

Featherweight Defenders: A formal model for virtual extension methods in Java

Brian Goetz and Robert Field, Oracle Corporation*

March 27, 2012

1 Introduction

As a means of modeling the semantics of *virtual extension methods* (also known as *defender methods*) in Java, we describe a lightweight model of Java in the style of *Featherweight Java* (Pierce et al.), called *Featherweight Defenders* (or FD).

This model retains many aspects of Java’s class and method inheritance. We drastically simplify method overloading and naming (classes and interfaces can have only one method, named `m()`, with no arguments), but retain single inheritance of classes and multiple inheritance of interfaces, covariant overriding, abstract and concrete class methods, and reabstraction of concrete class methods (where a concrete method is overridden with an abstract one).

We add in a new feature not yet present in Java, *virtual extension methods*, where an interface can provide a default implementation for a method. Interface methods may be abstract definitions (no default clause) or specify a concrete default implementation (`default k`). Concrete defaults may be overridden by abstract methods (by overriding the method without providing a new default), which is akin to reabstraction of concrete methods in classes.

The goal of this model is twofold: to capture the rules regarding well-formedness of types in the presence of extension methods, and to capture the runtime mapping of method names within a class to concrete method bodies which provide the implementation of that method for that class, which we call the *linkage* of the method `m()` in C .

We believe that this model includes the most significant inheritance features that are relevant to method linkage in the presence of extension methods. There are some additional features relevant to the implementation, such as bridge methods, that are beyond the scope of this model.

*The authors gratefully acknowledge the assistance of Sukyoung Ryu of KAIST in the refinement of this formal model.

$$\begin{aligned}
T & ::= \text{Object} \mid C \mid I \\
K & ::= \text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ [R m() \langle b \mid \text{abstract} \rangle] \} \\
L & ::= \text{interface } I \text{ extends } I_1, \dots, I_n \{ [R m() [\text{default} k]] \}
\end{aligned}$$

Figure 1: FD language syntax

2 Syntax

The metavariables C , D , and E (and their derivatives) range over class names and the metavariable I , J , and K range over interface names. The metavariables T , R , U , V , and W range over all types (classes and interfaces). The metavariable S ranges over sets of types. The metavariable k ranges over a set of nominal identifiers, and the metavariable b ranges over a set of method bodies. Figure 1 shows the syntatic forms for FD.

3 Preliminaries

Figure 2 shows the subtyping judgements for classes and interfaces.

$$\begin{aligned}
& \text{S-REFL} \frac{}{T <: T} \\
& \text{S-TRANS} \frac{T <: V \quad V <: W}{T <: W} \\
& \text{S-CLASS} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \}}{C <: D \quad \forall_i C <: I_i} \\
& \text{S-INTF} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ \dots \}}{\forall_i I <: I_i}
\end{aligned}$$

Figure 2: Basic subtyping rules

As in Featherweight Java, we use “lookup functions” (such as $mtype(T)$) and “marker predicates” (such as $T \text{ OK}$) in the inference rules to record information about types. These will be introduced as they are used by the inference rules.

4 Compile-time vs runtime

One of the goals of this model is to present a formal procedure for linking a method $m()$ in a class C among the many candidate choices contributed by superclasses and interfaces. In languages like Java, method linkage analysis is performed both at compile time (to ensure that methods link uniquely, that

required methods are implemented, and to reject source files that do not meet these requirements) and at run time (to perform linkage dynamically.) It is important to ensure that linkage decisions made at compile time and run time are consistent, but issues like separate compilation and dynamic linking pose challenges to this goal, since class files compiled separately may not provide a consistent view of the type hierarchy at run time, which may lead to runtime errors.

This model divides rules into two categories – typing-related (those beginning with T-) and linkage-related (those beginning with R-). A compiler would execute and enforce all the rules; the runtime would execute only the linkage rules. In a well-typed program, the compiler and runtime view of the world should be identical; in an inconsistently-typed program (say, due to separate compilation), it may still be possible to perform method linkage at runtime according to the R- rules. The treatment of separate compilation and dynamic linkage is outside the scope of this model.

5 Method typing

Figure 3 shows the typing rules for resolving the type of method $m()$ in classes and interfaces. It defines the lookup function $mtype(T)$, which describes the return type of the method $m()$ in T .

Central to the typing analysis is identifying whether a method $m()$ in type T (whether defined explicitly or inherited) can legally override all signatures of $m()$ from supertypes of T . Java supports covariant overrides – where a method in a subclass overrides a method in a supertype, but provides a narrower return type. If a method $m()$ in T overrides one or more methods from supertypes, the return type of $m()$ in T must be a subtype of each of the return types of $m()$ from supertypes. If there is such a type, we say $m()$ has a *well-defined return type* in T .

The value of $mtype(T)$ is a *set* of types, and is one of the following:

- The empty set $\{\}$, indicating that the method $m()$ is not a member of T ;
- The set $\{ U \}$, indicating that the method $m()$ has a well-defined return type in T , and that return type is U ;
- *undefined*, indicating that the method $m()$ does not have a well-defined return type in T . The compiler should reject types for which $mtype(T)$ is *undefined*, but having a well-defined return type is not sufficient to declare that T is well-formed (for example, T could have conflicting defaults or problems with covariant overrides.)

By defining $mtype$ in this way, we can explicitly distinguish between the case where we know that type T does not have a member $m()$ ($mtype(T) = \{\}$), and the case where we cannot assign a well-defined type for $m()$ in T ($mtype(T) = \textit{undefined}$). If we bind a symbol to $mtype(T)$ in the precondition of a typing judgement (such as the binding of S in T-CLASSNONE), and $mtype(T)$ is *undefined*, then the binding fails and therefore the judgement does not apply.

Having defined $mtype$ in this way, we define a relation for optional subtyping between a candidate type T and the return type of $m()$ in U :

$$T \tilde{<}: mtype(U) \equiv mtype(U) = \{ \} \vee [mtype(U) = \{ V \} \wedge T <: V]$$

Intuitively, if $T \tilde{<}: mtype(U)$, then T is a potentially valid return type for the method $m()$ in a subtype of U .

We extend this concept by defining the function lbi_{mtype} , which computes an *inclusive lower bound* under optional subtyping for a projection under $mtype$ of a set of types. (The inclusive lower bound for a set S is the lower bound of S if S contains its lower bound, and *undefined* otherwise.) We define $lbi_{mtype}(T_1, \dots, T_n)$ as follows:

$$lbi_{mtype}(T_1, \dots, T_n) = \begin{cases} \{ \} & \text{if } \forall_i mtype(T_i) = \{ \} \\ mtype(T_i) & \text{if } \exists_i \text{ such that } mtype(T_i) = \{ U \}, \text{ and} \\ & \forall_{j \neq i} U \tilde{<}: mtype(T_j) \\ \text{undefined} & \text{otherwise} \end{cases}$$

We use lbi_{mtype} to determine whether there is a consistent return type for the method $m()$ in T given the types of $m()$ in the supertypes of T . (Note that if $mtype(T_i)$ is *undefined* for any i , then $lbi_{mtype}(T_1, \dots, T_n)$ is also *undefined*.)

$$\text{T-CLASSEXP} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ T \ m() \langle b \mid \text{abstract} \rangle \} \quad T \tilde{<}: mtype(D) \quad \forall_i T \tilde{<}: mtype(I_i)}{mtype(C) = \{ T \}}$$

$$\text{T-CLASSIMPL} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \} }{mtype(C) = lbi_{mtype}(D, I_1, \dots, I_n)}$$

$$\text{T-INTEXP} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ T \ m() [\text{default } k] \} \quad \forall_i T \tilde{<}: mtype(I_i)}{mtype(I) = \{ T \}}$$

$$\text{T-INTIMPL} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ \} }{mtype(I) = lbi_{mtype}(I_1, \dots, I_n)}$$

$$\text{T-OBJECT} \frac{}{mtype(\text{Object}) = \{ \} \quad \text{Object OK}}$$

Figure 3: Method typing

6 Method inheritance

The rules so far only have to do with computing the membership of a method in a type, and the method’s return type in that type, and do not consider where the method implementation codes from. We now explore the rules for inheriting methods, including inheritance of bodies from superclasses and inheritance of default implementations from superinterfaces.

All classes (except the root class `Object`) have a single superclass. We treat a concrete method body and a declaration that the method is abstract in the same way (collectively, we call these a method declaration). A class inherits method declarations from its superclass, unless the method is explicitly given a new declaration in the subclass.

A key aspect of inheritance of default methods is *pruning* of less-specific default-providing interfaces from consideration in the linkage process. In Java, it is allowable for a class or interface to extend an interface both directly and indirectly, as in the following example:

```
interface A { Object m() default k }
interface B extends A { Object m() default l }
class C implements A, B { }
```

Here, `C` implements `A` both directly and indirectly. This idiom is common as a documentation device, and in Java 7 and earlier the additional declaration of `A` has no effect, because it is already implicit in the extension of `B`. This behavior should continue to hold true in the presence of extension methods.

The inheritance rules for extension methods calls for “redundant” inheritance from less-specific interfaces (such as `A` in the example above) to not be considered further in the inheritance decision, except inasmuch as the less-specific interface has already contributed to its subinterface.

We capture this pruning behavior by the function $prune(S)$:

$$prune(S) = \{ W \in S : \forall_{V \in S} V <: W \Rightarrow V = W \}$$

Intuitively, the rules for method linkage behave as follows:

- A method defined in a type takes precedence over methods defined in its supertypes (R-INTDEF, R-CLASSBODY).
- A method declaration (concrete or abstract) inherited from a superclass takes precedence over a default inherited from an interface (R-LINKIMPL).
- More specific default-providing interfaces take precedence over less specific ones (R-INTINH, R-CLASSINH).
- If we are to link `m()` to a default method from an interface, there must a unique most specific default-providing interface to link to (R-LINKDEF).

Figure 4 covers the rules for computing the candidates for inheriting methods from classes and interfaces. It defines the following lookup functions:

- $icand(T)$ (interface candidates), which indicates the set of interfaces which could provide a default implementation for `m()` in `T`. This is used both in

member inheritance calculations, as well as identifying the valid targets for a super-call.

- $mprov(C)$ (method provenance), which indicates the most specific superclass from which C could inherit a method declaration for $m()$. Like $mtype$, the value of $mprov(C)$ is a set, which is either empty (there is no superclass that provides a declaration for $m()$ in C), or a singleton set that contains the class providing the declaration for $m()$ in C .
- T **HasBody**, which indicates that T declares either a concrete method body (for classes) or a default implementation (for interfaces). (It is not inherited; it is strictly a property of the class or interface declaration, and will be used in the final linkage decision.)

$$\begin{array}{c}
\text{R-OBJECT} \frac{}{mprov(\text{Object}) = \{\} \quad icand(\text{Object}) = \{\}} \\
\\
\text{R-INTDEF} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ T m() [\text{default } k] \}}{icand(I) = \{ I \}} \\
\\
\text{R-INTINH} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ \} \\ S = \bigcup_i icand(I_i)}{icand(I) = \text{prune}(S)} \\
\\
\text{R-CLASSINH} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\ S = \bigcup_{U \in \{ D, I_1, \dots, I_n \}} icand(U)}{icand(C) = \text{prune}(S)} \\
\\
\text{R-CLASSBODY} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ T m() \langle b \mid \text{abstract} \rangle \}}{mprov(C) = \{ C \}} \\
\\
\text{R-CLASSINHBODY} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \}}{mprov(C) = mprov(D)} \\
\\
\text{R-HASBODY} \frac{\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ T m() b \}}{C \text{ HasBody}} \\
\\
\text{R-HASDEF} \frac{\text{interface } I \text{ extends } I_1, \dots, I_n \{ T m() \text{ default } k \}}{I \text{ HasBody}}
\end{array}$$

Figure 4: Inheritance candidates from classes and interfaces

$$\begin{array}{c}
\text{interface } I \text{ extends } I_1, \dots, I_n \{ \dots \} \\
\text{T-INTNoBODY} \frac{S = \text{mtype}(I) \quad \forall_i I_i \text{ OK} \quad \text{icand}(I) = \{ \}}{I \text{ OK}} \\
\\
\text{interface } I \text{ extends } I_1, \dots, I_n \{ \dots \} \\
\forall_i I_i \text{ OK} \quad \text{icand}(I) = \{ J \} \\
\text{T-INTInHBody} \frac{\text{mtype}(I) = \text{mtype}(J)}{I \text{ OK}}
\end{array}$$

Figure 5: Typechecking of interfaces

7 Typechecking of interfaces and classes

We now consider the rules for when an interface or class should be accepted as well-typed. We use the marker predicate $T \text{ OK}$ to indicate that T is well-typed.

Figure 5 shows the typechecking rules for interfaces. For an interface I , $I \text{ OK}$ means:

- All the supertypes of I are OK;
- If $m()$ is a member of I , it has a well-defined return type in I ;
- If $m()$ is a member of I , then there is at most one defender candidate for $m()$ in I , and if present, that defender candidate has the same return type as $m()$ in I .

Figure 6 shows the typechecking rules for classes. For a class C , $C \text{ OK}$ means:

- All the supertypes of C are OK;
- If $m()$ is a member of C , it has a well-defined return type in C ;
- If C inherits a method declaration from a superclass, it has the same return type as $m()$ in C .
- If C does not inherit a method declaration from a superclass, then there is at most one defender candidate for $m()$ in C , and if present, that defender candidate has the same return type as $m()$ in C .

Note that the typechecking rules do not reject classes which inherit abstract method definitions from supertypes and which do not provide an implementation. These classes are considered valid, but at runtime are identified as having no linkage. This is a consequence of simplifying the language model to not explicitly track classes as abstract vs concrete, and does not interfere with the goals of this model.

8 Method linkage

We are now able to define the linkage of $m()$ in a class C which may inherit its implementation from a superclass or from a default method in an interface.

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
D \text{ OK} \qquad \qquad \qquad \forall_i I_i \text{ OK} \\
\text{T-CLASSNONE} \frac{S = \text{mtype}(C) \quad \text{mprov}(C) = \{ \} \quad \text{icand}(C) = \{ \}}{C \text{ OK}}
\end{array}$$

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
D \text{ OK} \qquad \qquad \qquad \forall_i I_i \text{ OK} \\
\text{T-CLASSFROMSUPER} \frac{\text{mprov}(C) = \{ E \} \quad \text{mtype}(C) = \text{mtype}(E)}{C \text{ OK}}
\end{array}$$

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
D \text{ OK} \qquad \qquad \qquad \forall_i I_i \text{ OK} \\
\text{mprov}(C) = \{ \} \qquad \qquad \text{icand}(C) = \{ J \} \\
\text{T-CLASSFROMDEFAULT} \frac{\text{mtype}(C) = \text{mtype}(J)}{C \text{ OK}}
\end{array}$$

Figure 6: Typechecking of classes

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
\text{mprov}(C) = \{ T \} \qquad \qquad T \text{ HasBody} \\
\text{R-LINKIMPL} \frac{}{\text{mres}(C) = T}
\end{array}$$

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
\text{mprov}(C) = \{ \} \quad \text{icand}(C) = \{ J \} \quad J \text{ HasBody} \\
\text{R-LINKDEF} \frac{}{\text{mres}(C) = J}
\end{array}$$

Figure 7: Linkage

Figure 7 defines the lookup function $\text{mres}(C)$, which indicates the linkage of $\text{m}()$ in class C . Its value is the type from which $\text{m}()$ is inherited.

9 Superclass invocation linkage

Java provides a mechanism for a method overriding a method in a superclass to invoke the overridden method, using the `super.m()` syntax. We would wish to support a similar mechanism, where an interface or class overriding a default method from an interface I could invoke the overridden default as `I.super.m()`.

While the FD calculus does not provide syntax for method invocation, we can define a lookup function $\text{sres}(T, I)$ that would provide the desired linkage for a super-call to $\text{m}()$ from T through interface I , just as we did for $\text{mres}(C)$. Figure 8 shows the calculation of $\text{sres}(T, I)$. For making a super-call from T through interface I , I must be an immediate super-interface of T , I must have

$$\begin{array}{c}
\text{class } C \text{ extends } D \text{ implements } I_1, \dots, I_n \{ \dots \} \\
\text{R-CLASSSUPER} \frac{\text{icand}(I_i) = \{ J \} \quad J \in \text{icand}(C) \quad J \text{ HasBody}}{sres(C, I_i) = J}
\end{array}$$

$$\begin{array}{c}
\text{interface } I \text{ extends } I_1, \dots, I_n \{ \dots \} \\
S = \bigcup_k \text{icand}(I_k) \\
\text{R-INTFSUPER} \frac{\text{icand}(I_i) = \{ J \} \quad J \in \text{prune}(S) \quad J \text{ HasBody}}{sres(I, I_i) = J}
\end{array}$$

Figure 8: Super-call linkage

a well-defined default on its own, and that default must not be overridden by a more specific superinterface of T . (In other words, $I.super.m()$ must link to a method that was a candidate for inheritance by T if T were to not override $m()$.)