly, structural subtyping can be considered *intrinsic*, while nominal subtyping is *declarative*. Each kind of subtyping has its merits, but a formal model has not been developed for a language that integrates the two subtyping disciplines.

Structural subtyping offers a number of advantages. It is often more expressive than nominal subtyping, as subtyping relationships do not need to be declared ahead of time. It is compositional and intrinsic, existing outside of the mind of the programmer. This has the advantage of supporting unanticipated reuse—programmers don't have to plan for all possible scenarios. Additionally, structural subtyping is often more notationally succinct. Programmers can concisely express type requirements without having to define an entire subtyping hierarchy. In nominal systems, some situations may require multiple inheritance or an unnecessary proliferation of types; in a structural system, the desired subtyping properties just arise naturally from the base cases. Finally, structural subtyping is far superior in contexts where the structure of the data is of primary importance, such as in data persistent environments or distributed computing. In contrast, nominal subtyping can lead to unnecessary versioning problems: if some class C is

modified to C' (perhaps to add a method m), C' objects cannot be serialized and sent to a distributed process with the original definition C, even if C' is a strict extension of C.

As an example of the reuse benefits of structural subtyping, suppose class A has methods `foo()`, `a()` and `b()`, and class B has methods `foo()`, `x()` and `y()`. Suppose also that the code for A and B cannot be modified. In a language with structural subtyping, A and B share an implicit common interface { `foo` } and a programmer can write code against this interface. But, in a language with nominal subtyping, since the programmer did not plan ahead and create an `IFoo` interface and make A and B its subtypes, there is no way to write code that takes advantage of this commonality (short of using reflection). In contrast, with structural subtyping, if a class C is later added that contains method `foo()`, it too will share this implicit interface. If a programmer adds new methods to A, A's interface type will change automatically, without the programmer having to maintain the interface himself. If B or C also contain these new methods, the implicit combined interfaces will automatically exist.

Nominal subtyping also has its advantages. First, it allows the programmer to express and enforce design intent explicitly. A programmer's defined subtyping hierarchy serves as checked documentation that specifies how the various parts of a program are intended to work together. Second, explicit specification has the advantage of preventing "accidental" subtyping relationships, such as the standard example of `cowboy.draw()` and `circle.draw()`. Nominal subtyping also allows recursive types to be easily and transparently defined, since recursion can simply go through the declared names. Third, error messages are usually much more comprehensible, since, for the most part, every type in a type error is one that the programmer has defined explicitly. Finally, nominal subtyping enables efficient implementation of external dispatch.

External dispatch is provided by number of statically typed languages, such as Cecil [6, 7], MultiJava [8], among others. External methods increase the flexibility and evolvability of code because they do not fix in advance the set of methods of a class. Consider the example of a class hierarchy that represents AST nodes. (This motivating example is expanded further in Sec. 2.3.) Suppose this is part of a larger system, which includes an IDE for editing elements represented by this AST. Now suppose a programmer wishes to add new functionality to the IDE but cannot modify the original source code for the AST nodes. The new function provides the capability to jump from one node to a node that it references; this differs depending on what type of node is selected. Clearly, this functionality cannot be written without code that somehow performs dispatch on the AST class hierarchy.

In a language without external dispatch, the developer has limited choices. Usually, she must resort to manually writing `instanceof` tests, which is tedious and error-prone. In particular, if a new element is added to the AST hierarchy, the implementation will not behave correctly.

If the developers of the original class hierarchy *anticipated* this need and implemented the Visitor design pattern, it would then be easy to add new operations to the hierarchy, but then it would be difficult to add new classes. At best, Visitor trades one problem for another.

On the other hand, in a language with external dispatch, a programmer simply writes an *external method* that dispatches on the AST class hierarchy (i.e., separate from its

code). External dispatch makes it easy to adapt existing code to new interfaces, since new code can be added as an external method. Exhaustiveness checking rules for external methods ensure that when a new class is added to the hierarchy, in the absence of an inherited method, a new method must be added for that class.

Nominal subtyping enables efficient external dispatch since there is a name on which to tie the dispatch. Additionally, if external dispatch were allowed on structural types, the problem of accidental subtyping would be exacerbated, since overridden methods would apply wherever there was a structural match. Further, ambiguity problems could frequently arise, which would have to be manually resolved by the programmer. Consider, for example, a method m defined on objects with a `foo:int` field. If m is also later defined for objects with a `bar:char` field, m is now ambiguous—which method is called for an object with both fields?

In our language, Unity, we sidestep this issue—nominal and structural subtyping are integrated. This makes efficient external dispatch compatible with structural subtyping, but also gives programmers the benefits of both subtyping disciplines. Nominal subtyping gives programmers the ability to express explicit design intent, while structural subtyping makes interfaces easier to maintain and reuse.

Contributions. The contributions of this paper are as follows:

- A language design, Unity, that provides user-defined and structural subtyping relationships in a novel and uniform way. Unity combines the flexibility of external dispatch with the conceptual clarity of width and depth subtyping.
- A formalization of the design of Unity, along with proofs of type safety (Sec. 5).
- Examples that illustrate the expressiveness and flexibility of the language (Sec. 2), We contrast Unity with other languages in Sec. 2.1.
- A case study (Sec. 3) and an empirical study of several Java programs (Sec. 4).

2 Motivating Examples

We give, by example, the intuition behind Unity and illustrate combining structural subtyping with external methods. In Unity, an object type is a record type tagged with a *brand*. Brands induce the nominal subtyping relation, which we call "sub-branding."[1] Brands are nominal in that the user defines the sub-brand relationship, like the subclass relation in languages like Java, Eiffel, and C++.

When a brand β is defined, the programmer lists the fields that any objects tagged with β will include. In other words, if the user defines the brand `Point` as having the fieldss $\{$`x : int, y : int`$\}$, then any value tagged with `Point` must include at least the labels `x` and `y` (with `int` type)—but it may also contain additional fields, due to record subtyping. For instance, a programmer could create a colored point object with the expression `Point({x=0,y=1,color=blue})`. Subtyping takes into account both the nominal sub-brand relationship and the usual structural subtyping relationship (width and depth) on records.

[1] The name "brand" is borrowed from Strongtalk [3], which in turn borrowed it from Modula-3.

To integrate these two relationships, brand extension is constrained: the associated field types must be subtypes. In other words, a brand β_1 can be declared as a sub-brand of β_2 only if β_1's field type is a structural subtype of β_2's field type. As an example, suppose the brand `3DPoint` is defined as `3DPoint({x:int, y:int, z:int})`. `3DPoint` can be declared as a sub-brand of `Point`, since `{x:int, y:int, z:int}` is a sub-record of `{x:int, y:int}`. However, a brand `1DPoint({x:int})` cannot be a sub-brand of `Point` (since it lacks the `y` field), nor can `FloatingPoint(x:float, y:float})` (assuming `float` is not a subtype of `int`).

2.1 Example 1: A Window Toolkit

Fig. 1 contains a code excerpt for a windowing system and illustrates the novel features of Unity. The built-in brand `Top` is the root of the brand hierarchy, like `Object` in Java. To simplify the presentation, we include only the field `title`. `ScrollBar` is defined as a type alias using the `type` syntax. By default, a window does not have a scrollbar. The brands `Textbox` and `StaticText` extend `Window`, and also do not scroll by default.[2]

To add scrolling functionality, we have defined the function `scroll`, which operates on any `Window` (or sub-brand thereof) that has a `getScrollBar()` method. The type `Window({getScrollBar():ScrollBar})` classifies such an object. (We suppose here that the implementation of `scroll` need only access the scrollbar field and the fields of `Window`.) Note that the structural component of this type refers to another structural type, `ScrollBar`; structural types may be arbitrarily nested.

Let us assume that in a non-scrolling textbox, the user can only enter a fixed number of characters. Consequently, we define the brand `ScrollingTextbox` in order to change textbox functionality—in particular, the behavior when inserting a character. The `scroll` function is applicable to `ScrollingTextbox` since it is automatically a subtype of `Window({getScrollBar():ScrollBar})`.

In Unity, methods can be either internal (defined within a brand), or external (defined outside the brand). To allow sound modular checking of external methods (see Sec. 5), only internal methods are permitted to be abstract; external methods must be concrete. The method `insertChar` has been defined as an external method. This method is applicable to a `Textbox` or `ScrollingTextbox` that has a `getCurrentPos` method. `Textbox` does not have an internal `getCurrentPos` method, so it has been added as an external method. The method `getCurrentPos`, in turn, is only applicable to a `Textbox` that has a `pos:int` field. This illustrates the structural constraints that can be put on a method. For a method m, a programmer can specify a set of fields and methods that must be present in m's receiver.

Since a textbox that scrolls allows the user to enter more text than the window size permits, a new sub-brand had to be defined so that its implementation of `insertChar` could be overridden. If other sub-brands of `Window` (such as `StaticText`) do not need to change their existing behavior when a scrollbar is added, no new sub-brands need be

[2] Note that all fields must be listed by the subtypes of `Window`; this design decision is merely to simplify our core calculus.

```
abstract brand Window ({title : string}) extends Top
concrete brand Textbox ({title : string, text : string}) extends Window
concrete brand StaticText ({title : string, text : string}) extends Window
concrete brand ScrollingTextbox({title : string, text : string, s : ScrollBar};
                method getScrollbar() : ScrollBar = this.s)
                extends Textbox

type ScrollBar = Top(getMaximum():int, setMaximum(x:int) : unit,
                     getValue():int, setValue(x:int) : unit)


let scroll = λw : Window({getScrollBar() : ScrollBar}) .
    ... // code that performs the scrolling operation


method insertChar Textbox({getCurrentPos() : int}) : unit =
         λc : char .  ...   // insert a character only if it will fit in the window

method insertChar ScrollingTextbox({getCurrentPos() : int}) : unit =
         λc : char .  ...   // insert the character, scrolling if necessary

method getCurrentPos(Textbox({pos:int})) : int = ...
```

Subtyping relationships

```
Window ({title : string, s : ScrollBar}) ≤ Window ({title : string})

Textbox ({...}) ≤ Window ({title : string})
ScrollingTextbox ({...}) ≤ Textbox ({...})
ScrollingTextbox({...}) ≤ Window({title : string, s : ScrollBar})

StaticText({...}) ≤ Window ({title : string})
StaticText({..., s : ScrollBar}) ≤ Window(title : string, s : ScrollBar)
```

Fig. 1. Unity code for a windowing system. Nominal subtyping allows the brand `ScrollingTextbox` to change the behavior of `insertChar` through tag dispatch, while structural subtyping allows the `scroll` function to apply to any window with an `s : ScrollBar` field. `ScrollBar` is defined as a type alias using the `type` syntax. In the subtyping relationships, some field names are elided with "...."

defined. Scrolling functionality can be added to these types by including a `ScrollBar` field and a `getScrollBar()` method, and the `scroll` function is then applicable.

This example demonstrates both the use of functions (i.e., lambda expressions), and methods. The difference between the two is that functions do not perform dispatch (that is, they cannot be overridden), but they can be defined at any scope. Methods can be overridden, but they can only be defined at the top-level.

These brand definitions induce subtyping relationships, shown below the code listing in Fig. 1. Interestingly, `ScrollingTextbox({...})` is a subtype of both `Window({getScrollBar():ScrollBar})` and of `Textbox({...})`, but we have avoided both multiple inheritance and the problems typically associated with it. The

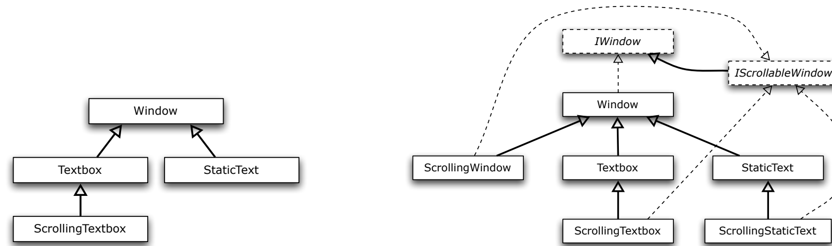

(a) The windowing example as implemented in Unity. Depicted here are the brands that must be defined in order to obtain the desired subtyping relationships.

(b) The same example implemented in Java. Dashed rectangles are interfaces; solid rectangles are classes. Dashed lines indicate the `implements` relationship and solid lines indicate `extends`.

Fig. 2. For the windowing example, the user-defined subtyping relationships necessary in Unity vs. those necessary in Java.

type `Window({getScrollBar():ScrollBar})` is a conceptual interface that exists without being named.

The example illustrates the two kinds of extensibility that Unity provides: structural extensibility and brand extensibility.

- **Structural types** can be used to create *structural method constraints*—methods that an object must have in order to conform to that type. They can be also be used to *create a new type that adds fields to an existing brand*, without defining new behavior for the resulting type. So, a `ScrollBar` can be added to a `StaticText` object without defining a new brand, as the existing functionality of the static text box does not need to change if a scrollbar is added.
- **Brand extension** creates a new brand that can be used in dispatch; as a consequence, programs can *define new behavior for the newly defined brand*. Here, `ScrollingTextbox` is defined as an extension of `Textbox` because the behavior of `insertChar` is different depending on whether or not the text box has a scrollbar attached to it. Design intent is preserved because whenever different behavior is required (such as with `Cowboy.draw()` and `Circle.draw()`), nominal subtyping must be used.

Additionally, we see here the synergy between structural subtyping and external dispatch. Structural subtyping can be used to *specify the constraints* of a method, and external methods can be used to make existing brands *conform* to those constraints.

2.2 Comparison to Other Systems

Here we compare Unity to closely related systems. See Sec. 6 for other related work.

Java. In Java-like languages, expressing this example would be unwieldy. A common way to express the necessary constraints would involve first defining two interfaces, `IWindow` and `IScrollableWindow`. `ScrollBar` would also have to be an interface.

If a programmer wished to allow the possibility of adding a scrollbar to a window class, even without changing any other functionality, he would have to define a subclass that also implemented `IScrollableWindow`. In this example, we would define the classes `ScrollingWindow`, `ScrollingTextbox`, and `ScrollingStaticText`, though only `ScrollingTextbox` needs to change any functionality; see the class diagram Fig. 2(b). Contrast this with the brand structure of the Unity program depicted in Fig. 2(a). In Unity, only the types associated with dispatch need to be defined, in contrast to Java.

The Java equivalent of the `scroll` function could be a static function of some helper class and would take an object of type `IScrollableWindow`. Of course, if a programmer defined a new scrolling window class with the correct `getScrollBar()` method, but forgot to implement `IScrollableWindow`, the `scroll` function could not be used on objects of that class. (This situation often arises in Java programs, particularly when one wishes to use library code, the developers of which are not even aware of the interface that they should implement.)

There are other oddities in the Java version. The Java class `ScrollingWindow` is semantically analogous to the Unity type `Window(getScrollBar():ScrollBar)`, but `ScrollingTextbox` and `ScrollingStaticText` are *not* subclasses of `ScrollingWindow`, while the corresponding Unity types *are* subtypes of `Window(getScrollBar():ScrollBar)`. To have such subtyping relationships would require multiple inheritance in a language like Java, while the Unity code works with just single inheritance. This illustrates the lack of expressiveness that is inherent in languages that require the programmer to name all relevant subtyping relationships.

Traits. A language with traits [25] would provide a much cleaner solution than that of Java, but would still lack the expressiveness of Unity's structural subtyping. This is because traits are mainly designed to solve issues of implementation inheritance (especially multiple inheritance) that are largely orthogonal to the ones we are considering. In this example, the same subtyping hierarchy would have to be created as in the Java example, but the `scroll` function could be written for `IScrollableWindow` with the appropriate dispatch. (A static method could always be written in Java, but it would not perform dispatch on subtypes.) This would enable some code reuse, but would still require creating a number of types.

Mixins. In a language with mixins [2], the programmer would create a mixin class `ScrollableWindow` that consists of the fields of `Window` along with `s : ScrollBar` and the code for the function `scroll`. The code for the `ScrollableWindow` would then be mixed into `StaticText` to create a `ScrollingStaticText` and into `Textbox` to create `ScrollingTextbox`. The behavior of `insertChar` would then be specialized for `ScrollingTextbox`.

With mixins, the same number of eventual classes would be created as in Java, but creating them becomes easier because of mixin construction. In contrast with Unity, the code for `scroll` cannot be reused unless the mixin `ScrollableWindow` is used,

which restricts its flexibility. This can pose a problem when interoperating with classes that were created in isolation from the mixin. Mixins also require up-front planning; methods and fields cannot be added after-the-fact.

Structural subtyping. Languages which support structural subtyping, such as Moby [10], O'Caml [14], and PolyTOIL [4], would elegantly express all of the desired subtyping relationships, but these languages allow only internal dispatch—that is, all methods must be defined inside the class definition. In our language, `insertChar` can be an external method; it need not reside inside the definitions of `Textbox` and `ScrollingTextbox`. It would be non-trivial to add support for external dispatch or multimethods to these types of languages.

Cecil. Cecil fully supports external and multimethod dispatch [6, 7]. Cecil's powerful, but very complex, type system can express most of the necessary relationships (though new classes do need to be defined for `ScrollingWindow` and `ScrollingStaticText`). To write the `scroll` function, a programmer would have to use bounded quantification and a "where" clause constraint, the latter being type-checked via a constraint solver. That is, in psuedo-code, the argument to `scroll` would have type:

$$\text{for all T where T} <: \text{Window and signature getScrollBar(T) : ScrollBar}$$

Here, the type `ScrollBar` would have to be a class, rather than a structural type as in Unity, due to two major shortcomings of `where` clauses: they cannot be nested and can only occur on top level methods. Additionally, `where` clauses cannot be used to constrain the method's receiver. In Unity, on the other hand, structural types are compositional and can therefore be nested within another type (e.g., `ScrollBar` in Fig. 1), can occur at any level in the program (e.g., the lambda expression `scroll`), and can be used to constrain a method's receiver (e.g., method `insertChar`).

Virtual classes. Some of the required relationships could be expressed elegantly using virtual classes [15] or nested inheritance [21], but only with advance planning. To express this example, the programmer would create a class `Base` containing the virtual classes `Window`, `Textbox`, and `StaticText`. Then, she would create a subclass of `Base`, called `Scroll`, that contained its own `Window`. This definition of `Window` would add a field for a scrollbar. Additionally, `Scroll` would have a virtual class `ScrollingTextbox` which would include the new definition of `insertChar`. The programmer would not to create a new class `ScrollingStaticText` since the new definition of `Window` would automatically apply to `StaticText` (i.e., `Scroll.StaticText` would automatically have a scrollbar).

The virtual classes solution is elegant, but if the programmer did not plan ahead and redefine `Window` in the `Scroll` class, there would not be a way to describe this type. Essentially, virtual classes make it very easy to reuse code across related classes (an advantage of virtual classes and nested inheritance over Unity), but cannot easily express the structural types of Unity.

```
abstract brand AstNode( ;{abstract method compile : () → unit}) extends Top

concrete brand PlusNode ({n1 : AstNode(), n2 : AstNode()};
   method compile() : unit = compilePlus(this); /* compile PlusNode */)
   extends AstNode
concrete brand Num ({val : int}; method compile() : unit = ... /* compile Num*/)
   extends AstNode
concrete brand Var ({s : Symbol}; method compile() : unit = ... /* compile Var */)
   extends AstNode

// highlight the text corresponding to 'node' in the text editor,
// using the location specified by the 'loc' field
let highlightNode = λnode : AstNode(loc : Location). ...

// AST nodes with debug information
concrete brand DebugPlusNode ({n1 : AstNode(), n2 : AstNode(), loc : Location};
   method compile() : unit = compilePlus(this); outputLocation(out, this.loc))
   extends PlusNode
concrete brand DebugNum ({val : int, loc : Location};
   method compile() : unit = ... /* compile DebugNum */) extends Num
concrete brand DebugVar ({s : Symbol, loc : Location, varName : string};
   method compile() : unit = ... /* compile DebugVar */) extends Var
```

Fig. 3. Example 2: AST Nodes in an IDE. The top portion is the code before changes to add debug information to the AST. The function `highlightNode` makes use of structural information and the external dispatch in `compile` changes its behavior for the declared `Debug*` sub-brands.

2.3 Example 2: AST Nodes in an IDE

In this section, we describe another example to show other ways in which Unity can be used. Suppose we have an integrated development environment that includes an editor and a compiler. The top portion of Fig. 3 contains an excerpt of a simplified version of the code for such a system. Here, the brands `PlusNode`, `Num` and `Var` define a simple abstract syntax tree. The internal method `compile` performs compilation on an `AstNode`.

One can use structural subtyping to create AST nodes with additional information, such as a node with a `loc` field specifying the file location of the code corresponding to the node. Additional functions are available for such nodes, such as the function `highlightNode` that highlights a node's source code in the text editor.

We did not have to define a new brand for AST nodes that include file location information. Whether or not a node contains file information, functions that operate over AST nodes need not change their behavior, so in this case structural subtyping suffices.

Suppose now that the programmer wishes to add "debug" versions of these AST nodes that contain additional output information for compiling in debug mode. For example, a `DebugNum` has a `Location` field, while `DebugVar` includes a `Location` field as well as a string representation of the variable name. The newly added code is listed in the bottom portion of Fig. 3.

Since each of these brands have been defined as extensions, they may also customize the behavior of `compile` to output this additional information when compiling. Additionally, since all of the `Debug*` brands have a `Location` field, the function `highlightNode` can be used on objects of this type.

This example again illustrates the expressiveness that achieved by combining nominal and structural subtyping; `highlightNode` makes use of additional structural information, while `compile` relies on nominal dispatch to behave differently in different situations.

2.4 Real-World Examples

The following real-world examples illustrate the gains in flexibility that could be achieved through structural subtyping.

Eclipse SWT. In the Eclipse SWT (Simple Windowing Toolkit), many classes (such as `Button`, `Label`, `Link`, etc.) have the methods `getText` and `setText`, that set the main text for the control, such as a button's text, the text in a textbox, etc. However, there is no common `IText` interface. Many classes—13 in total—also support adding an image through the `getImage` and `setImage` method, but again there is no interface that captures this. A programmer may wish to write a method that sets the image of any control by retrieving the image from an image registry. Given the current API, such a method would have to rely on runtime reflection, with no guarantee of successful method invocation at compile time.

Eclipse JDT. In the JDT (Java Development Tools), there are 8 classes (including `IMethod`, `IType`, `IField`) that have the method `getElementName`, but there is no `IElement` interface with this method. With structural subtyping, these classes implicitly share an interface, and code could be written that is polymorphic over the exact class type. For instance, a tree view of an AST may wish to display packages, methods, and fields in a uniform way. With the current hierarchy it is not possible to simply call the `getElementName` of the object, since these classes do not have an explicit interface with this method.

3 Case Study: Optional Methods in Java

In this section, we describe the tradeoffs that a library designer must make when using a language that has only nominal subtyping. The design of the Java collections library illustrates that designers would rather circumvent the type system than have a proliferation of types. We believe this situation often occurs with nominal subtyping, but because of structural subtyping, such a situation need not occur in Unity.

In the Java collections library, the interface `java.util.Collection` has several "optional" methods: `add`, `addAll`, `clear`, `remove`, `removeAll`, and `retainAll`. Many of the abstract classes implementing `Collection` (e.g., `AbstractCollection`, `AbstractList`, `AbstractSet`) throw an `UnsupportedOperationException` when

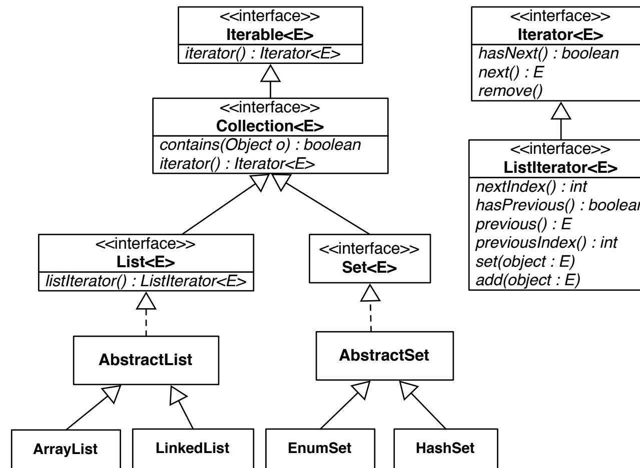

Fig. 4. A portion of the Java collections framework. A few methods are highlighted in most interfaces. Type parameters are elided in classes.

those methods are called. There are a total of 30 optional methods in `java.util.*`, and `java.lang.Iterator` has an additional optional method. The methods were designed this way to avoid an explosion of interfaces such as `MutableCollection`, `ImmutableCollection`, etc., and a corresponding increase in the number of sub-interfaces (e.g., `MutableList`, `ImmutableList`, etc.) [18].

Let us consider a Java collections framework without the optional methods. Figure 4 shows a relevant portion of the current Java collections hierarchy. Figure 5 show new interfaces that capture the distinction of mutability directly in the hierarchy—doing away with optional operations. The interface Collection<E> represents a collection that is modifiable, while the new interface ReadableCollection<E> represents a collection that only contains read operations. Accordingly, its `iterator()` method returns a `ReadIterator`, which is defined without a `remove()` operation. There are now two new `ListIterator` interfaces, for fixed-size lists, modifiable lists, read-only lists. These correspond to the FixedSizeList<E> and ReadableList<E> interfaces in Figure 5. (The interface FixedSizeList<E> has been added because selective overriding of methods in `AbstractList` would yield such a type, as noted in the documentation.) The hierarchy for `Set` is similar to that of `List` (though simpler, since there are no fixed-size sets, and no set-specific iterator).

In a language with structural subtyping, such as Unity, not all interesting combinations of structural types have to be declared in advance (though in a library setting they might be, for consistency's sake). For a language with type abbreviations, the key idea is that a type abbreviation would simply be syntactic sugar for a set of methods, which could be given an abbreviation with a different name in another part of the system. Additionally, the subtyping relationships between all the interfaces would not need to be defined in advance. Finally, as a side note, the notational overhead in defining type ab-

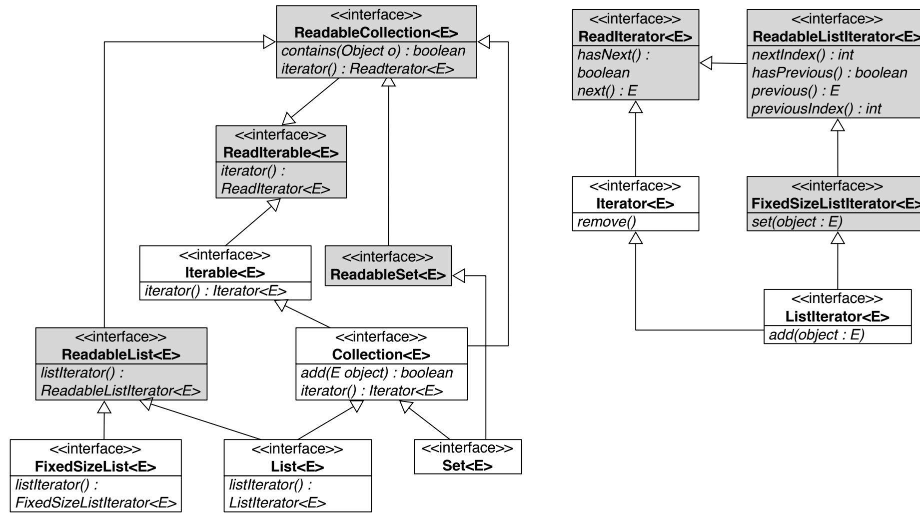

Fig. 5. Refactored `AbstractList` and `AbstractSet` classes, along with new interfaces to remove optional methods. A few methods of most interfaces are highlighted for List and Set; for iterators, all methods, except for inherited methods, are shown. New interfaces have a gray background.

breviations would be potentially far lower than that of defining a Java interface, which has a relatively high notational cost (due, in part to the nominal nature of interfaces).

For this example, in Unity, the new interfaces (shown in gray), would not necessarily need to be defined by the library author, unless specifically needed. This eases the task of library development, as the library author does not need to anticipate which supertypes of the given interfaces would be useful for clients.

Thus, with a combination of nominal and structural subtyping, we need not sacrifice static type safety in order to overcome the shortcomings of a purely nominal type system.

4 Empirical Analysis

To determine if there are potential cases where structural subtyping would be useful in a real system, we ran an analysis of 15 open-source Java programs. The analysis searches for common method signatures that are not related through inheritance. A "common method" is any method declaration where there exists in another class a method declaration with an identical name and the same signature, but the method is not present in any common supertype of the two.

For instance, in Apache Collections, four buffer classes had the methods `increment` and `decrement`, but these were not contained in a common superclass or super-interface. The results of the analysis are in Fig. 6. Tomcat, a servlet container, had the greatest percentage of common methods, 28.4%. Ant, the software build system, was close behind with 28.1%. Even the programs with the smallest number of

Fig. 6. Results of empirical analysis. For each program, the total number of methods, the number of common methods, the percentage of common methods compared to the total, the average number of methods in each common method group, and the number of common method groups are displayed. "Total methods" includes interface methods, abstract methods, and overriding implementations of the same method. The results suggest that many Java programs have potential uses of structural subtyping.

	Total methods	Common methods	% Common	Average methods/group	Total common groups
Tomcat	14678	4172	28.4%	3.2	1288
Ant	9178	2577	28.1%	3.5	727
JHotDraw	5149	1193	23.2%	2.8	428
Smack	3921	881	22.5%	3.3	270
Struts	3783	772	20.4%	2.7	291
Apache Forrest	164	28	17.1%	2.2	13
Cayenne	9243	1545	16.7%	2.8	553
Log4j	1950	312	16.0%	3.1	102
OpenFire	8135	1300	16.0%	2.8	470
Apache Collections	3762	584	15.5%	2.8	211
Derby	24521	3575	14.6%	2.5	1402
Lucene	2472	331	13.4%	2.5	134
jEdit	5845	699	12.0%	2.6	271
Apache HttpClient	1818	217	11.9%	2.6	83
Areca	3565	423	11.9%	2.6	163

common methods had a significant number of them; Areca, a backup program, had 11.9% common methods.

Inspecting the common methods, we found several cases where a structural type could be useful. For instance, in Smack, a Jabber client library, there were 6 classes with the common method `String getElementName()`. There were also 30 classes with the method `String toXML()` which might also be a method that clients might wish to call in a uniform manner. In JHotDraw, a GUI framework, there were 9 classes that had `addPropertyChangeListener` and `removePropertyChangeListener` methods. In Log4j, a logging library, there were 4 classes with the method `int getBufferSize()`, and 8 classes with the method `setOption`. In the Apache Collections library, nearly all the common methods appeared to be potentially useful. For instance, there were 5 iterator decorator classes with `getIterator` and `setIterator` methods. 4 bag classes had a method `getBag`, 4 buffer classes had a method `getBuffer`, and 4 classes had the method `getComparator`. Additionally, 7 classes had the method `int size()` and 5 classes had the method `int indexOf(Object)`.

Note that we did not closely examine the implementations of these methods to determine if they were semantically performing similar actions, as we were unfamiliar with the codebase. So, it is possible that the methods coincidentally had similar names but were performing different actions. In future work, we plan to study one application in depth to see how it can benefit from structural subtyping.

Programs	$p ::=$	$decl$ in $p \mid e \mid e; p$
Declarations	$decl ::=$	$brand\text{-}decl \mid ext\text{-}decl$
Brand declaration	$brand\text{-}decl ::=$	mod brand $\beta(\tau; \overline{m\text{-}decl})$ extends β
Modifiers	$mod ::=$	abstract $\mid$ concrete
Method declaration	$m\text{-}decl ::=$	abstract method $m\ (\overline{m : \tau}) : \tau$
		$\mid$ method $m\ (\overline{m : \tau}) : \tau = e$
External method	$ext\text{-}decl ::=$	method $m\ \beta(\overline{m : \tau_m}) : \tau = e$
Expressions	$e ::=$	$() \mid x \mid \lambda x{:}\tau.\, e \mid e\, e \mid \widehat{\beta}(e) \mid \{\overline{\ell = e}\} \mid e.\ell \mid e.m$
Types	$\tau ::=$	unit $\mid \tau \to \tau \mid \tau \wedge \tau \mid \beta(\overline{m : \tau}) \mid \{\overline{\ell : \tau}\} \mid \tau \Rightarrow \tau$
Values	$v ::=$	$() \mid \widehat{\beta}(v) \mid \{\overline{\ell = v}\} \mid \lambda x{:}\tau.\, e \mid$
Contexts	$\Gamma ::=$	$\cdot \mid \Gamma, x : \tau$
	$\Sigma ::=$	$\cdot \mid \Sigma, mod\ \beta(\tau; \overline{mod\ m : \tau})$ extends β
	$\Delta ::=$	$\cdot \mid \Delta, \widehat{\beta}(\overline{m = e})$ extends $\widehat{\beta}$

Conventions

$\widehat{\beta} \equiv$ tag value corresponding to β
fieldType$_\Sigma(\beta) = \tau$ if $\beta(\tau; \overline{m : \tau})$ extends $\beta' \in \Sigma$
modifier$_\Sigma(\beta) = mod$ if $mod\ \beta(\tau; \overline{m : \tau})$ extends $\beta' \in \Sigma$
M ranges over $\overline{m : \tau}$

Fig. 7. Unity grammar

Overall, however, we found the results to be promising, and they suggest that Java programs could indeed benefit from structural subtyping. If a programmer wished to write code that called a common method, he could easily do so by using a type—which exists implicitly—consisting of that method. In contrast, in Java and other languages with nominal subtyping, programmers would have to explicitly create interfaces. And, in some cases, the interface would contain only one method, which seems an unnecessary overhead.

5 Formal System

The Unity grammar is presented in Fig. 7. The language is a lambda calculus extended with values tagged with brands. Methods can be defined on a brand and the usual dispatch semantics apply. Brand and method declarations are top-level. To define brands, the brand construct is used. A brand can be either abstract or concrete. Objects cannot be created from abstract brands (similar to Java's abstract classes). We use the metavariables β and θ to range over brand names. The metavariable M ranges over a list of (method : type) pairs.

One of the valid expression forms for a program is an expression followed by a program ($e; p$). In this last construct, the expression is evaluated and will be type correct according to the definitions that preceded it.

When a brand is defined, a name is given for it, as well as the brand's *field type* (usually a record); this is the type of the fields of the brand. The programmer initializes the field value when an object is created.

In Unity, a method is either internal or external. In the former case, the method is defined along with the brand, like method declarations in Java-like languages. To allow modular exhaustiveness checking of external methods, external methods cannot be abstract; a method body must be provided for every external method. We have taken this rule from MultiJava [8].

When a method m is defined on a brand β, a set of methods is specified—the methods that must exist within β (either internal or external) before m can be invoked.[3] For example, in Fig. 1, the function `insertChar` required that its receiver have a `getCurrentPos` method.

To simplify the formal system, methods take only one argument: the this parameter. Additional parameters may be specified using lambda expressions.[4]

If β is a brand name, then $\widehat{\beta}$ is the tag value corresponding to β. In other words, $\beta(\overline{m : \tau})$ is a type, and $\widehat{\beta}$ is its run-time analogue.

To create objects, the expression form $\widehat{\beta}(e)$ is used. This creates an object that is tagged with $\widehat{\beta}$. Methods are called using $e.m$, while function application is written $e_1 e_2$.

Our language includes a limited form of intersection types. Our motivation for including these is to make external methods available to objects that were defined before the external method was created. Section 5.1 describes this in more detail.

Σ is the subtyping context; it stores the user-declared sub-branding relationships. Δ is the corresponding run-time context. Δ contains a strict subset of the information in Σ—it does not contain whether a brand is abstract or concrete, and it does not keep track of the field type or methods associated with each brand. We assume the existence of a special brand Top that is not defined in Σ or Δ, but that may be extended by user-defined brands. Since every brand must have a super-brand, the brand subtype hierarchy is a tree rooted at Top.

Like other object calculi, Unity is purely functional so as to simplify the system. State is orthogonal to the issues we are considering; our design should be easily adaptable to a language with imperative features.

5.1 Static Semantics

Here we describe the subtyping and typing judgements shown in Figures 8, 9 and 10. Auxiliary judgements are in Fig. 11

[3] For simplicity and to support information hiding, types cannot contain field constraints as in example 1, but this is not a fundamental limitation of the system.

[4] Note that if the body of a method is a lambda expression, it does not perform dispatch. To perform dispatch, the body of the method should be another method call. In this way, asymmetric multimethods (multimethods where the order of parameters is used in dispatch) can easily be encoded in our system. To encode a method m with body e that dispatches on β_1 and β_2, method m in β_1 dispatches to method m in β_2, the body of which is e.

$$\boxed{\Sigma \vdash \tau_1 \leq \tau_2}$$

$$\frac{}{\Sigma \vdash \tau \leq \tau} \qquad \frac{\Sigma \vdash \tau_1 \leq \tau_2 \qquad \Sigma \vdash \tau_2 \leq \tau_3}{\Sigma \vdash \tau_1 \leq \tau_3} \qquad \frac{\begin{array}{cc}\Sigma \vdash \beta_1 \sqsubseteq \beta_2 & \Sigma \vdash M_1 \leq M_2 \\ \Sigma \vdash \beta_1(M_1)\ \mathbf{type} & \Sigma \vdash \beta_2(M_2)\ \mathbf{type}\end{array}}{\Sigma \vdash \beta_1(M_1) \leq \beta_2(M_2)}$$

$$\frac{\Sigma \vdash \sigma_1 \leq \tau_1 \qquad \Sigma \vdash \tau_2 \leq \sigma_2}{\Sigma \vdash \tau_1 \to \tau_2 \leq \sigma_1 \to \sigma_2} \qquad \frac{\Sigma \vdash \tau \leq \sigma_1 \qquad \Sigma \vdash \tau \leq \sigma_2}{\Sigma \vdash \tau \leq \sigma_1 \wedge \sigma_2} \qquad \frac{}{\Sigma \vdash \tau_1 \wedge \tau_2 \leq \tau_1}$$

$$\frac{}{\Sigma \vdash \tau_1 \wedge \tau_2 \leq \tau_2} \qquad \frac{\{\ell_i : \tau_i\ ^{i \in 1..n}\} \text{ is a permutation of } \{\ell_j : \tau_j\ ^{j \in 1..n}\}}{\Sigma \vdash \{\ell_i : \tau_i\ ^{i \in 1..n}\} \leq \{\ell_j : \tau_j\ ^{j \in 1..n}\}}$$

$$\frac{n > m}{\Sigma \vdash \{\ell_i : \tau_i\ ^{i \in 1..n}\} \ \leq\ \{\ell_j : \tau_j\ ^{j \in 1..m}\}} \qquad \frac{\Sigma \vdash \tau_i \leq \sigma_i \ (i \in 1..n)}{\Sigma \vdash \{\ell_i : \tau_i\}\ ^{i \in 1..n} \leq \{\ell_i : \sigma_i\}\ ^{i \in 1..n}}$$

$$\frac{\Sigma \vdash \beta_1 \sqsubseteq \beta_2}{\Sigma \vdash \beta_1(M_1) \wedge \beta_2(M_2) \leq \beta_1(M_1 \wedge M_2)}$$

$$\frac{\Sigma \vdash \beta_1 \sqsubseteq \beta_2 \qquad \Sigma \vdash M_2 \leq M_1 \qquad \Sigma \vdash \sigma_1 \leq \sigma_2}{\Sigma \vdash \beta_1(M_1) \Rightarrow \sigma_1 \ \leq\ \beta_2(M_2) \Rightarrow \sigma_2} \qquad \frac{\Sigma \vdash \{\overline{m : \tau}\} \leq \{\overline{n : \sigma}\}}{\Sigma \vdash \overline{m : \tau} \leq \overline{n : \sigma}}$$

Fig. 8. Unity subtyping judgement

Subtyping. Subtyping comprises two parts: the sub-brand judgement ($\sqsubseteq$) and the subtype judgement ($\leq$). The latter is shown in Fig. 8. The first judgement is not on types, but brands, which are a component of a type but not themselves a type. The sub-brand judgement is just the reflexive, transitive closure of the declared extends relation.

The subtype judgement ($\leq$) uses the sub-brand judgement in the third subtyping rule, which states that an object type $\beta_1(M_1)$ is a subtype of $\beta_2(M_2)$ when β_1 is a sub-brand of β_2 ($\beta_1 \sqsubseteq \beta_2$) and M_1 is a sub-record of M_2 ($M_1 \leq M_2$). There are additional conditions that $\beta_1(M_1)$ **type** and $\beta_2(M_2)$ **type**, which ensures that these are valid types. The relevant type formation rule here is:

$$\frac{\begin{array}{cc}\overline{m} \text{ distinct} & \Gamma \mid \Sigma \vdash \overline{\sigma}\ \mathbf{type} \\ \Sigma \vdash \beta \ \mathsf{extends}\ \beta_2 & \mathit{override}_{\Sigma}(\overline{m : \sigma}, \beta_2)\end{array}}{\Gamma \mid \Sigma \vdash \beta(\overline{m : \sigma})\ \mathbf{type}}$$

This rule checks that the given labels and types are a sub-record of the required fields for the brand. This ensures that a brand type always contains at least the labels it was defined to have. There is an additional check that the methods are valid overrides (*override* is defined in Fig. 11). The rest of the rules for the type formation judgement are straightforward; the full judgement appears in Appendix A.

Our language includes a limited form of intersection types, à la Davies and Pfenning; the rules for intersection types are borrowed from their work [9].

There is also a subtyping rule for a list of (method : type) pairs; it simply applies the record subtyping rule. The remaining subtyping rules are the standard reflexivity, transitivity, and function subtyping rules.

$$\boxed{\Sigma \vdash p\ \mathbf{ok}}$$

$$\frac{\begin{array}{c}\beta \notin \Sigma \qquad \tau \leq \text{fieldType}_\Sigma(\beta') \\ \Sigma \vdash \beta.\overline{m\text{-}decl} : (\overline{mod_m m : \tau}) \qquad \Sigma' = \Sigma, mod\ \beta(\tau; \overline{mod_m\ m : \tau})\ \mathsf{extends}\ \beta' \\ \Sigma' \vdash \beta.(\tau; \overline{m\text{-}decl})\ \mathbf{ok} \qquad override_\Sigma(\overline{m : \tau}, \beta') \\ mod = \mathsf{concrete\ implies}\ methods_{\Sigma'}(\beta) = \mathsf{concrete}\ \overline{n : \sigma} \qquad \Sigma' \vdash p\ \mathbf{ok}\end{array}}{\Sigma \vdash mod\ \mathsf{brand}\ \beta(\tau; \overline{m\text{-}decl})\ \mathsf{extends}\ \beta'\ \mathsf{in}\ p\ \mathbf{ok}}\ (\text{TP-BRAND})$$

$$\frac{\begin{array}{c}\Sigma = \{mod\ \beta_1(\sigma; M')\ \mathsf{extends}\ \beta_2\}, \Sigma_0 \\ m \notin M' \qquad \Sigma' = \{mod\ \beta_1(\sigma; M', m : \beta_1(M) \Rightarrow \tau)\ \mathsf{extends}\ \beta_2\}, \Sigma_0 \\ override_\Sigma(\beta_1(M) \Rightarrow \tau, \beta_2) \\ \mathsf{this} : \beta_1(M), \mathsf{fields} : \sigma \mid \Sigma' \vdash e : \tau \qquad \Sigma' \vdash p\ \mathbf{ok}\end{array}}{\Sigma \vdash \mathsf{method}\ m\ \beta_1(M) : \tau = e\ \mathsf{in}\ p\ \mathbf{ok}}\ (\text{TP-EXT-METHOD})$$

$$\frac{\cdot \mid \Sigma \vdash e : \tau}{\Sigma \vdash e\ \mathbf{ok}}\ (\text{TP-EXPR1}) \qquad \frac{\cdot \mid \Sigma \vdash e : \tau \qquad \Sigma \vdash p\ \mathbf{ok}}{\Sigma \vdash e; p\ \mathbf{ok}}\ (\text{TP-EXPR2})$$

Fig. 9. Unity typing judgement for programs

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

$$\frac{x : \tau \in \Gamma}{\Gamma \mid \Sigma \vdash x : \tau} \qquad \frac{}{\Gamma \vdash () : \mathsf{unit}} \qquad \frac{\Sigma \vdash \tau_1\ \mathbf{type} \qquad \Gamma, x : \tau_1 \mid \Sigma \vdash e : \tau_2}{\Gamma \mid \Sigma \vdash \lambda x{:}\tau_1.\, e : \tau_1 \rightarrow \tau_2}$$

$$\frac{\Gamma \mid \Sigma \vdash e_1 : \tau_1 \rightarrow \tau_2 \qquad \Gamma \mid \Sigma \vdash e_2 : \tau_1}{\Gamma \mid \Sigma \vdash e_1\, e_2 : \tau_2} \qquad \frac{\Gamma \mid \Sigma \vdash e : \sigma \qquad \Sigma \vdash \sigma \leq \tau}{\Gamma \mid \Sigma \vdash e : \tau}$$

$$\frac{\begin{array}{c}\mathsf{concrete}\ \beta(\tau) \in \Sigma \\ \Gamma \mid \Sigma \vdash e : \tau' \qquad \Sigma \vdash \tau' \leq \tau \qquad methods_\Sigma(\beta) = \overline{mod_m\ m : \sigma}\end{array}}{\Gamma \mid \Sigma \vdash \widehat{\beta}(e) : \beta(\overline{m : \sigma})}\ (\text{TP-NEW-OBJ})$$

$$\frac{\Gamma \mid \Sigma \vdash \overline{e} : \overline{\tau}}{\Gamma \mid \Sigma \vdash (\overline{\ell = e}) : \{\overline{\ell : \tau}\}} \qquad \frac{\Gamma \mid \Sigma \vdash e : \{\ell_i : \tau_i\ ^{i \in 1..n}\}}{\Gamma \mid \Sigma \vdash e.\ell_k : \tau_k}$$

$$\frac{\begin{array}{c}\Gamma \mid \Sigma \vdash e : \beta(M) \qquad m_k : \tau_{m_k} \in (M \wedge methods_\Sigma(\beta)) \\ \tau_{m_k} = \beta'(\overline{n : \sigma}) \Rightarrow \tau \qquad \beta(M \wedge methods_\Sigma(\beta)) \leq \beta'(\overline{n : \sigma})\end{array}}{\Gamma \mid \Sigma \vdash e.m_k : \tau}\ (\text{TP-INVOKE})$$

Fig. 10. Unity typing judgement for expressions

Typing rules. Full typing rules for typechecking programs and expressions appear in Figs. 9 and 10, respectively. Auxiliary judgements are defined in Fig. 11. The interesting rules are TP-BRAND, TP-EXT-METHOD, TP-NEW-OBJ and TP-INVOKE; the others are standard.

The rule TP-BRAND (Fig. 9) ensures that a brand declaration is well-formed. The newly defined brand must contain at least the labels and fields of the supertype (possibly

$$\boxed{\Sigma \vdash \textit{m-decl} : \tau}$$

$$\frac{\Sigma \vdash \beta(\overline{m : \sigma}) \Rightarrow \tau \ \mathbf{type}}{\Sigma \vdash \beta.\mathit{mod}\ \mathsf{method}\ m_1(\overline{m : \sigma}) : \tau = e \ : \ \mathit{mod}\ m_1 : \beta(\overline{m : \sigma}) \Rightarrow \tau}$$

$$\boxed{\Sigma \vdash \beta.(\tau; \textit{m-decl})\ \mathbf{ok}}$$

$$\frac{\mathsf{this} : \beta(\overline{m : \sigma}), \mathsf{fields} : \tau \mid \Sigma \vdash e : \tau'}{\Sigma \vdash \beta.(\tau; \mathsf{method}\ m_1(\overline{m : \sigma}) : \tau' = e)\ \mathbf{ok}}$$

$$\boxed{\mathit{methods}_\Sigma(\beta) = \overline{\mathit{mod}\ m : \tau}}$$

$$\frac{}{\mathit{methods}_\Sigma(\mathsf{Top}) = \cdot}$$

$$\frac{\Sigma \vdash \beta_1(\tau; \overline{\mathit{mod}_1\ m : \tau_m}, M)\ \mathsf{extends}\ \beta_2 \qquad \mathit{methods}_\Sigma(\beta_2) = \overline{\mathit{mod}_2\ m : \sigma_m}, M'}{\mathit{methods}_\Sigma(\beta_1) = \overline{\mathit{mod}\ m : \tau_m}, M, M'}$$

$$\boxed{\mathit{override}_\Sigma(m : \tau, \beta)}$$

$$\frac{\mathit{methods}_\Sigma(\beta) = (\mathit{mod}\ m : \sigma, M)\ \text{implies}\ \tau \leq \sigma}{\Sigma \vdash \mathit{override}_\Sigma(m : \tau, \beta)}$$

$$\boxed{M_1 \wedge M_2}$$

$$(m_i : \tau_{m_i}{}^{i \in 1..n}, M) \wedge (m_i : \tau'_{m_i}{}^{i \in 1..n}, M') \stackrel{\text{def}}{=} \left(m_i : (\tau_{m_i} \wedge \tau'_{m_i})^{i \in 1..n}, M, M'\right)$$
$$\text{where } m_i, M, M' \text{ are mutually exclusive}$$

Fig. 11. Unity typechecking auxiliary judgements

with refined types); this is checked via the condition $\tau \leq \text{fieldType}_\Sigma(\beta')$. Note that if a field type is a record, then subtypes must list all the labels of the parent. Aside from simplifying the calculus, this sidesteps issues of variable shadowing while allowing subtypes to refine the type of a particular label. The rule also checks that the methods given are valid overrides of the methods of the super-brand, and, in the case of concrete classes, that all methods are concrete.

This rule and the type formation rule for brands described above illustrate the need for both a sub-brand and subtype judgement. The context Σ stores information about the fields and methods of a brand; these are retrieved via fieldType_Σ and $\mathit{methods}_\Sigma$ (called by $\mathit{override}_\Sigma$), respectively. Additionally, without a runtime component to the nominal hierarchy, there would not be a way to perform dispatch, which we describe in Sect. 5.2.

The rule TP-EXT-METHOD checks external method definitions. The existing brand definitions are updated by adding the new external method via the new context Σ'. The rule also checks that the method types of the external method defined on sub-brands are in the subtype relation, which ensures that the context Σ' is well-formed.

The rule TP-NEW-OBJ (Fig. 10) checks the correctness of the object creation expression. The rule checks that the brand has been defined as concrete, and that the given record labels are a subtype of the required record labels.

The rule TP-INVOKE typechecks method invocations. The method being called must be contained in either the set of methods in the brand's type, M, or in the set of methods of the brand ($\mathit{methods}_\Sigma(\beta)$). Additionally, the methods in the brand's type, combined with the methods of the brand (via intersection) must be a subtype of the method's required methods. Adding the intersection condition increases expressiveness over having the rule just consider the methods of the brand, since the type might

$$\boxed{p \mid \Delta \longmapsto p' \mid \Delta'}$$

$$\frac{\overline{m\text{-}decl} \longmapsto \overline{m = e}}{\begin{array}{c} mod\ \mathsf{brand}\ \beta_1(\tau; \overline{m\text{-}decl})\ \mathsf{extends}\ \beta_2\ \mathsf{in}\ p \mid \Delta \longmapsto \\ p \mid \Delta, (\beta_1(\overline{m = e})\ \mathsf{extends}\ \beta_2) \end{array}} \text{ (E-BRAND-DECL)}$$

$$\frac{\Delta = \{\beta(\overline{m = e})\ \mathsf{extends}\ \beta'\}, \Delta_0}{\begin{array}{c} \mathsf{method}\ m_1\ \beta(\overline{m : \tau_m}) : \sigma = e_1\ \mathsf{in}\ p \mid \Delta \longmapsto \\ p \mid \{\beta(\overline{m = e}, m_1 = e_1)\ \mathsf{extends}\ \beta'\}, \Delta_0 \end{array}} \text{ (E-EXT-DECL)}$$

$$\frac{e \longmapsto_\Delta e'}{e \mid \Delta \longmapsto e' \mid \Delta} \qquad \frac{e \longmapsto_\Delta e'}{e; p \mid \Delta \longmapsto e'; p \mid \Delta} \qquad \frac{p \mid \Delta \longmapsto p' \mid \Delta'}{v; p \mid \Delta \longmapsto p' \mid \Delta'}$$

$$\boxed{e \longmapsto_\Delta e'}$$

$$\frac{e \longmapsto_\Delta e'}{e.m \longmapsto_\Delta e'.m} \qquad \frac{mbody_\Delta(m, \widehat{\beta}) = e}{\widehat{\beta}(v).m \longmapsto_\Delta [\widehat{\beta}(v)/\mathsf{this}, v/\mathsf{fields}]\, e} \text{ (E-INVOKE)}$$

$$\frac{e \longmapsto_\Delta e'}{\widehat{\beta}(e) \longmapsto_\Delta \widehat{\beta}(e')} \qquad \frac{e_1 \longmapsto_\Delta e_1'}{e_1\, e_2 \longmapsto_\Delta e_1'\, e_2} \qquad \frac{e_2 \longmapsto_\Delta e_2'}{v_1\, e_2 \longmapsto_\Delta v_1\, e_2'} \qquad \frac{}{(\lambda x{:}\tau.\, e)\, v \longmapsto_\Delta [v/x]\, e}$$

$$\frac{e_k \longmapsto_\Delta e_k'}{\begin{array}{c} \{\ldots, \ell_k = e_k, \ldots\} \longmapsto_\Delta \\ \{\ldots, \ell_k = e_k', \ldots\} \end{array}} \qquad \frac{e \longmapsto_\Delta e'}{e.\ell \longmapsto_\Delta e'.\ell} \qquad \frac{}{\{\ell_i = v_i\ ^{i \in 1..n}\}.\ell_k \longmapsto_\Delta v_k}$$

Fig. 12. Unity evaluation judgement

have methods defined on a sub-brand. For example, within the body of the function $\lambda x : \mathsf{Top}(\texttt{toString} : \texttt{()} \to \mathsf{string}).\, e$, the type of x contains the method `toString`. If we suppose that `toString` is not defined for the brand Top, then x's type contains methods that are not defined in the brand itself.

5.2 Dynamic Semantics

Most of the evaluation rules for Unity are standard; the evaluation judgement is in Fig. 12 and auxiliary judgements are in Fig. 13.

The interesting evaluation rules are E-BRAND-DECL and E-EXT-DECL, which evaluate brand definitions and external method definitions, respectively. To evaluate a brand definition, the method definitions are evaluated to the method body and the rest of the program is evaluated with the extended context. Similarly, E-EXT-DECL updates the context with new method definitions for the brand, then evaluates the rest of the program with the new context.

The auxiliary function $mbody_\Delta(m, \widehat{\beta})$ finds the appropriate method body for a method m, starting at the tag $\widehat{\beta}$. This function is used by the rule E-INVOKE, which within the method body returned by $mbody_\Delta$, substitutes the object for `this` and the

$\boxed{mbody_{\Delta}(m, \widehat{\beta}) = e}$

$$\frac{\widehat{\beta}_1(m_0 = e_0, \overline{m' = e'_m}) \text{ extends } \widehat{\beta}_2 \in \Delta}{mbody_{\Delta}(m_0, \widehat{\beta}_1) = e_0} \qquad \frac{\widehat{\beta}_1(\overline{m = e}) \text{ extends } \widehat{\beta}_2 \in \Delta \quad m_0 \notin \overline{m} \quad mbody_{\Delta}(m_0, \widehat{\beta}_2) = e_0}{mbody_{\Delta}(m_0, \widehat{\beta}_1) = e_0}$$

$\boxed{m\text{-}decl \longmapsto m = e}$

$$\frac{}{\text{abstract method } m(\overline{m : \sigma_m}) : \tau \longmapsto \cdot} \qquad \frac{}{\text{method } m(\overline{m : \sigma_m}) : \tau = e \longmapsto m = e}$$

Fig. 13. Unity evaluation auxiliary judgements

field value of the object for `fields`. Method declarations are evaluated in a straightforward manner; all of the type information is discarded (so in the case of abstract methods, the entire declaration is discarded), leaving just the method body.

5.3 Type Safety

The full proof of type safety is provided in a companion technical report [17]. We summarize the main results here. First, we provide the definition of a well-formed context:

Definition 1 (Well-formed context).
The context Σ is *well-formed*, iff the following conditions hold:

1. there is exactly one entry for each brand β.
2. if $mod\ \beta_1(\tau; M)$ extends $\beta_2 \in \Sigma$, then
   (a) $\beta_2(M)$ **type**
   (b) $\tau \leq \text{fieldType}_{\Sigma}(\beta_2)$
   (c) if mod = concrete, then $methods_{\Sigma}(\beta_1) = \text{concrete } \overline{n : \tau}$.

Our theorems on type safety assume a correspondence between the static brand definition context Σ and the runtime context Δ. This ensures that the runtime context, which does not contain type information, is consistent with the static typing context. Formally, this correspondence is defined as follows:

Definition 2 (Models relation on contexts). The definition of $\Sigma \vdash \Delta$ (Σ *models* Δ) is given by the following inference rules:

$$\frac{}{\cdot \vdash \cdot} \qquad \frac{\begin{array}{c}\Sigma \vdash \Delta \\ \Sigma' = \Sigma, mod\ \beta_1(\tau; \{\text{concrete } m_i : \beta_1(M_i) \Rightarrow \tau'_i{}^{\ i \in 1..n}\}, \overline{\text{abstract } n : \sigma}) \text{ extends } \beta_2 \\ \text{this} : \beta_1(M_i), \text{fields} : \tau \mid \Sigma' \vdash e_i : \tau_i \ (i \in 1..n)\end{array}}{\Sigma' \vdash \Delta, \widehat{\beta}_1(m_i = e_i{}^{\ i \in 1..n}) \text{ extends } \widehat{\beta}_2}$$

Type safety is proved using the standard progress and preservation theorems. For progress, we prove a lemma that states that if we have a well-typed value whose type contains a method m_k, then a runtime context consistent with the static context must contain a method body for m_k:

Lemma 1. *If* $\Gamma \mid \Sigma \vdash \widehat{\beta}(v) : \tau$ *and* $\Sigma \vdash \tau \leq \beta'(M)$*, where* $\Sigma \vdash \Delta$ *and* $m_k \in M$*, then* $mbody_{\Delta}(m_k, \widehat{\beta})$ *is defined.*

The lemma is stated in this way so that the subsumption case is easy to prove. The lemma is proved by induction on the typing derivation. The interesting case is that of TP-NEW-OBJ, which uses the definition of a well-formed context and that of $\Sigma \vdash \Delta$.

Theorem 1 (Progress [programs]). *If* $\cdot \mid \Sigma \vdash p$ **ok***, for some* Σ*, then either* p *is a value or, for* Δ *such that* $\Sigma \vdash \Delta$*, there exist* p' *and* Δ' *such that* $p \mid \Delta \longmapsto p' \mid \Delta'$*.*

This theorem is proved by appealing to an auxiliary lemma that proves progress for expressions and a standard canonical forms lemma. The interesting case is that of method invocation, which is proved using Lemma 1.

Preservation is slightly more difficult to prove. We first prove the following lemma by induction on the typing derivation. The lemma states that the body of a method is well-typed if the static context Σ models the runtime context Δ.

Lemma 2. *If* $\Gamma \mid \Sigma \vdash \widehat{\theta}(v) : \sigma$ *and* $\sigma \leq \beta(m_0 : \beta'(M_0) \Rightarrow \tau, M)$ *and* $\Sigma \vdash \Delta$ *and* $mbody_{\Delta}(m_0, \widehat{\theta}) = e_0$*, then* $\mathsf{this} : \beta'(M_0), \mathsf{fields} : fieldType_{\Sigma}\theta \mid \Sigma \vdash e_0 : \tau$*.*

Theorem 2 (Preservation [programs]). *If* $\Gamma \mid \Sigma \vdash p$ **ok** *and* $\Sigma \vdash \Delta$ *and* $p \mid \Delta \longmapsto p' \mid \Delta'$*, then there exists a* Σ' *such that* $\Sigma' \vdash \Delta'$ *where* $\Gamma \mid \Sigma' \vdash p'$ **ok***.*

We prove this theorem using of a preservation theorem on expressions, a standard substitution lemma, and Lemma 2 above.

5.4 Modularity

Our typechecking rules are modular; each rule relies only on information in the context up to the current program point, rather than requiring a global dictionary of brand definitions. Our exhaustiveness checks are modular because external method definitions cannot be abstract (enforced by the grammar); otherwise, information about all brand definitions would be required.

Since our language does not include modules, our ambiguity checks are not modular in the strictest sense of the term, as they depend on all definitions up to the current program point. However, our system could be easily extended with additional rules to support modular ambiguity checking. Millstein and Chambers have developed such rules and have also defined several levels of modular typechecking [19]. Our current system is compatible with their broadest notion of modular typechecking, the so-called "most-extending module" approach, exemplified by their language System E. To perform the most modular form of typechecking, however, we would require that all implementations of an external method be in the same module. Further, external methods would be forbidden from overriding internal methods (currently permitted in our system). These checks correspond to the restriction $M1$ in Dubious [19] and restriction $R3$ in MultiJava [8].

A related issue is that of information hiding, a form of which our language supports. A brand's field value can only be accessed by the brand's methods, effectively making them private. It would be possible to extend this further and disallow external methods from accessing fields, or allow marking some internal methods as private.

5.5 Polymorphism and Recursive Types

We have designed an extension Unity$_\alpha$ with polymorphism (described in [17]), but we have omitted this feature in this version of Unity since we discovered that polymorphism was orthogonal to the issues surrounding nominal and structural subtyping. In Unity$_\alpha$, the syntax is extended as follows:

$$\begin{aligned}
\textit{brand-decl} &::= \textit{mod}\ \textsf{brand}\ \forall \overline{T}.\, \beta\langle \overline{T} \rangle(\tau; \overline{\textit{m-decl}})\ \textsf{extends}\ \beta\langle\overline{\tau}\rangle \\
\textit{ext-decl} &::= \textsf{method}\ m\ \forall\overline{T'}.\, \beta\langle\overline{T'}\rangle(\overline{m:\tau}) : \tau = e \\
e &::= \ \dots \mid \widehat{\beta}[\overline{\tau}] \mid \Lambda T.\, e \mid e[e] \\
\tau &::= \ \dots \mid X \mid \beta\langle\overline{\tau}\rangle(\overline{m:\tau}) \mid \forall T.\, \tau
\end{aligned}$$

The sub-brand judgement is on parameterized brands (i.e. $\beta\langle\overline{\tau}\rangle$) and, aside from a new rule for $\forall T.\, \tau$ types, the subtype judgement is essentially the same.

We have also created an extension that includes structural recursive types [17]. Unity as presented in this paper supports only nominal recursive types; when defining a brand β, the name β can be used in the components of its definition. Adding structural recursive types was relatively straightforward; we simply added standard iso-recursive μ types to the language, along with a fold and unfold operation. In this system, it is possible to express types such as:

$$\mu X.\textsf{Top}(\texttt{clone()} : X)$$

which specifies that the result of the clone function is the type itself being defined. The advantage to structural recursive types is that structural object interfaces, such as `ScrollBar` in Example 1, can be specified as pure structural types (using the Top brand) while still being self-referential.

6 Related Work

Type Systems. At the FOOL/WOOD '07 workshop, we presented the predecessor of this version of Unity [16]. Here, we have extended that work by adding methods and information hiding to our core calculus, providing additional examples, and including a case study.

Researchers have recently considered the problem of integrating nominal and structural subtyping. Reppy and Turon have addressed the problem in the context of type-checking traits [24]. Their resulting type system is a hybrid of nominal and structural subtyping. However, in their system, structural types are second-class; they apply to trait functions but not to expressions or ordinary functions. Consequently, there is less expressiveness as compared with Unity: it is not possible to constrain the argument of a function to have particular members, for example.

After our initial workshop proposal, Odersky et al. independently implemented a similar language feature, validating the practical importance of our work. In Scala, type refinements allow a nominal type to include additional structural information [22]. Scala type refinements have many similarities with the language Whiteoak, an extension of Java with structural types [12]. Like Scala, in Whiteoak, by using intersection types, a type can include both structural and nominal aspects.

Scala and Whiteoak differ from Unity in that they do not have external methods, nor do they allow structural constraints to be placed on a method's receiver. Also, the language designs have neither been formalized nor proved sound.

Ostermann has designed a language that seeks to enhance the expressiveness of nominal subtyping to gain some of the benefits of structural subtyping [23]. Ostermann has identified an additional important benefit of nominal subtyping—that of blame assignment: i.e., who accepts responsibility for maintaining a subtype relation, the user or the designer of a component? The language design is much more expressive than a purely nominal system; it is possible to, for example, create subtypes of a class type without inheriting its implementation, and declare supertypes of an existing type. But, this comes at the cost of a subtyping relation that is not transitive, which may prove counter-intuitive to programmers. The programmer must manually provide a set of "witness" types so that the typechecker can apply subsumption. Therefore, it is unclear whether this approach is practical.

Bono et al. have also proposed a type system that includes both nominal and structural aspects, but their system does not fully integrate the two disciplines [1]. The system only uses structural typing when typechecking uses of the `this` variable, making their system considerably less expressive than ours.

The language MOBY is in many ways similar to Unity, as it supports structural subtyping and a form of tag subtyping through its inheritance-based subtyping mechanism, which is similar to our sub-branding [10, 11]. This allows expressing many useful subtyping constraints, but MOBY's class types are not integrated with object types in the same way as in Unity. For instance, in MOBY, it is not possible to express the constraint that an object should have a particular class *and* should have some particular methods (that are not defined in the class itself). Additionally, the object-oriented core of MOBY supports only internal dispatch. MOBY does include "tagtypes" that are very similar to our brands. These can be used to support downcasts or to encode multimethods, but they are disjoint from the object-oriented core of the system.

Strongtalk presents a structural type system for Smalltalk and also supports named subtyping relationships through its "brand" mechanism [3]. However, it is not possible to define subtyping on brands. Additionally, since it is a type system for Smalltalk, it supports only the single dispatch model.

Modula-3's type system has structural types with branding, but not structural sub*typing* [20]. That is, its type system will treat two record types as equivalent if they have the same structure but different type aliases, but does not recognize one as a subtype of the other if it has additional fields. The object-oriented part of the language uses nominal subtyping.

In the C++ concepts proposal, concepts can be either nominal or structural [13]. However, concepts apply only to template constraints, not to the subtyping relation.

External and Multimethod Dispatch. External and multimethod dispatch has been extensively studied, but in the context of either dynamically typed languages, or languages with a purely nominal type system. Cecil is one of the first languages to include statically checked multimethods, but performs a whole-program analysis to ensure that multimethods are exhaustive and unambiguous [6, 7]. As previously mentioned (Sect. 2.2),

Cecil contains "where" clauses that can model some aspects of structural types, but they can only appear on top-level methods and cannot be nested, in contrast to Unity.

More recent work has focused on modular typechecking of external methods and multimethods, as well as the problem of integrating external methods into existing languages; this includes the Dubious calculus (System M) and MultiJava [19, 8]. We have built on these existing techniques for modular typechecking of external methods.

The language $\lambda\&$ [5] includes multimethod dispatch and includes structural subtyping on methods, similar to Unity. However, the subtyping hierarchy on classes uses only nominal subtyping, in contrast to Unity.

Acknowledgements We would like to thank Karl Crary and William Lovas for helpful discussion and feedback on our language, and Kevin Bierhoff for comments on an earier version of this paper. This research was supported in part by the U.S. Department of Defense, Army Research Office grant number DAAD19-02-1-0389 entitled "Perpetually Available and Secure Information Systems," and NSF CAREER award CCF-0546550.

References

[1] Viviana Bono, Ferruccio Damiani, and Elena Giachino. Separating Type, Behavior, and State to Achieve Very Fine-grained Reuse. In *Electronic proceedings of FTfJP'07 (http://www.cs.ru.nl/ftfjp/)*, 2007.

[2] G. Bracha and W. Cook. Mixin-based inheritance. In *ECOOP '90*, 1990.

[3] Gilad Bracha and David Griswold. Strongtalk: typechecking Smalltalk in a production environment. In *OOPSLA '93*, pages 215–230, 1993.

[4] Kim B. Bruce, Angela Schuett, Robert van Gent, and Adrian Fiech. PolyTOIL: A type-safe polymorphic object-oriented language. *ACM Trans. Program. Lang. Syst.*, 25(2):225–290, 2003.

[5] Giuseppe Castagna, Giorgio Ghelli, and Giuseppe Longo. A calculus for overloaded functions with subtyping. *Inf. Comput.*, 117(1):115–135, 1995.

[6] Craig Chambers. Object-oriented multi-methods in Cecil. In *ECOOP '92*, 1992.

[7] Craig Chambers and the Cecil Group. The Cecil language: specification and rationale, version 3.2. Available at `http://www.cs.washington.edu/research/projects/cecil/`, February 2004.

[8] Curtis Clifton, Todd Millstein, Gary T. Leavens, and Craig Chambers. MultiJava: Design rationale, compiler implementation, and applications. *ACM Trans. Program. Lang. Syst.*, 28(3):517–575, 2006.

[9] Rowan Davies and Frank Pfenning. Intersection types and computational effects. In *ICFP '00*, pages 198–208, 2000.

[10] Kathleen Fisher and John Reppy. The design of a class mechanism for Moby. In *PLDI '99*, pages 37–49, 1999.

[11] Kathleen Fisher and John Reppy. Inheritance-based subtyping. *Inf. Comput.*, 177(1):28–55, 2002.

[12] Joseph Gil and Itay Maman. Whiteoak. Available at `http://ssdl-wiki.cs.technion.ac.il/wiki/index.php/Whiteoak`, 2008.

[13] Douglas Gregor, Jaakko Järvi, Jeremy Siek, Bjarne Stroustrup, Gabriel Dos Reis, and Andrew Lumsdaine. Concepts: Linguistic support for generic programming in C++. In *Proceedings of OOPSLA '06*, pages 291–310. ACM Press, October 2006.

[14] Xavier Leroy, Damien Doligez, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. The Objective Caml system, release 3.09. Available at `http://caml.inria.fr/pub/docs/manual-ocaml/index.html`, 2004.
[15] O. L. Madsen and B. Moller-Pedersen. Virtual classes: a powerful mechanism in object-oriented programming. In *OOPSLA '89*, pages 397–406, 1989.
[16] Donna Malayeri and Jonathan Aldrich. Combining structural subtyping and external dispatch. In *FOOL/WOOD'07*, January 2007. Available at `http://foolwood07.cs.uchicago.edu/program.html`.
[17] Donna Malayeri and Jonathan Aldrich. Integrating Nominal and Structural Subtyping. Technical Report CMU-CS-08-120, School of Computer Science, Carnegie Mellon University, May 2008.
[18] Sun Microsystems. Java collections API design FAQ. Available at `http://java.sun.com/j2se/1.4.2/docs/guide/collections/designfaq.html`, 2003.
[19] Todd D. Millstein and Craig Chambers. Modular statically typed multimethods. *Inf. Comput.*, 175(1):76–118, 2002.
[20] Greg Nelson, editor. *Systems programming with Modula-3*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1991.
[21] Nathaniel Nystrom, Stephen Chong, and Andrew C. Myers. Scalable extensibility via nested inheritance. In *OOPSLA '04*, pages 99–115, 2004.
[22] Martin Odersky. The Scala language specification. Available at `http://www.scala-lang.org/docu/files/ScalaReference.pdf`, 2007.
[23] K. Ostermann. Nominal and Structural Subtyping in Component-Based Programming. *Journal of Object Technology*, 7(1), 2008.
[24] John Reppy and Aaron Turon. Metaprogramming with traits. In *ECOOP 2007*, July-August 2007.
[25] Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Andrew Black. Traits: Composable units of behavior. In *ECOOP '03*, 2003.

A Formal System

Well-formed judgement

$$\boxed{\Sigma \vdash \tau\ \textbf{type}}$$

$$\frac{}{\Sigma \vdash \mathsf{unit}\ \textbf{type}} \qquad \frac{\Sigma \vdash \tau_1\ \textbf{type} \qquad \Sigma \vdash \tau_2\ \textbf{type}}{\Sigma \vdash \tau_1 \to \tau_2\ \textbf{type}} \qquad \frac{\Sigma \vdash \tau_1\ \textbf{type} \qquad \Sigma \vdash \tau_2\ \textbf{type}}{\Sigma \vdash \tau_1 \wedge \tau_2\ \textbf{type}}$$

$$\frac{\begin{array}{c}\overline{m}\ \text{distinct} \qquad \Sigma \vdash \overline{\sigma}\ \textbf{type} \\ \Sigma \vdash \beta\ \mathsf{extends}\ \beta_2 \qquad \Sigma \vdash \mathit{override}(\overline{m : \sigma}, \beta_2)\end{array}}{\Sigma \vdash \beta(\overline{m : \sigma})\ \textbf{type}} \qquad \frac{\overline{\ell}\ \text{distinct} \qquad \Sigma \vdash \overline{\tau}\ \textbf{type}}{\Sigma \vdash \{\overline{\ell : \tau}\}\ \textbf{type}}$$

$$\frac{\Sigma \vdash \beta(m_i : \theta_i(\overline{n_i : \sigma_i}) \Rightarrow \tau_i'^{\ i \in 1..n})\ \textbf{type} \qquad \Sigma \vdash \beta \sqsubseteq \theta_i \ \ (i \in 1..n) \qquad \Sigma \vdash \tau_2\ \textbf{type}}{\Sigma \vdash \beta(m_i : \theta_i(\overline{n_i : \sigma_i}) \Rightarrow \tau_i'^{\ i \in 1..n}) \Rightarrow \tau_2\ \textbf{type}}$$