

Type-Safe Feature-Oriented Product Lines

Sven Apel[†], Christian Kästner[‡], Armin Größlinger[†], and
Christian Lengauer[†]

[†] Department of Informatics and Mathematics, University of Passau
{apel, groesslinger, lengauer}@uni-passau.de

[‡] School of Computer Science, University of Magdeburg
kaestner@iti.cs.uni-magdeburg.de



Technical Report, Number MIP-0909
Department of Informatics and Mathematics
University of Passau, Germany
June 2009

Type-Safe Feature-Oriented Product Lines

Sven Apel[†], Christian Kästner[‡], Armin Größlinger[†], and Christian Lengauer[†]

[†] Department of Informatics and Mathematics, University of Passau
{apel, groesslinger, lengauer}@uni-passau.de

[‡] School of Computer Science, University of Magdeburg
kaestner@iti.cs.uni-magdeburg.de

Abstract. A *feature-oriented product line* is a family of programs that share a common set of features. A *feature* implements a stakeholder’s requirement, represents a design decision and configuration option and, when added to a program, involves the introduction of new structures, such as classes and methods, and the refinement of existing ones, such as extending methods. With feature-oriented decomposition, programs can be generated, solely on the basis of a user’s selection of features, by the composition of the corresponding feature code. A key challenge of feature-oriented product line engineering is how to guarantee the correctness of an *entire* feature-oriented product line, i.e., of all of the member programs generated from different combinations of features. As the number of valid feature combinations grows progressively with the number of features, it is not feasible to check all individual programs. The only feasible approach is to have a type system check the entire code base of the feature-oriented product line. We have developed such a type system on the basis of a formal model of a feature-oriented Java-like language. We demonstrate that the type system ensures that *every* valid program of a feature-oriented product line is well-typed and that the type system is complete.

1 Introduction

Feature-oriented programming (FOP) aims at the modularization of programs in terms of features [55,13]. A *feature* implements a stakeholder’s requirement and is typically an increment in program functionality [55,13]. Contemporary feature-oriented programming languages and tools such as AHEAD [13], Xak [2], CaesarJ [48], Classbox/J [14], FeatureHouse [9], and FeatureC++ [10] provide a variety of mechanisms that support the specification, modularization, and composition of features. A key idea is that a feature is implemented by a distinct code unit, called a *feature module*. When added to a base program, it introduces new structures, such as classes and methods, and refines existing ones, such as extending methods [43,11]. A program that is decomposed into features is called henceforth a *feature-oriented program*.¹

¹ Typically, feature-oriented decomposition is orthogonal to class-based or functional decomposition [63,49,61]. A multitude of modularization and composition mechanisms [17,26,24,45,46,56,65] have been developed in order to allow programmers to decompose a program along multiple dimensions [63]. Feature-oriented languages and tools provide a significant subset of these mechanisms [11].

Beside the decomposition of programs into features, the concept of a feature is useful for distinguishing different, related programs thus forming a *software product line* [35,21]. Typically, programs of a common domain share a set of features but also differ in other features. For example, suppose an email client for mobile devices that supports the protocols IMAP and POP3 and another client that supports POP3, MIME, and SSL encryption. With a decomposition of the two programs into the features IMAP, POP3, MIME, and SSL, both programs can share the code of the feature POP3. Since mobile devices have only limited resources, unnecessary features should be removed.

With feature-oriented decomposition, programs can be generated solely on the basis of a user's selection of features by the composition of the corresponding feature modules. Of course, not all combinations of features are legal and result in correct programs [12]. A *feature model* describes which features can be composed in which combinations, i.e., which programs are *valid* [35,21]. It consists of an (ordered) set of features and a set of constraints on feature combinations [21,12]. For example, our email client may have different rendering engines for HTML text, e.g., the Mozilla engine or the Safari engine, but only one at a time. A set of feature modules along with a feature model is called a *feature-oriented product line* [12].

An important question is how the correctness of feature-oriented programs, in particular, and product lines, in general, can be guaranteed. A first problem is that contemporary feature-oriented languages and tools usually involve a code generation step during composition in which the code is transformed into a lower-level representation. In previous work, we have addressed this problem by modeling feature-oriented mechanisms directly in the formal syntax and semantics of a core language, called *Feature Featherweight Java (FFJ)*. The type system of FFJ ensures that the composition of feature modules is type-safe [8].

In this paper, we address a second problem: How can the correctness of an *entire* feature-oriented product line be guaranteed? A naive approach would be to type-check all valid programs of a product line using a type checker like the one of FFJ [8]. However, this approach does not scale; already for 34 implemented optional features, a variant can be generated for every person on the planet. Noticing this problem, Czarnecki and Pietroszek [22] and Thaker et al. [64] suggested the development of a type system that checks the entire code base of the feature-oriented product line, instead of all individual feature-oriented programs. In this scenario, a type checker must analyze *all* feature modules of a product line on the basis of the feature model. We will show that, with this information, the type checker can ensure that *every* valid program variant that can be generated is type-safe. Specifically, we make the following contributions:

- We provide a condensed version of FFJ, which is in many respects more elegant and concise than its predecessor [8].
- We develop a formal type system that uses information about features and constraints on feature combinations in order to type-check a product line without generating every program.
- We prove *correctness* by proving that every program generated from a well-formed product line is well-formed, as long as the feature selection satisfies the constraints of the product line. Furthermore, we prove *completeness* by proving that the well-

typedness of all programs of a product line guarantees that the product line is well-typed as a whole.

- We offer an implementation of FFJ, including the proposed type system, which can be downloaded for evaluation and for experiments with further feature-oriented language and typing mechanisms.

Our work differs in many respects from previous and related work (see Section 5 for a comprehensive discussion). Most notably, Thaker et al. have implemented a type system for feature-oriented product lines and conducted several case studies [64]. We take their work further with a formalization and a correctness and completeness proof.

Furthermore, our work differs in many respects from previous work on modeling and type-checking feature-oriented and related programming mechanisms. Most notably, we model the feature-related mechanisms directly in FFJ's syntax and semantics, without any transformation to a lower-level representation, and we stay very close to the syntax of contemporary feature-oriented languages and tools (see Section 5). We begin with a brief introduction to FFJ.

2 Feature-Oriented Programs in FFJ

In this section, we introduce the language FFJ. Originally, FFJ was designed for feature-oriented programs [8,7]. We extend FFJ in Section 3 to support feature-oriented product lines, i.e., to support the representation of multiple alternative program variants at a time.

2.1 An Overview of FFJ

FFJ is a lightweight feature-oriented language that has been inspired by *Featherweight Java (FJ)* [32]. As with FJ, we have aimed at minimality in the design of FFJ. FFJ provides basic constructs like classes, fields, methods, and inheritance and only a few new constructs capturing the core mechanisms of feature-oriented programming. But, so far, FFJ's type system has not supported the development of feature-oriented product lines. That is, the feature modules written in FFJ are interpreted as a single program. We will change this in Section 3.

An FFJ program consists of a set of classes and refinements. A *refinement* extends a class that has been introduced previously. Each class and refinement is associated with a feature. We say that a feature *introduces* a class or *applies* a refinement to a class. Technically, the mapping between classes/refinements and the features they belong to can be established in different ways, e.g., by extending the language with modules representing features [48,14,23] or by grouping classes and refinements that belong to a feature in packages or directories [13,10].

Like in FJ, each class declares a superclass, which may be the class `Object`. Refinements are defined using the keyword `refines`. The semantics of a refinement applied to a class is that the refinement's members are added to and merged with the member of the refined class. This way, a refinement can *add* new fields and methods to the class and *override* existing methods (declared by `overrides`).

On the left side in Figure 1, we show an excerpt of the code of a basic email client, called EMAILCLIENT, (top) and a feature, called SSL, (bottom) in FFJ. The feature SSL adds the class SSL (Lines 7–10) to the email client’s code base and refines the class Trans in order to encrypt outgoing messages (Lines 11–15). To this effect, the refinement of Trans adds a new field `key` (Line 12) and overrides the method `send` of class Trans (Lines 13-15).




Fig. 1. A feature-oriented email client supporting SSL encryption.

Typically, a programmer applies multiple refinements to a class by composing a sequence of features. This is called a *refinement chain*. A refinement that is applied immediately before another refinement in the chain is called its *predecessor*. The order of the refinements in a refinement chain is determined by their composition order. On the right side in Figure 1, we depict the refinement and inheritance relationships of our email example.

Fields are unique within the scope of a class and its inheritance hierarchy and refinement chain. That is, a refinement or subclass is not allowed to add a field that has already been defined in a predecessor in the refinement chain or in a superclass. For example, a further refinement of Trans would not be allowed to add a field `key`, since `key` has been introduced by a refinement of feature SSL already. With methods, this is different. A refinement or subclass may add new methods (overloading is prohibited) and override existing methods. In order to distinguish the two cases, FFJ expects the programmer to declare whether a method overrides an existing method (using the modifier `overrides`). For example, the refinement of Trans in feature SSL overrides the method `send` introduced by feature MAIL; for subclasses, this is similar.

The distinction between method introduction and overriding allows the type system to check (1) whether an introduced method inadvertently replaces or occludes an existing method with the same name and (2) whether, for every overriding method, there is a proper method to be overridden. Apart from the modifier `overrides`, a method in FFJ is similar to a method in FJ. That is, a method body is an expression (prefixed with

return) and not a sequence of statements. This is due to the functional nature of FFJ and FJ. Furthermore, overloading of methods (introducing methods with equal names and different argument types) is not allowed in FFJ (and FJ).

As shown in Figure 1, refinement chains grow from left to right and inheritance hierarchies from top to bottom. When looking up a method body, FFJ traverses the combined inheritance and refinement hierarchy of an object and selects the right-most and bottom-most body of a method declaration or method refinement that is compatible. This kind of lookup is necessary since we model features *directly* in FFJ, instead of generating and evaluating FJ code [40]. First, the FFJ calculus looks for a method declaration in the refinement chain of the object's class, starting with the last refinement back to the class declaration itself. The first body of a matching method declaration is returned. If the method is not found in the class' refinement chain or in its own declaration, the methods in the superclass (and then the superclass' superclass, etc.) are searched, each again from the most specific refinement of the class declaration itself. The field lookup works similarly, except that the entire inheritance and refinement hierarchy is searched and the fields are accumulated in a list. In Figure 2, we illustrate the processes of method body and field lookup schematically.




Fig. 2. Order of method body and field lookup in FFJ.

2.2 Syntax of FFJ

Before we go into detail, let us explain some notational conventions. We abbreviate lists in the obvious ways:

- \overline{C} is shorthand for C_1, \dots, C_n
- $\overline{C} f$ is shorthand for $C_1 f_1, \dots, C_n f_n$
- $\overline{C} f;$ is shorthand for $C_1 f_1; \dots; C_n f_n;$
- $\overline{t} : \overline{C}$ is shorthand for $t_1 : C_1, \dots, t_n : C_n$
- $\overline{C} <: \overline{D}$ is shorthand for $C_1 <: D_1 \dots C_n <: D_n$
- ...

Note that, depending on the context, blanks, commas, or semicolons separate the elements of a list. The context will make clear which separator is meant. The symbol •

denotes the empty list and lists of field declarations, method declarations, and parameter names must not contain duplicates. We use the metavariables A – E for class names, f – h for field names, and m for method names. Feature names are denoted by Greek letters.

In Figure 3, we depict the syntax of FFJ in extended Backus-Naur-Form. An FFJ program consists of a set of class and refinement declarations. A class declaration L declares a class with the name C that inherits from a superclass D and consists of a list $\overline{C} f$; of fields and a list \overline{M} of method declarations.² A refinement declaration R consists of a list $\overline{C} f$; of fields and a list \overline{M} of method declarations.

$L ::=$ <i>class declarations:</i> class C extends D { $\overline{C} f$; \overline{M} }	$t ::=$ <i>terms:</i> x <i>variable</i> $t.f$ <i>field access</i> $t.m(\overline{t})$ <i>method invocation</i> new $C(\overline{t})$ <i>object creation</i> $(C)t$ <i>cast</i>
$R ::=$ <i>refinement declarations:</i> refines class C { $\overline{C} f$; \overline{M} }	$v ::=$ <i>values:</i> new $C(\overline{v})$ <i>object creation</i>
$M ::=$ <i>method declarations:</i> [overrides] C $m(\overline{C} x)$ { return t ; }	

Fig. 3. Syntax of FFJ in extended BNF.

A method m expects a list $\overline{C} x$ of arguments and declares a body that returns only a single expression t of type C . Using the modifier *overrides*, a method declares that it intends to override another method with the same name and signature. Where we want to distinguish methods that override others and methods that do not override others, we call the former *method introductions* and the latter *method refinements*.

Finally, there are five forms of terms: the variable, field access, method invocation, object creation, and type cast, which are taken from FJ without change. The only values are object creations whose arguments are values as well.

2.3 FFJ’s Class Table

Declarations of classes and refinements can be looked up via a class table CT . The compiler fills the class table during the parser pass. In contrast to FJ, class and refinement declarations are identified not only by their names but, additionally, by the names of the enclosing features. For example, in order to retrieve the declaration of class **Trans**, introduced by feature **MAIL**, in our example of Figure 1, we write $CT(\text{MAIL.Trans})$; in order to retrieve the refinement of class **Trans** applied by feature **SSL**, we write $CT(\text{SSL.Trans})$. We call $\Phi.C$ the *qualified type* of class C in feature Φ . In FFJ, class and refinement declarations are unique with respect to their qualified types. This property is ensured because of the following sanity conditions: a feature is not allowed

² The concept of a class constructor is unnecessary in FFJ and FJ [54]. Its omission simplifies the syntax, semantics, and type rules significantly without loss of generality.

- to introduce a class or refinement twice inside a single feature module and
- to refine a class that the feature has just introduced.

These are common sanity conditions in feature-oriented languages and tools [13,10,9].

As for FJ, we impose further sanity conditions on the class table and the inheritance relation:

- $CT(\Phi.C) = \text{class } C \dots \text{ or refines class } C \dots$ for every qualified type $\Phi.C \in \text{dom}(CT)$; Feature `Base` plays the same role for features as `Object` plays for classes; it is a symbol denoting the empty feature at which lookups terminate.
- $\text{Base.Object} \notin \text{dom}(CT)$;
- for every class name C appearing anywhere in CT , we have $\Phi.C \in \text{dom}(CT)$ for at least one feature Φ ; and
- the inheritance relation contains no cycles (incl. self-cycles).

2.4 Refinement in FFJ

Information about the refinement chain of a class can be retrieved using the refinement table RT . The compiler fills the refinement table during the parser pass. $RT(C)$ yields a list of all features that either introduce or refine class C . The leftmost element of the result list is the feature that introduces the class C and, then, from left to right, the features are listed that refine class C in the order of their composition. In our example of Figure 1, $RT(\text{Trans})$ yields the list `EMAILCLIENT, SSL`. There is only a single sanity condition for the refinement table:

- $RT(C) = \bar{\Phi}$ for every type $C \in \text{dom}(CT)$, with $\bar{\Phi}$ being the features that introduce and refine class C .

In Figure 4, we show two functions for the navigation of the refinement chain that rely on RT . Function *last* returns, for a class name C , a qualified type $\Psi_n.C$, in which Ψ_n refers to the feature that applies the final refinement to class C ; if a class is not refined at all, Ψ_n refers to the feature that introduces class C . Function *pred* returns, for a qualified type $\Phi.C$, another qualified type $\Psi_n.C$, in which Ψ_n refers to the feature that introduces or refines class C and that is the immediate predecessor of Φ in the refinement chain; if there is no predecessor, `Base.Object` is returned.

Navigating along the refinement chain

$$\frac{RT(C) = \bar{\Psi}}{\text{last}(C) = \Psi_n.C} \quad \frac{RT(C) = \bar{\Psi}, \Phi, \bar{\Omega} \quad \bar{\Psi} \neq \bullet}{\text{pred}(\Phi.C) = \Psi_n.C} \quad \frac{RT(C) = \Phi, \bar{\Omega}}{\text{pred}(\Phi.C) = \text{Base.Object}}$$

Fig. 4. Refinement in FFJ.

2.5 Subtyping in FFJ

In Figure 5, we show the subtype relation of FFJ. The subtype relation $<:$ is defined by one rule each for reflexivity and transitivity and one rule for relating the type of a class to the type of its immediate superclass. It is not necessary to define subtyping over qualified types because only classes (not refinements) declare superclasses and there is only a single declaration per class.

$$\begin{array}{c}
 \textit{Subtyping} \qquad \boxed{C <: D} \\
 C <: C \qquad \frac{C <: D \quad D <: E}{C <: E} \qquad \frac{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \dots \}}{C <: D}
 \end{array}$$

Fig. 5. Subtyping in FFJ.

2.6 Auxiliary Definitions of FFJ

In Figure 6, we show the auxiliary definitions of FFJ. Function *fields* searches the refinement chain from right to left and accumulates the fields into a list (using the comma as concatenation operator). If there is no further predecessor in the refinement chain, i.e., we have reached a class declaration, then the refinement chain of the superclass is searched (see Figure 2). If `Base.Object` is reached, the empty list is returned (denoted by \bullet).

Function *mbody* looks up the most specific and most refined body of a method m . A body consists of the formal parameters \bar{x} of a method and the actual term t representing the content. The search is like in *fields*. First, the refinement chain is searched from right to left and, then, the superclasses' refinement chains are searched, as illustrated in Figure 2. Note that `[overrides]` means that a given method declaration may (or may not) have the modifier. This way, we are able to define uniform rules for method introduction and method refinement. Function *mtype* yields the signature $\bar{B} \rightarrow B_0$ of a declaration of method m . The lookup is like in *mbody*.

Predicate *introduce* is used to check whether a class has been introduced by multiple features and whether a field or method has been introduced multiple times in a class. Precisely, it states, in the case of classes, whether C has not been introduced by any feature other than Φ and whether a method m or a field f has not been introduced by $\Phi.C$ or in any of its predecessors or superclasses. To evaluate it, we check, in the case of classes, whether $CT(\Psi.C)$ yields a class declaration or not, for any feature Ψ different from Φ , in the case of methods, whether *mtype* yields a signature or not and, in the case of fields, whether f is defined in the list of fields returned by *fields*.

Predicate *refine* states whether, for a given refinement, a proper class has been declared previously in the refinement chain. The predicate *override* states whether a method m has been introduced before in some predecessor of $\Phi.C$ and whether the previous declaration of m has the given signature.

Field lookup

$$\boxed{fields(\Phi.C) = \overline{C} f}$$

$$fields(\text{Base.Object}) = \bullet$$

$$\frac{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} f; \overline{M} \}}{fields(\Phi.C) = fields(\text{last}(D)), \overline{C} f} \quad \frac{CT(\Phi.C) = \text{refines class } C \{ \overline{C} f; \overline{M} \}}{fields(\Phi.C) = fields(\text{pred}(\Phi.C)), \overline{C} f}$$

Method body lookup

$$\boxed{mbody(m, \Phi.C) = (\overline{x}, t)}$$

$$\frac{[\text{overrides}] B \ m(\overline{B} \ \overline{x}) \{ \text{return } t; \} \in \overline{M}}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} f; \overline{M} \}} \quad \frac{m \text{ is not defined in } \overline{M}}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} f; \overline{M} \}} \\ \frac{}{mbody(m, \Phi.C) = (\overline{x}, t)} \quad \frac{}{mbody(m, \Phi.C) = mbody(m, \text{last}(D))}$$

$$\frac{[\text{overrides}] B \ m(\overline{B} \ \overline{x}) \{ \text{return } t; \} \in \overline{M}}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} f; \overline{M} \}} \quad \frac{m \text{ is not defined in } \overline{M}}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} f; \overline{M} \}} \\ \frac{}{mbody(m, \Phi.C) = (\overline{x}, t)} \quad \frac{}{mbody(m, \Phi.C) = mbody(m, \text{pred}(\Phi.C))}$$

Method type lookup

$$\boxed{mtype(m, \Phi.C) = \overline{C} \rightarrow C}$$

$$\frac{B_0 \ m(\overline{B} \ \overline{x}) \{ \text{return } t; \} \in \overline{M}}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} f; \overline{M} \}} \quad \frac{m \text{ is not defined in } \overline{M}}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} f; \overline{M} \}} \\ \frac{}{mtype(m, \Phi.C) = \overline{B} \rightarrow B_0} \quad \frac{}{mtype(m, \Phi.C) = mtype(m, \text{last}(D))}$$

$$\frac{B_0 \ m(\overline{B} \ \overline{x}) \{ \text{return } t; \} \in \overline{M}}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} f; \overline{M} \}} \quad \frac{m \text{ is not defined in } \overline{M}}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} f; \overline{M} \}} \\ \frac{}{mtype(m, \Phi.C) = \overline{B} \rightarrow B_0} \quad \frac{}{mtype(m, \Phi.C) = mtype(m, \text{pred}(\Phi.C))}$$

Valid class introduction

$$\boxed{\text{introduce}(\Phi.C)}$$

$$\frac{\nexists \Psi : (CT(\Psi.C) = \text{class } C \dots \wedge \Phi \neq \Psi)}{\text{introduce}(\Phi.C)}$$

Valid field introduction

$$\boxed{\text{introduce}(f, \Phi.C)}$$

$$\frac{fields(\Phi.C) = \overline{E} \ \overline{h} \quad f \notin \overline{h}}{\text{introduce}(f, \Phi.C)}$$

Valid method introduction

$$\boxed{\text{introduce}(m, \Phi.C)}$$

$$\frac{(m, \Phi.C) \notin \text{dom}(mtype)}{\text{introduce}(m, \Phi.C)}$$

Valid class refinement

$$\boxed{\text{refine}(\Phi.C)}$$

$$\frac{RT(C) = \overline{\Psi}, \Phi, \overline{\Omega} \quad CT(\Psi_1.C) = \text{class } C \dots}{\text{refine}(\Phi.C)}$$

Valid method overriding

$$\boxed{\text{override}(m, \Phi.C, \overline{C} \rightarrow C_0)}$$

$$\frac{mtype(m, \Phi.C) = \overline{B} \rightarrow B_0 \quad \overline{C} = \overline{B} \quad C_0 = B_0}{\text{override}(m, \Phi.C, \overline{C} \rightarrow C_0)}$$

Fig. 6. Auxiliary definitions of FFIJ.

2.7 Evaluation of FFJ Programs

Each FFJ program consists of a class table and a term.³ The term is evaluated using the evaluation rules shown in Figure 7. The evaluation terminates when a value, i.e., a term of the form $\text{new } C(\bar{v})$, is reached. Note that we use a *direct semantics* of class refinement [40]. That is, the field and method lookup mechanisms incorporate all refinements when a class is searched for fields and methods. An alternative, which is discussed in Section 5, would be a *flattening semantics*, i.e., to merge a class in a preprocessing step with all of its refinements into a single declaration.

$$\frac{\text{fields}(\text{last}(C)) = \bar{C} \bar{f}}{(\text{new } C(\bar{v})) . f_i \longrightarrow v_i} \quad (\text{E-PROJNEW})$$

$$\frac{\text{mbody}(m, \text{last}(C)) = (\bar{x}, t_0)}{(\text{new } C(\bar{v})) . m(\bar{u}) \longrightarrow [\bar{x} \mapsto \bar{u}, \text{this} \mapsto \text{new } C(\bar{v})] t_0} \quad (\text{E-INVKNW})$$

$$\frac{C <: D}{(D)(\text{new } C(\bar{v})) \longrightarrow \text{new } C(\bar{v})} \quad (\text{E-CASTNEW})$$

$$\frac{t_0 \longrightarrow t'_0}{t_0 . f \longrightarrow t'_0 . f} \quad (\text{E-FIELD})$$

$$\frac{t_0 \longrightarrow t'_0}{t_0 . m(\bar{t}) \longrightarrow t'_0 . m(\bar{t})} \quad (\text{E-INVKRECV})$$

$$\frac{t_i \longrightarrow t'_i}{v_0 . m(\bar{v}, t_i, \bar{t}) \longrightarrow v_0 . m(\bar{v}, t'_i, \bar{t})} \quad (\text{E-INVKARG})$$

$$\frac{t_i \longrightarrow t'_i}{\text{new } C(\bar{v}, t_i, \bar{t}) \longrightarrow \text{new } C(\bar{v}, t'_i, \bar{t})} \quad (\text{E-NEWARG})$$

$$\frac{t_0 \longrightarrow t'_0}{(C)t_0 . f \longrightarrow (C)t'_0 . f} \quad (\text{E-CAST})$$

Fig. 7. Evaluation of FFJ programs.

Using the subtype relation $<:$ and the auxiliary functions *fields* and *mbody*, the evaluation of FFJ is fairly simple. The first three rules are most interesting (the remaining rules are just congruence rules). Rule E-PROJNEW describes the projection of a field from an instantiated class. A projected field f_i evaluates to a value v_i that has been passed as argument to the instantiation. Function *fields* is used to look up the fields of the given class. It receives $\text{last}(C)$ as argument since we want to search the entire refinement chain of class C from right to left (cf. Figure 2).

Rule E-PROJINVK evaluates a method invocation by replacing the invocation with the method's body. The formal parameters of the method are substituted in the body for

³ The refinement table is not relevant for evaluation.

the arguments of the invocation; the value on which the method is invoked is substituted for this. The function *mbody* is called with the last refinement of the class C in order to search the refinement chain from right to left and return the most specific method body (cf. Figure 2).

Rule E-CASTNEW evaluates an upcast by simply removing the cast. Of course, the premise must be that the cast is really an upcast and not a downcast or an incorrect cast.

2.8 Type Checking FFJ Programs

The type relation of FFJ consists of the type rules for terms and the well-formedness rules for classes, refinements, and methods, shown in Figures 8 and 9.

<i>Term typing</i>	$\Gamma \vdash t : C$
$\frac{x : C \in \Gamma}{\Gamma \vdash x : C}$	(T-VAR)
$\frac{\Gamma \vdash t_0 : C_0 \quad \text{fields}(\text{last}(C_0)) = \overline{C} \text{ f}}{\Gamma \vdash t_0.f_i : C_i}$	(T-FIELD)
$\frac{\Gamma \vdash t_0 : C_0 \quad \Gamma \vdash \overline{t} : \overline{C} \quad \text{mtype}(m, \text{last}(C_0)) = \overline{D} \rightarrow C \quad \overline{C} <: \overline{D}}{\Gamma \vdash t_0.m(\overline{t}) : C}$	(T-INVK)
$\frac{\Gamma \vdash \overline{t} : \overline{C} \quad \text{fields}(\text{last}(C)) = \overline{D} \text{ f} \quad \overline{C} <: \overline{D}}{\Gamma \vdash \text{new } C(\overline{t}) : C}$	(T-NEW)
$\frac{\Gamma \vdash t_0 : D \quad D <: C}{\Gamma \vdash (C)t_0 : C}$	(T-UCAST)
$\frac{\Gamma \vdash t_0 : D \quad C <: D \quad C \neq D}{\Gamma \vdash (C)t_0 : C}$	(T-DCAST)
$\frac{\Gamma \vdash t_0 : D \quad C \not<: D \quad D \not<: C \quad \text{stupid warning}}{\Gamma \vdash (C)t_0 : C}$	(T-SCAST)

Fig. 8. Term typing in FFJ.

Term Typing Rules. A term typing judgment is a triple consisting of a typing context Γ , a term t , and a type C (see Figure 8).

Rule T-VAR checks whether a free variable is contained in the typing context. Rule T-FIELD checks whether a field access $t_0.f$ is well-typed. Specifically, it checks whether f is declared in the type of t_0 and whether the type f equals the type of the entire term. Rule T-INVK checks whether a method invocation $t_0.m(\overline{t})$ is well-typed. To this end, it checks whether the arguments \overline{t} of the invocation are subtypes of the types of the

Method typing

$\boxed{M \text{ OK} \dashv \Phi.C}$

$$\frac{\overline{x : \overline{B}}, \text{this} : C \vdash t_0 : E_0 \quad E_0 <: B_0}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} \overline{f}; \overline{M} \} \quad \text{introduce}(m, \text{last}(D))} B_0 m(\overline{B} \overline{x}) \{ \text{return } t_0; \} \text{OK} \dashv \Phi.C$$

$$\frac{\overline{x : \overline{B}}, \text{this} : C \vdash t_0 : E_0 \quad E_0 <: B_0}{CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C} \overline{f}; \overline{M} \} \quad \text{override}(m, \text{last}(D), \overline{B} \rightarrow B_0)} \text{overrides } B_0 m(\overline{B} \overline{x}) \{ \text{return } t_0; \} \text{OK} \dashv \Phi.C$$

$$\frac{\overline{x : \overline{B}}, \text{this} : C \vdash t_0 : E_0 \quad E_0 <: B_0}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} \overline{f}; \overline{M} \} \quad \text{introduce}(m, \text{pred}(\Phi.C))} B_0 m(\overline{B} \overline{x}) \{ \text{return } t_0; \} \text{OK} \dashv \Phi.C$$

$$\frac{\overline{x : \overline{B}}, \text{this} : C \vdash t_0 : E_0 \quad E_0 <: B_0}{CT(\Phi.C) = \text{refines class } C \{ \overline{C} \overline{f}; \overline{M} \} \quad \text{override}(m, \text{pred}(\Phi.C), \overline{B} \rightarrow B_0)} \text{overrides } B_0 m(\overline{B} \overline{x}) \{ \text{return } t_0; \} \text{OK} \dashv \Phi.C$$

Class typing

$\boxed{L \text{ OK} \dashv \Phi}$

$$\frac{\forall f \in \overline{f} : \text{introduce}(f, \text{last}(D)) \quad \text{introduce}(\Phi.C) \quad \overline{M} \text{ OK} \dashv \Phi.C}{\text{class } C \text{ extends } D \{ \overline{C} \overline{f}; \overline{M} \} \text{OK} \dashv \Phi}$$

Refinement typing

$\boxed{R \text{ OK} \dashv \Phi}$

$$\frac{\forall f \in \overline{f} : \text{introduce}(f, \text{pred}(\Phi.C)) \quad \text{refine}(\Phi.C) \quad \overline{M} \text{ OK} \dashv \Phi.C}{\text{refines class } C \{ \overline{C} \overline{f}; \overline{M} \} \text{OK} \dashv \Phi}$$

Fig. 9. Well-formedness rules of FFJ.

formal parameters of m and whether the return type of m equals the type of the entire term. Rule T-NEW checks whether an object creation $\text{new } C(\bar{t})$ is well-typed in that it checks whether the arguments \bar{t} of the instantiation of C are subtypes of the types \bar{D} of the fields of C and whether C equals the type of the entire term. The rules T-UCAST, T-DCAST, and T-SCAST check whether casts are well-typed. In each rule, it is checked whether the type C the term t_0 is cast to is a subtype, supertype, or unrelated type of the type of t_0 and whether C equals the type of the entire term.⁴

Well-Formedness Rules. In Figure 9, we show FFJ’s well-formedness rules of classes, refinements, and methods.

The typing judgments of classes and refinements are binary relations between a class or refinement declaration and a feature, written $L \text{ OK} \dashv \Phi$ and $R \text{ OK} \dashv \Phi$. The rule of classes checks whether all methods are well-formed in the context of the class’ qualified type. Moreover, it checks whether none of the fields of the class declaration is introduced multiple times in the combined inheritance and refinement hierarchy and whether there is no feature other than Φ that introduces a class C (using *introduce*). The well-formedness rule of refinements is analogous, except that the rule checks whether a corresponding class has been introduced before (using *refine*).

The typing judgment of methods is a binary relation between a method declaration and the qualified type that declares the method, written $M \text{ OK} \dashv \Phi.C$. There are four different rules for methods (from top to bottom in Figure 9)

1. that do not override another method and that are declared by classes,
2. that override another method and that are declared by classes,
3. that do not override another method and that are declared by refinements,
4. that override another method and that are declared by refinements.

All four rules check whether the type E_0 of the method body is a subtype of the declared return type B_0 of the method declaration. For methods that are being introduced, it is checked whether no method with an identical name has been introduced in a superclass (Rule 1) or in a predecessor in the refinement chain (Rule 3). For methods that override other methods, it is checked whether a method with identical name and signature exists in the superclass (Rule 2) or in a predecessor in the refinement chain (Rule 4).

Well-Typed FFJ Programs. Finally, an FFJ program, consisting of a term, a class table, and a refinement table, is well-typed if

- the term is well-typed (checked using FFJ’s term typing rules),
- all classes and refinements stored in the class table are well-typed (checked using FFJ’s well-formedness rules), and
- the class and refinement tables are well-formed (ensured by the corresponding sanity conditions).

⁴ Rule T-SCAST is needed only for the small step semantics of FFJ (and FJ) in order to be able to formulate and prove the type preservation property. FFJ (and FJ) programs whose type derivation contains this rule (i.e., the premise *stupid warning* appears in the derivation) are not further considered (cf. [32]).

Type Soundness of FFJ. The type system of FFJ is sound. We can prove this using the standard theorems of preservation and progress [66]:

THEOREM 2.1 (Preservation) If $\Gamma \vdash t : C$ and $t \longrightarrow t'$, then $\Gamma \vdash t' : C'$ for some $C' < C$.

THEOREM 2.2 (Progress) Suppose t is a well-typed term.

1. If t includes $\text{new } C_0(\bar{t}).f_i$ as a subterm, then $\text{fields}(\text{last}(C_0)) = \overline{C} \bar{f}$ for some \overline{C} and \bar{f} .
2. If t includes $\text{new } C_0(\bar{t}).m(\bar{u})$ as a subterm, then $\text{mbody}(m, \text{last}(C_0)) = (\bar{x}, t_0)$ and $|\bar{x}| = |\bar{u}|$ for some \bar{x} and t_0 .

We provide the proofs of the two theorems in Appendix A.

3 Feature-Oriented Product Lines in FFJ_{PL}

In this section, our goal is to define a type system for feature-oriented product lines – a type system that checks whether all valid combinations of features yield well-typed programs. In this scenario, the features in question may be optional or mutually exclusive so that different combinations are possible that form different feature-oriented programs. Since there may be plenty of valid combinations, type checking all of them individually is usually not feasible.

In order to provide a type system for feature-oriented product lines, we need information about which combinations of features are valid, i.e., which features are mandatory, optional, or mutually exclusive, and we need to adapt the subtype and type rules of FFJ to check that there are no combinations/variants that lead to ill-typed terms. The type system guarantees that every program derived from a well-typed product line is a well-typed FFJ program. FFJ together with the type system for checking feature-oriented product lines is henceforth called FFJ_{PL}.

3.1 An Overview of Feature-Oriented Product Lines

A feature-oriented product line is made up of a set of feature modules and a feature model. The feature modules contains the features' implementation and the feature model describes how the feature modules can be combined. In contrast to the feature-oriented programs of Section 2, typically, some features are optional and some are mutually exclusive (Also other relations such as disjunction, negation, and implication are possible [12]; they are broken down to mandatory, optional, and mutually exclusive features, as we will explain.). Generally, in a *derivation step*, a user selects a valid subset of features from which, subsequently, a feature-oriented program is derived. In our case, derivation means assembling the corresponding feature modules for a given set of features. In Figure 10, we illustrate the process of *program derivation*.

Typically, a wide variety of programs can be derived from a product line [21,19]. The challenge is to define a type system that guarantees, on the basis of the feature modules and the feature model, that all valid programs are well-typed. Once a program is derived from such a product line, we can be sure that it is well-typed and we can evaluate it using the standard evaluation rules of FFJ (see Section 2.7).




Fig. 10. The process of deriving programs from a product line.

3.2 Managing Variability – Feature Models

The aim of developing a product line is to manage the *variability* of a set of programs developed for a particular domain and to facilitate the *reuse* of feature implementations among the programs of the domain. A *feature model* captures the variability by (explicitly or implicitly) defining an ordered set of all features of a product line and their legal feature combinations. A well-defined feature order is essential for field and method lookup (see Section 3.6).

Different approaches to product line engineering use different representations of feature models to define legal feature combinations. The simplest approach is to enumerate all legal feature combinations. In practice, commonly different flavors of tree structures are used, sometimes in combination with additional propositional constraints, to define legal combinations [21,12], as illustrated in Figure 10.

For our purpose, the actual representation of legal feature combinations is not relevant. In FFJ_{PL} , we use the feature model only to check whether feature and/or specific program elements are present in certain circumstances. A design decision of FFJ_{PL} is to abstract from the concrete representation of the underlying feature model and rather to provide an interface to the feature model. This has two benefits: (1) we do not need to struggle with all the details of the formalization of feature models, which is well understood by researchers [12,22,64,23] and outside the scope of this paper, and (2) we are able to support different kinds of feature model representations, e.g., a tree structures, grammars, or propositional formulas [12]. The interface to the feature model is simply a set of functions and predicates that we use to ask questions like “may (or may not) feature A be present together with feature B” or “is program element m present in every variant in which also feature A is present”, i.e., “is program element m always *reachable* from feature A”.

3.3 Challenges of Type Checking

Let us explain the challenges of type checking by extending our email example, as shown in Figure 11. Suppose our basic email client is refined to process incoming text messages (feature TEXT, Lines 1–8). Optionally, it is enabled to process HTML

messages, using either Mozilla’s rendering engine (feature MOZILLA, Lines 9–12) or Safari’s rendering engine (feature SAFARI, Lines 13–16). To this end, the features MOZILLA and SAFARI override the method `render` of class `Display` (Line 11 and 15) in order to invoke the respective rendering engines (field `renderer`, Lines 10 and 14) instead of the text printing function (Line 7).

```

                                                                    Feature TEXT
1  refines class Trans {
2    Unit receive(Msg msg) {
3      return /* do something... */ new Display().render(msg);
4    }
5  }
6  class Display {
7    Unit render(Msg msg) { /* display message in text format */
8  }
                                                                    Feature MOZILLA
9  refines class Display {
10   MozillaRenderer renderer;
11   overrides Unit render(Msg m) { /* render HTML message using the Mozilla engine */
12 }
                                                                    Feature SAFARI
13 refines class Display {
14   SafariRenderer renderer;
15   overrides Unit render(Msg m) { /* render HTML message using the Safari engine */
16 }

```

Fig. 11. A feature-oriented email client using Mozilla’s and Safari’s rendering engines.

The first thing to observe is that the features MOZILLA and SAFARI rely on class `Display` and its method `render` introduced by feature TEXT. In order to guarantee that every derived program is well-formed, the type system checks whether `Display` and `render` are *always reachable* from the features MOZILLA and SAFARI, i.e., whether, in every program variant that contains MOZILLA and SAFARI, also feature TEXT is present.

The second thing to observe is that the features MOZILLA and SAFARI both add a field `renderer` to `Display` (Lines 10 and 14), both of which have different types. In FFJ, a program with both feature modules would not be a well-typed program because the field `renderer` is introduced twice. However, Figure 11 is not intended to represent a single feature-oriented program but a feature-oriented product line; the features MOZILLA and SAFARI are mutually exclusive, as defined in the product line’s feature model (stated earlier), and the type system has to take this fact into account.

Let us summarize the key challenges of type checking product lines:

- A global class table contains classes and refinements of all features of a product line, even if some features are optional or mutually exclusive so that they are present only in *some* derived programs. That is, a single class can be introduced by multiple features as long as the features are mutually exclusive. This is also the case for multiple introductions of methods and fields, which may even have different types.

- The presence of types, fields, and methods depends on the presence of the features that introduce them. A reference from the elements of a feature to a type, a field projection, or a method invocation is valid if the referenced element is always reachable from the referring feature, i.e., in every variant that contains the referring feature.
- Like references, an extension of a program element, such as a class or method refinement, is valid only if the extended program element is always reachable from the feature that applies the refinement.
- Refinements of classes and methods do not necessarily form linear refinement chains. There may be alternative refinements of a single class or method that exclude one another, as explained below.

3.4 Collecting Information on Feature Modules

For type checking, the FFJ_{PL} compiler collects various information on the feature modules of the product line. Before the actual type checking is performed, the compiler fills three tables with information: the class table (CT), the introduction table (IT), and the refinement table (RT).

The class table CT of FFJ_{PL} is like the one of FFJ and has to satisfy the same sanity conditions except that (1) there may be multiple declarations of a class (or field or method), as long as they are defined in mutually exclusive features, and (2) there may be cycles in the inheritance hierarchy, but no cycles for each set of classes which are reachable from any given feature.

The introduction table IT maps a type to a list $\bar{\Phi}$ of (mutually exclusive) features that introduce the type. The features returned by IT are listed in the order prescribed by the feature model. In our example of Figure 11, a call of $IT(\text{Display})$ would return a list consisting only of the single feature TEXT . Likewise, the introduction table maps field and method names, in combination with their declaring classes, to features. For example, a call of $IT(\text{Display.renderer})$ would return the list MOZILLA, SAFARI . The sanity conditions for the introduction table are straightforward:

- $IT(\mathbf{C}) = \bar{\Phi}$ for every type $\mathbf{C} \in \text{dom}(CT)$, with $\bar{\Phi}$ being the features that introduce class \mathbf{C} .
- $IT(\mathbf{C}.f) = \bar{\Phi}$ for every field f contained in some class $\mathbf{C} \in \text{dom}(CT)$, with $\bar{\Phi}$ being the features that introduce field f .
- $IT(\mathbf{C}.m) = \bar{\Phi}$ for every method m contained in some class $\mathbf{C} \in \text{dom}(CT)$, with $\bar{\Phi}$ being the features that introduce method m .

Much like in FFJ , in FFJ_{PL} there is a refinement table RT . A call of $RT(\mathbf{C})$ yields a list of all features that either introduce *or* refine class \mathbf{C} , which is different from the introduction table that returns only the features that introduce class \mathbf{C} . As with IT , the features returned by RT are listed in the order prescribed by the feature model. The sanity condition for FFJ_{PL} 's refinement table is identical to the one of FFJ , namely:

- $RT(\mathbf{C}) = \bar{\Phi}$ for every type $\mathbf{C} \in \text{dom}(CT)$, with $\bar{\Phi}$ being the features that introduce and refine class \mathbf{C} .

3.5 Feature Model Interface

As said before, in FFJ_{PL} , we abstract from the concrete representation of the feature model and define instead an interface consisting of proper functions and predicates. There are two kinds of questions we want to ask about the feature model, which we explain next.

First, we would like to know which features are *never* present together, which features are *sometimes* present together, and which features are *always* present together. To this end, we define two predicates, *never* and *sometimes*, and a function *always*. Predicate $\text{never}(\bar{\Omega}, \Phi)$ indicates that feature Φ is never reachable in the context $\bar{\Omega}$, i.e., there is no valid program variant in which the features $\bar{\Omega}$ and feature Φ are present together. Predicate $\text{sometimes}(\bar{\Omega}, \Phi)$ indicates that feature Φ is sometimes present when the features $\bar{\Omega}$ are present, i.e., there are variants in which the features $\bar{\Omega}$ and feature Φ are present together and there are variants in which they are not present together. Function $\text{always}(\bar{\Omega}, \Phi)$ is used to evaluate whether feature Φ is always present in the context $\bar{\Omega}$ (either alone or within a group of alternative features). There are three cases: if feature Φ is always present in the context, *always* returns the feature again ($\text{always}(\bar{\Omega}, \Phi) = \Phi$); if feature Φ is not always present, but would be together with a certain group of mutually exclusive features $\bar{\Psi}$ (i.e., one of the group is always present), *always* returns all features of this group ($\text{always}(\bar{\Omega}, \Phi) = \Phi, \bar{\Psi}$). If a feature is not present at all, neither alone nor together with other mutually exclusive features, *always* returns the empty list ($\text{always}(\bar{\Omega}, \Phi) = \bullet$). The above predicates and function provide all information we need to know about the features' relationships. They are used especially for field and method lookup.

Second, we would like to know whether a specific program element is always present when a given set of features is present. This is necessary to ensure that references to program elements are always valid (i.e., not dangling). We need two sources of information for that. First, we need to know all features that introduce the program element in question (determined using the introduction table) and, second, we need to know which combinations of features are legal (determined using the feature model). For the field *renderer* of our example, the introduction table would yield the features MOZILLA and SAFARI and, from the feature model, it follows that MOZILLA and SAFARI are mutually exclusive, i.e., $\text{never}(\text{MOZILLA}, \text{SAFARI})$. But it can happen that none of the two features is present, which can invalidate a reference to the field. The type system needs to know about this situation.

To this end, we introduce a predicate *validref* that expresses that a program element is always reachable from a set of features. For example, $\text{validref}(\bar{\Omega}, \mathbf{C})$ holds if type \mathbf{C} is always reachable from the context $\bar{\Omega}$, $\text{validref}(\bar{\Omega}, \mathbf{C}.f)$ holds if field f of class \mathbf{C} is always reachable from the context $\bar{\Omega}$, and $\text{validref}(\bar{\Omega}, \mathbf{C}.m)$ holds if method m of class \mathbf{C} is always reachable from the context $\bar{\Omega}$. Applying *validref* to a list of program elements means that the conjunction of the predicates for every list element is taken. Finally, when we write $\text{validref}(\bar{\Omega}, \mathbf{C}) \dashv \bar{\Psi}$, we mean that program element \mathbf{C} is always reachable from a context $\bar{\Omega}$ in a subset $\bar{\Psi}$ of features of the product line.

In our prototype, we have implemented the above functions and predicates using a SAT solver that reasons about propositional formulas representing constraints on legal

feature combinations (see Section 4), as proposed by Batory [12] and Czarnecki and Pietroszek [22].

3.6 Refinement in FFJ_{PL}

In Figure 12, we show the functions *last* and *pred* for the navigation along the refinement chain. The two functions are identical to the ones of FFJ (cf. Figure 4). However, in FFJ_{PL} , there may be alternative declarations of a class and, in the refinement chain, refinement declarations may even precede class declarations, as long as the declaring features are mutually exclusive. Let us illustrate refinement in FFJ_{PL} by means of the example shown in Figure 13. Class \mathbf{C} is introduced in the features Φ_1 and Φ_3 . Feature Φ_2 refines class \mathbf{C} introduced by feature Φ_1 and feature Φ_4 refines class \mathbf{C} introduced by feature Φ_3 . Feature Φ_1 and Φ_2 are never present when feature Φ_3 or Φ_4 are present and vice versa. A call of $RT(\mathbf{C})$ would return the list Φ_1, \dots, Φ_4 , a call of $last(\mathbf{C})$ would return the qualified type $\Phi_4.\mathbf{C}$, and a call of $pred(\Phi_4.\mathbf{C})$ would return the qualified type $\Phi_3.\mathbf{C}$ and so on.

Navigating along the refinement chain

$$\frac{RT(\mathbf{C}) = \bar{\Psi}}{last(\mathbf{C}) = \Psi_n.\mathbf{C}} \quad \frac{RT(\mathbf{C}) = \bar{\Psi}, \Phi, \bar{\Omega} \quad \bar{\Psi} \neq \bullet}{pred(\Phi.\mathbf{C}) = \Psi_n.\mathbf{C}} \quad \frac{RT(\mathbf{C}) = \Phi, \bar{\Omega}}{pred(\Phi.\mathbf{C}) = \text{Base.Object}}$$

Fig. 12. Refinement in FFJ_{PL} .




Fig. 13. Multiple alternative refinements.

3.7 Subtyping in FFJ_{PL}

The subtype relation is more complicated in FFJ_{PL} than in FFJ. The reason is that a class may have multiple declarations in different features, each declaring possibly different superclasses, as illustrated in Figure 14. That is, when checking whether a class is a subtype of another class, we need to check whether the subtype relation holds in *all* alternative inheritance paths that may be reached from a given context. For example, `FooBar` is a subtype of `BarFoo` because `BarFoo` is a superclass of `FooBar` in every

program variant (since $always(\Phi_1, \Phi_2) = \Phi_2, \Phi_3$); but `FooBar` is not a subtype of `Foo` and `Bar` because, in both cases, a program variant exists in which `FooBar` is not a (indirect) subclass of the class in question.




Fig. 14. Multiple inheritance chains in the presence of alternative features.

In Figure 15, we show the subtype relation of FFJ_{PL} . The subtype relation $C <: E \dashv \bar{\Omega}$ is read as follows: in the context $\bar{\Omega}$, type C is a subtype of type E , i.e., type C is a subtype of type E in every variant in which also the features $\bar{\Omega}$ are present. The first rule in Figure 15 covers reflexivity and terminates the recursion over the inheritance hierarchy. The second rule states that class C is a subtype of class E if at least one declaration of C is always present (tested with *validref*) and if every of C 's declarations that may be present together with $\bar{\Omega}$ (tested with *sometimes*) declares some type D as its supertype and D is a subtype of E in the context $\bar{\Omega}$. That is, E must be a direct or indirect supertype of D in *all* variants in which the features $\bar{\Omega}$ are present. Additionally, supertype D must be always reachable from the context $(\bar{\Omega}, \Psi)$. When traversing the inheritance hierarchy, in each step, the context is extended by the feature that introduces the current class in question, e.g., $\bar{\Omega}$ is extended with Ψ .

Interestingly, the second rule subsumes the two FFJ rules for transitivity and direct superclass declaration because some declarations of C may declare E directly as its superclass and some declarations may declare another superclass D that is, in turn, a subtype of E , and the rule must be applicable to both cases simultaneously.




Fig. 15. Subtyping in FFJ_{PL} .

Applied to our example of Figure 14, we have $\text{FooBar} <: \text{FooBar} \dashv \Phi_1$ because of the reflexivity rule. We also have $\text{FooBar} <: \text{BarFoo} \dashv \Phi_1$ because FooBar is reachable from feature Φ_1 and every feature that introduces FooBar , namely Φ_1 , contains a corresponding class declaration that declares BarFoo as FooBar 's superclass, and BarFoo is always reachable from Φ_1 . However, we have $\text{FooBar} \not<: \text{Foo} \dashv \Phi_1$ and $\text{FooBar} \not<: \text{Bar} \dashv \Phi_1$ because FooBar 's immediate superclass BarFoo is not always a subtype of Foo respectively of Bar .

3.8 Auxiliary Definitions of FFJ_{PL}

Extending FFJ toward FFJ_{PL} makes it necessary to add and modify some auxiliary functions. The most complex changes concern the field and method lookup mechanisms.

Field Lookup. The auxiliary function *fields* collects the fields of a class including the fields of its superclasses and refinements. Since alternative class or refinement declarations may introduce alternative fields (or the same field with identical or alternative types), *fields* may return different fields for different feature selections. Since we want to type-check all valid variants, *field* returns multiple field lists (i.e., a list of lists) that cover all possible feature selections. Each inner list contains field declarations collected in an alternative path of the combined inheritance and refinement hierarchy.

For legibility, we separate the inner lists using the delimiter ‘ \circ ’. For example, looking up the fields of class FooBar in the context of feature Φ_1 (Figure 14) yields the list $A\ a, D\ d, E\ e \circ B\ b, D\ d, E\ e$ because the features Φ_2 and Φ_3 are mutually exclusive and one of them is present in each variant in which also Φ_1 is present. For readability, we use the metavariables \mathcal{F} and \mathcal{G} when referring to inner field lists. We abbreviate a list of lists of fields $\mathcal{F}_1 \circ \dots \circ \mathcal{F}_n$ by $\overline{\mathcal{F}}$. Analogously, $\overline{\mathcal{F}}$ is shorthand for $\mathcal{F}_{11} \circ \dots \circ \mathcal{F}_{n1} \circ \dots \circ \mathcal{F}_{1m} \circ \dots \circ \mathcal{F}_{nm}$.

Function *fields* receives a qualified type $\Phi.C$ and a context of selected features $\overline{\Omega}$. If we want all possible field lists, the context is empty. If we want only field lists for a subset of feature selections, e.g., only the fields that can be referenced from a term in a specific feature module, we can use the context to specify one or more features of which we know that they must be selected.

The basic idea of FFJ_{PL} 's field lookup is to traverse the combined inheritance and refinement hierarchy much like in FFJ . There are four situations that are handled differently:

1. The field lookup returns the empty list when it reaches Base.Object .
2. The field lookup ignores all fields that are introduced by features that are never present in a given context.
3. The field lookup collects all fields that are introduced by features that are always present in a given context. References to these fields are always valid.
4. The field lookup collects all fields that are introduced by features that may be present in a given context but that are not always present. In this case, a special

marker @ is added to the fields in question because we cannot guarantee that a reference to this field is safe in the given context.⁵ It is up to the type system to decide, based on the marker, whether this situation may provoke an error (e.g., the type system ignores the marker when looking for duplicate fields but reports an error when type checking object creations).

5. A special situation occurs when the field lookup identifies a group of alternative features. In such a group each feature is optional and excludes every other feature of the group and at least one feature of the group is always present in a given context. Once the field lookup identifies a group of alternative features, we split the result list, each list containing the fields of a feature of the group and the fields of the original list.

<i>Field lookup</i>	$fields(\bar{\Omega}, \Phi.C) = \overline{C f}$
$fields(\bar{\Omega}, \Phi.Object) = \bullet$	(FL-1)
$\frac{never(\bar{\Omega}, \Phi)}{fields(\bar{\Omega}, \Phi.C) = fields(\bar{\Omega}, pred(\Phi.C))}$	(FL-2)
$\frac{\begin{array}{l} sometimes(\bar{\Omega}, \Phi) \quad always(\bar{\Omega}, \Phi) = \Phi \\ CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C f}; \overline{M} \} \end{array}}{fields(\bar{\Omega}, \Phi.C) = append(fields(\bar{\Omega}, last(D)), \overline{C f})}$	(FL-3.1)
$\frac{\begin{array}{l} sometimes(\bar{\Omega}, \Phi) \quad always(\bar{\Omega}, \Phi) = \Phi \\ CT(\Phi.C) = \text{refines class } C \{ \overline{C f}; \overline{M} \} \end{array}}{fields(\bar{\Omega}, \Phi.C) = append(fields(\bar{\Omega}, pred(\Phi.C)), \overline{C f})}$	(FL-3.2)
$\frac{\begin{array}{l} sometimes(\bar{\Omega}, \Phi) \quad always(\bar{\Omega}, \Phi) = \bullet \\ CT(\Phi.C) = \text{class } C \text{ extends } D \{ \overline{C f}; \overline{M} \} \end{array}}{fields(\bar{\Omega}, \Phi.C) = append(fields(\bar{\Omega}, last(D)), \overline{C f}@)}$	(FL-4.1)
$\frac{\begin{array}{l} sometimes(\bar{\Omega}, \Phi) \quad always(\bar{\Omega}, \Phi) = \bullet \\ CT(\Phi.C) = \text{refines class } C \{ \overline{C f}; \overline{M} \} \end{array}}{fields(\bar{\Omega}, \Phi.C) = append(fields(\bar{\Omega}, pred(\Phi.C)), \overline{C f}@)}$	(FL-4.2)
$\frac{\begin{array}{l} sometimes(\bar{\Omega}, \Phi) \quad always(\bar{\Omega}, \Phi) = \bar{\Psi} \\ fields(\bar{\Omega}, \Phi.C) = fields((\bar{\Omega}, \Psi_1), \Phi.C) \circ \dots \circ fields((\bar{\Omega}, \Psi_n), \Phi.C) \end{array}}{fields(\bar{\Omega}, \Phi.C) = fields((\bar{\Omega}, \Psi_1), \Phi.C) \circ \dots \circ fields((\bar{\Omega}, \Psi_n), \Phi.C)}$	(FL-5)

Fig. 16. Field lookup in FFJ_{PL}.

In order to distinguish the different cases, we use the predicates and functions defined in Section 3.5 (especially *never*, *sometimes*, and *always*). The definition of func-

⁵ Note that the marker @ is generated during type checking, so we do not include it in the syntax of FFJ.

tion *fields*, shown in Figure 16, follows the intuition described above: Once `Base.Object` is reached, the recursion terminates (FL-1). When a feature is never reachable in the given context, *fields* ignores this feature and resumes with the previous one (FL-2). When a feature is mandatory (i.e., always present in a given context), the fields in question are added to each alternative result list, which were created in Rule FL-5 (FL-3.1 and FL-3.2).⁶ When a feature is optional, the fields in question, annotated with the marker $\textcircled{}$, are added to each alternative result list (FL-4.1 and FL-4.2). When a feature is part of an alternative group of features, we cannot immediately decide how to proceed. We split the result list in multiple lists (by means of multiple recursive invocations of *fields*), in which we add one of the alternative features to each context passed to an invocation of *fields* (FL-5).

$$\begin{array}{c}
\textit{Method type lookup} \qquad \boxed{mtype(\overline{\Omega}, m, \Phi.C) = \overline{B} \rightarrow B_0} \\
\\
mtype(\overline{\Omega}, m, \text{Base.Object}) = \bullet \qquad \text{(ML-1)} \\
\\
\frac{B_0 \ m(\overline{B} \ x) \ \{ \dots \} \in \overline{M} \quad \textit{sometimes}(\overline{\Omega}, \Phi) \quad CT(\Phi.C) = \text{class } C \text{ extends } D \ \{ \overline{C} \ f; \overline{M} \}}{mtype(\overline{\Omega}, m, \Phi.C) = mtype(\overline{\Omega}, m, \textit{pred}(\Phi.C)), mtype(\overline{\Omega}, m, \textit{last}(D)), \overline{B} \rightarrow B_0} \text{(ML-2)} \\
\\
\frac{B_0 \ m(\overline{B} \ x) \ \{ \dots \} \in \overline{M} \quad \textit{sometimes}(\overline{\Omega}, \Phi) \quad CT(\Phi.C) = \text{refines class } C \ \{ \overline{C} \ f; \overline{M} \}}{mtype(\overline{\Omega}, m, \Phi.C) = mtype(\overline{\Omega}, m, \textit{pred}(\Phi.C)), \overline{B} \rightarrow B_0} \text{(ML-3)} \\
\\
\frac{(\overline{m} \text{ is not defined in } \overline{M} \vee \textit{never}(\overline{\Omega}, \Phi)) \quad CT(\Phi.C) = \text{class } C \text{ extends } D \ \{ \overline{C} \ f; \overline{M} \}}{mtype(\overline{\Omega}, m, \Phi.C) = mtype(\overline{\Omega}, m, \textit{pred}(\Phi.C)), mtype(\overline{\Omega}, m, \textit{last}(D))} \text{(ML-4)} \\
\\
\frac{(\overline{m} \text{ is not defined in } \overline{M} \vee \textit{never}(\overline{\Omega}, \Phi)) \quad CT(\Phi.C) = \text{refines class } C \ \{ \overline{C} \ f; \overline{M} \}}{mtype(\overline{\Omega}, m, \Phi.C) = mtype(\overline{\Omega}, m, \textit{pred}(\Phi.C))} \text{(ML-5)}
\end{array}$$

Fig. 17. Method Lookup in FFJ_{PL} .

Method Type Lookup. Like in field lookup, in method lookup, we have to take alternative definitions of methods into account. But the lookup mechanism is simpler than in *fields* because the order of signatures found in the combined inheritance and refinement hierarchy is irrelevant for type checking. Hence, function *mtype* yields a simple list $\overline{B} \rightarrow B_0$ of signatures for a given method name m . For example, calling $mtype(\Phi_1, m, \Phi_1.C)$ in the context of Figure 14 yields the list $D \rightarrow A, B \rightarrow B$.

⁶ Function *append* adds to each inner list of a list of field lists a given field. Its implementation is straightforward and omitted for brevity.

In Figure 17, we show the definition of function *mtype*. For `Base.Object`, the empty list is returned (ML-1). If a class that is sometimes reachable introduces a method in question (ML-2), its signature is added to the result list and all possible predecessors in the refinement chain (using *pred*) and all possible subclasses are searched (using *last*). Likewise, if a refinement that is sometimes reachable introduces a method with the name searched (ML-3), its signature is added to the result list and all possible predecessors in the refinement chain are searched (using *pred*). If a class or refinement does not declare a corresponding method (ML-4 and ML-5) or the a class is never reachable, the search proceeds with the possible superclasses or predecessors.

The current definition of function *mtype* returns possibly many duplicate signatures. A straightforward optimization would be to remove duplicates before using the result list, which we omitted for simplicity.

$$\begin{array}{c}
\text{Valid class introduction} \quad \boxed{\text{introduce}(\bar{\Omega}, \Phi, \mathbf{C})} \\
\frac{\# \Psi : \left(\begin{array}{l} CT(\Psi, \mathbf{C}) = \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{C}} \text{ f}; \bar{\mathbf{M}} \} \\ \Psi \neq \Phi \quad \text{sometimes}(\bar{\Omega}, \Psi) \end{array} \right)}{\text{introduce}(\bar{\Omega}, \Phi, \mathbf{C})} \\
\\
\text{Valid field introduction} \quad \boxed{\text{introduce}(\bar{\Omega}, \mathbf{f}, \Phi, \mathbf{C})} \\
\frac{\forall \bar{\mathbf{E}} \bar{\mathbf{h}} \in \text{fields}(\bar{\Omega}, \Phi, \mathbf{C}) : \mathbf{f} \notin \bar{\mathbf{h}}}{\text{introduce}(\bar{\Omega}, \mathbf{f}, \Phi, \mathbf{C})} \\
\\
\text{Valid method introduction} \quad \boxed{\text{introduce}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C})} \\
\frac{\text{mtype}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C}) = \bullet}{\text{introduce}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C})} \\
\\
\text{Valid class refinement} \quad \boxed{\text{refine}(\bar{\Omega}, \Phi, \mathbf{C})} \\
\frac{RT(\mathbf{C}) = \bar{\Psi}, \Phi, \bar{\Pi} \quad \text{validref}(\bar{\Omega}, \mathbf{C}) \dashv \bar{\Psi}}{\text{refine}(\bar{\Omega}, \Phi, \mathbf{C})} \\
\\
\text{Valid method overriding} \quad \boxed{\text{override}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C}, \bar{\mathbf{C}} \rightarrow \mathbf{C}_0)} \\
\frac{RT(\mathbf{C}) = \bar{\Psi}, \Phi, \bar{\Pi} \quad \text{validref}(\bar{\Omega}, \mathbf{C}, \mathbf{m}) \dashv \bar{\Psi}, \Phi \\ \forall \bar{\mathbf{B}} \rightarrow \mathbf{B}_0 \in \text{mtype}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C}) : \bar{\mathbf{C}} = \bar{\mathbf{B}} \wedge \mathbf{C}_0 = \mathbf{B}_0}{\text{override}(\bar{\Omega}, \mathbf{m}, \Phi, \mathbf{C}, \bar{\mathbf{C}} \rightarrow \mathbf{C}_0)}
\end{array}$$

Fig. 18. Valid introduction, refinement, and overriding in FFJ_{PL} .

Valid Introduction, Refinement, and Overriding. In Figure 18, we show predicates for checking the validity of introduction, refinement, and overriding in FFJ_{PL} . Predicate *introduce* indicates whether a class with the qualified type $\Phi.C$ has not been introduced by any other feature Ψ that may be present in the context $\bar{\Omega}$. Likewise, *introduce* holds if a method m or a field f has not been introduced by a qualified type $\Phi.C$ (including possible predecessors and superclasses) that may be present in the given context $\bar{\Omega}$. To this end, it checks either whether *mtype* yields the empty list or whether f is not contained in every inner list returned by *fields*.

For a given refinement, predicate *refine* indicates whether a proper class, which is always reachable in the given context, has been declared previously in the refinement chain. We write $\text{validref}(\bar{\Omega}, C) \dashv \bar{\Psi}$ in order to state that a declaration of class C has been introduced in the set $\bar{\Psi}$ of features, which is only a subset of the features of the product line, namely the features that precede the feature that introduces class C . Predicate *override* indicates whether a declaration of method m has been introduced (and is always reachable) in some feature introduced by before the feature that refines m and whether every possible declaration of m in any predecessor of a $\Phi.C$ has the same signature.

3.9 Type Relation of FFJ_{PL}

The type relation of FFJ_{PL} consists of type rules for terms and well-formedness rules for classes, refinements, and methods, shown in Figure 19 and Figure 20.

Term Typing Rules. A term typing judgment in FFJ_{PL} is a quadruple, consisting of a typing context Γ , a term t , a list of types \bar{C} , and a feature Φ that contains the term (see Figure 19). A term can have multiple types in a product line because there may be multiple declarations of classes, fields, and methods. The list \bar{C} contains all possible types a term can have.

Rule T-VAR $_{PL}$ is standard and does not refer to the feature model. It yields a list consisting only of the type of the variable in question.

Rule T-FIELD $_{PL}$ checks whether a field access $t_0.f$ is well-typed in every possible variant in which also Φ is present. Based on the possible types \bar{E} of the term t_0 the field f is accessed from, the rule checks whether f is always reachable from Φ (using *validref*). Note that this is a key mechanism of FFJ_{PL} 's type system. It ensures that a field, being accessed, is definitely present in every valid program variant in which the field access occurs – without generating all these variants. Furthermore, all possible fields of all possible types \bar{E} are assembled in a nested list $\overline{\mathcal{F}, C f, \bar{G}}$ in which $C f$ denotes a declaration of the field f ; the call of $\text{fields}(\Phi, \overline{\text{last}(\bar{E})})$ is shorthand for $\text{fields}(\Phi, \text{last}(\bar{E}_1)) \dots \text{fields}(\Phi, \text{last}(\bar{E}_n))$, in which the individual result lists are concatenated. Finally, the list of all possible types $C_{11}, \dots, C_{n1}, \dots, C_{1m}, \dots, C_{nm}$ of field f becomes the list of types of the overall field access. Note that the result list may contain duplicates, which could be eliminated for optimization purposes.

Rule T-INVK $_{PL}$ checks whether a method invocation $t_0.m(\bar{t})$ is well-typed in every possible variant in which also Φ is present. Based on the possible types \bar{E} of the term t_0 the method m is invoked on, the rule checks whether m is always reachable from Φ

Term typing

$\Gamma \vdash t : \overline{C} \dashv \Phi$

$$\frac{x : C \in \Gamma}{\Gamma \vdash x : C \dashv \Phi} \quad (\text{T-VAR}_{PL})$$

$$\frac{\forall E \in \overline{E} : \text{validref}(\Phi, E.f) \quad \Gamma \vdash t_0 : \overline{E} \dashv \Phi \quad \text{fields}(\Phi, \text{last}(\overline{E})) = \overline{\mathcal{F}, C f, \mathcal{G}}}{\Gamma \vdash t_0.f : C_{11}, \dots, C_{n1}, \dots, C_{1m}, \dots, C_{nm} \dashv \Phi} \quad (\text{T-FIELD}_{PL})$$

$$\frac{\forall E \in \overline{E} : \text{validref}(\Phi, E.m) \quad \forall \overline{C} \in \overline{\overline{C}}, \forall \overline{D} \in \overline{\overline{D}} \in \overline{\overline{\overline{D}}} : \overline{C} <: \overline{D} \dashv \Phi \quad \Gamma \vdash t_0 : \overline{E} \dashv \Phi \quad \Gamma \vdash t : \overline{C} \dashv \Phi \quad \text{mtype}(\Phi, m, \text{last}(\overline{E})) = \overline{\overline{D}} \rightarrow \overline{B}}{\Gamma \vdash t_0.m(\hat{t}) : B_{11}, \dots, B_{n1}, \dots, B_{1m}, \dots, B_{nm} \dashv \Phi} \quad (\text{T-INVK}_{PL})$$

$$\frac{\text{validref}(\Phi, C) \quad \forall \overline{Dg} \in \overline{\mathcal{F}}, \forall \overline{C} \in \overline{\overline{C}} : \overline{C} <: \overline{D} \dashv \Phi \quad \Gamma \vdash t : \overline{C} \dashv \Phi \quad \text{fields}(\Phi, \text{last}(C)) = \overline{\mathcal{F}} \quad @ \notin \overline{\mathcal{F}}}{\Gamma \vdash \text{new } C(\hat{t}) : C \dashv \Phi} \quad (\text{T-NEW}_{PL})$$

$$\frac{\text{validref}(\Phi, C) \quad \Gamma \vdash t_0 : \overline{E} \dashv \Phi \quad \forall E \in \overline{E} : (E <: C \dashv \Phi \vee C <: E \dashv \Phi)}{\Gamma \vdash (C)t_0 : C \dashv \Phi} \quad (\text{T-UDCAST}_{PL})$$

$$\frac{\text{validref}(\Phi, C) \quad \text{stupid warning} \quad \Gamma \vdash t_0 : \overline{E} \dashv \Phi \quad \exists E \in \overline{E} : (C \not<: E \dashv \Phi \wedge E \not<: C \dashv \Phi)}{\Gamma \vdash (C)t_0 : C \dashv \Phi} \quad (\text{T-SCAST}_{PL})$$

Fig. 19. Term typing in FFJ_{PL} .

Method typing

$\boxed{M \text{ OK} \dashv \Phi.C}$

$$\frac{\overline{x} : \overline{B}, \text{this} : C \vdash t_0 : \overline{E} \dashv \Phi \quad \forall E \in \overline{E} : E <: B_0 \dashv \Phi \quad \text{validref}(\Phi, \overline{B}) \quad \text{introduce}(\Phi, m, \text{last}(D)) \quad CT(\Phi.C) = \text{class } C \text{ extends } D \{ C f; \overline{M} \}}{B_0 m(\overline{B} x) \{ \text{return } t_0; \} \text{ OK} \dashv \Phi.C}$$

$$\frac{\overline{x} : \overline{B}, \text{this} : C \vdash t_0 : \overline{E} \dashv \Phi \quad \forall E \in \overline{E} : E <: B_0 \dashv \Phi \quad \text{validref}(\Phi, \overline{B}) \quad \text{override}(\Phi, m, \text{last}(D), \overline{B} \rightarrow B_0) \quad CT(\Phi.C) = \text{class } C \text{ extends } D \{ C f; \overline{M} \}}{\text{overrides } B_0 m(\overline{B} x) \{ \text{return } t_0; \} \text{ OK} \dashv \Phi.C}$$

$$\frac{\overline{x} : \overline{B}, \text{this} : C \vdash t_0 : \overline{E} \dashv \Phi \quad \forall E \in \overline{E} : E <: B_0 \dashv \Phi \quad \text{validref}(\Phi, \overline{B}) \quad \text{introduce}(\Phi, m, \text{pred}(\Phi.C)) \quad CT(\Phi.C) = \text{refines class } C \{ C f; \overline{M} \}}{B_0 m(\overline{B} x) \{ \text{return } t_0; \} \text{ OK} \dashv \Phi.C}$$

$$\frac{\overline{x} : \overline{B}, \text{this} : C \vdash t_0 : \overline{E} \dashv \Phi \quad \forall E \in \overline{E} : E <: B_0 \dashv \Phi \quad \text{validref}(\Phi, \overline{B}) \quad \text{override}(\Phi, m, \text{pred}(\Phi.C), \overline{B} \rightarrow B_0) \quad CT(\Phi.C) = \text{refines class } C \{ C f; \overline{M} \}}{\text{overrides } B_0 m(\overline{B} x) \{ \text{return } t_0; \} \text{ OK} \dashv \Phi.C}$$

Class typing

$\boxed{L \text{ OK} \dashv \Phi}$

$$\frac{\forall f \in \overline{f} : \text{validref}(\Phi, D) \quad \text{validref}(\Phi, \overline{C}) \quad \text{introduce}(\Phi, f, \text{last}(D)) \quad \text{introduce}(\Phi, \Phi.C) \quad \overline{M} \text{ OK} \dashv \Phi.C}{\text{class } C \text{ extends } D \{ C f; \overline{M} \} \text{ OK} \dashv \Phi}$$

Refinement typing

$\boxed{R \text{ OK} \dashv \Phi}$

$$\frac{\forall f \in \overline{f} : \text{validref}(\Phi, \overline{C}) \quad \text{introduce}(\Phi, f, \text{pred}(\Phi.C)) \quad \text{refine}(\Phi, \Phi.C) \quad \overline{M} \text{ OK} \dashv \Phi.C}{\text{refines class } C \{ C f; \overline{M} \} \text{ OK} \dashv \Phi}$$

Fig. 20. Well-formedness rules of FFJ_{PL} .

(using *validref*). As with field access, this check is essential. It ensures that in generated programs only methods are invoked that are also present. Furthermore, all possible signatures of m of all possible types \bar{E} are assembled in the nested list $\bar{D} \rightarrow \bar{B}$ and it is checked that all possible lists \bar{C} of argument types of the method invocation are subtypes of all possible lists \bar{D} of parameter types of the method (this implies that the lengths of the two lists must be equal). A method invocation has multiple types assembled in a list that contains all result types of method m determined by *mtype*. As with field access, duplicates should be eliminated for optimization purposes.

Rule T-NEW_{PL} checks whether an object creation `new C(\bar{t})` is well-typed in every possible variant in which also Φ is present. Specifically, it checks whether there is a declaration of class C always reachable from Φ . Furthermore, all possible field combinations of C are assembled in the nested list \bar{F} , and it is checked whether all possible combinations of argument types passed to the object creation are subtypes of the types of all possible field combinations (this implies that the number of arguments types must equal the number of field types). The fields of the result list must not be annotated with the marker $@$ since optional fields may not be present in every variant and references may become invalid (see field lookup).⁷ An object creation has only a single type C .

Rules T-UDCAST_{PL} and T-SCAST_{PL} check whether casts are well-typed in every possible variant in which also Φ is present. This is done by checking whether the type C the term t_0 is cast to is always reachable from Φ and whether this type is a subtype, supertype, or unrelated type of all possible types \bar{E} the term t_0 can have. We have only a single rule T-UDCAST_{PL} for up- and downcasts because the list \bar{E} of possible types may contain super- and subtypes of C simultaneously. If there is a type in the list which leads to a stupid case, we flag a *stupid warning*. A cast yields a list containing only a single type C .

Well-Formedness Rules. In Figure 20, we show the well-formedness rules of classes, refinements, and methods.

Like in FFJ, the typing judgment of classes and refinements is a binary relation between a class or refinement declaration and a feature. The rule of classes checks whether all methods are well-formed in the context of the class' qualified type. Moreover, it checks whether the class declaration is unique in the scope of the enclosing feature Φ , i.e., whether no other feature, that may be present together with feature Φ , introduces a class with an identical name (using *introduce*). Furthermore, it checks whether the superclass and all field types are always reachable from Φ (using *validref*). Finally, it checks whether none of the fields of the class declaration have been introduced before (using *introduce*). The well-formedness rule of refinements is analogous, except that the rule checks that there is *at least one* class declaration reachable that is refined and that has been introduced before the refinement (using *refine*).

The typing judgment of methods is a binary relation between a method declaration and the qualified type that declares the method. Like in FFJ, there are four different rules for methods (from top to bottom in Figure 20)

1. that do not override another method and that are declared by classes,

⁷ The treatment of $@$ is semiformal but simplifies the rule.

2. that override another method and that are declared by classes,
3. that do not override another method and that are declared by refinements,
4. that override another method and that are declared by refinements.

All four rules check whether all possible types \bar{E} of the method body are subtypes of the declared return type B_0 of the method and whether the argument types \bar{B} are always reachable from the enclosing feature Φ (using *validref*).

For methods that are introduced, it is checked, using *introduce*, whether no method with identical name has been introduced in any possible superclass (Rule 1) or in any possible predecessor in the refinement chain (Rule 3). For methods that override other methods, it is checked, using *override*, whether a method with identical name and signature exists in any possible superclass (Rule 2) or in any possible predecessor in the refinement chain (Rule 4).

Well-Typed FFJ_{PL} Product Lines. An FFJ_{PL} product line, consisting of a term, a class table, an introduction table, and a refinement table, is well-typed if

- the term is well-typed (checked using FFJ_{PL}'s term typing rules),
- all classes and refinements stored in the class table are well-formed (checked using FFJ_{PL}'s well-formedness rules), and
- the class, introduction, and refinement tables are well-formed (ensured by the corresponding sanity conditions).

3.10 Type Safety of FFJ_{PL}

Type checking in FFJ_{PL} is based on information contained in the class table, introduction table, refinement table, and feature model. The first three are filled by the compiler that has parsed the code base of the product line. The feature model is supplied directly by the user (or tool). The compiler determines which class and refinement declarations belong to which features. The classes and refinements of the class table are checked using their well-formedness rules which, in turn, use the well-formedness rules for methods and the term typing rules for method bodies. Several rules use the introduction and refinement tables in order to map types, fields, and methods to features and the feature model to navigate along refinement chains and to check the presence of program elements.

What does type safety mean in the context of a product line? The product line itself is never evaluated; rather, different programs are derived that are then evaluated. Hence, the property we are interested in is that *all* programs that can be derived from a well-typed product line are in turn well-typed. Furthermore, we would like to be sure that *all* FFJ_{PL} product lines, from which only well-typed FFJ programs can be derived, are well-typed. We formulate the two properties as the two theorems *Correctness of FFJ_{PL}* and *Completeness of FFJ_{PL}*.

Correctness

THEOREM 3.1 (*Correctness of FFJ_{PL}*) Given a well-typed FFJ_{PL} product line pl (including with a well-typed term t , well-formed class, introduction, and refinement tables CT , IT , and RT , and a feature types model FM), every program that can be derived with a valid feature selection fs is a well-typed FFJ program (cf. Figure 10).

$$\frac{pl = (t, CT, IT, RT, FM) \quad pl \text{ is well-typed} \quad fs \text{ is valid in } FM}{derive(pl, fs) \text{ is well-typed}}$$

Function *derive* collects the feature modules from a product line according to a user's selection fs , i.e., non-selected feature modules are removed from the derived program. After this derivation step, the class table contains only classes and refinements stemming from the selected feature modules. We define a *valid feature selection* to be a list of features whose combination does not contradict the constraints implied by the feature model.

The proof idea is to show that the type derivation tree of an FFJ_{PL} product line is a superimposition of multiple so-called *type derivation slices*. As usual, the type derivation proceeds from the root (i.e., an initial type rule that checks the term and all classes and refinements of the class table) to the leaves (type rules that do not have a premise) of the type derivation tree. Each time a term has multiple types, e.g., a method has different alternative return types, which is caused by multiple mutually exclusive method declarations, the type derivation splits into multiple branches. With *branch* we refer only to positions in which the type derivation tree is split into multiple subtrees in order to type check multiple mutually exclusive term definitions. Each subtree from the root of the type derivation tree along the branches toward a leaf is a type derivation slice. Each slice corresponds to the type derivation of a feature-oriented program.

Let us illustrate the concept of a type derivation slice by a simplified example. Suppose the application of an arbitrary type rule to a term t somewhere in the type derivation. Term t has multiple types \bar{C} due to different alternative definitions of t 's subterms. For simplicity, we assume here that t has only a single subterm t_0 , like in the case of a field access ($t = t_0.f$), in which the overall term t has multiple types depending on t_0 's and f 's types; the rule can be easily extended to multiple subterms by adding a predicate per subterm. The type rule ensures the well-typedness of all possible variants of t on the basis of the variants of t 's subterm t_0 . Furthermore, the type rule checks whether a predicate *predicate* (e.g., $C <: D$) holds for each variant of the subterm with its possible types \bar{E} , written *predicate*($t_0 : E_i$). The possible types \bar{C} of the overall term follow in some way from the possible types \bar{E} of its subterm. Predicate *validref* is used to check whether all referenced elements and types are present in all valid variants, including different combinations of optional features. For the general case, this can be written as follows:

$$\frac{\begin{array}{c} predicate(t_0 : E_1) \quad predicate(t_0 : E_2) \quad \dots \quad predicate(t_0 : E_n) \\ t_0 : \bar{E} \quad always(\dots) \end{array}}{\Gamma \vdash t : \bar{C} \dashv \Phi} \text{(T-*}_{PL}\text{)}$$

The different uses of *predicate* in the premise of an FFJ_{PL} type rule correspond to the branches in the type derivation that denote alternative definitions of subterms.

Hence, the premise of the FFJ_{PL} type rule is the conjunction of the different premises that cover the different alternative definitions of the subterms of a term.

The proof strategy is as follows. Assuming that the FFJ_{PL} type system ensures that each slice is a valid FFJ type derivation (see Lemma B.1 in Appendix B.1) and that each valid feature selection corresponds to a single slice (since alternative features have been removed; see Lemma B.2 in Appendix B.1), each feature-oriented program that corresponds to a valid feature selection is guaranteed to be well-typed. Note that multiple valid feature selections may correspond to the same slice because of the presence of optional features. It follows that, for every valid feature selection, we derive a well-formed FFJ program – since its type derivation is valid – whose evaluation satisfies the properties of progress and preservation (see Appendix A). In Appendix B, we describe the proof of Theorem 3.1 in more detail.

Completeness

THEOREM 3.2 (*Completeness of FFJ_{PL}*) Given an FFJ_{PL} product line pl (including a well-typed term t , well-formed class, introduction, and refinement tables CT , IT , and RT , and a feature model FM), and given that *all* valid feature selections fs yield well-typed FFJ programs, according to Theorem 3.1, pl is a well-typed product line according to the rules of FFJ_{PL} .

$$\frac{pl = (t, CT, IT, RT, FM) \quad \forall fs : (fs \text{ is valid in } FM \Rightarrow \text{derive}(pl, fs) \text{ is well-typed})}{pl \text{ is well-typed}}$$

The proof idea is to examine three basic cases and to generalize subsequently: (1) pl has only mandatory features; (2) pl has only mandatory features except a single optional feature; (3) pl has only mandatory features except two mutually exclusive features. All other cases can be formulated as combinations of these three basic cases. To this end, we divide the possible relations between features into three disjoint sets: (1) a feature is reachable from another feature in *all* variants, (2) a feature is reachable from another feature in *some*, but *not in all*, variants, (3) two features are mutually exclusive. From these three possible relations, we can prove the three basic cases in isolation and, subsequently, construct a general case that can be phrased as a combination of the three basic cases. The description of the general case and the reduction finish the proof of Theorem 3.2. In Appendix B, we describe the proof of Theorem 3.2 in detail.

4 Implementation & Discussion

We have implemented FFJ and FFJ_{PL} in Haskell, including the program evaluation and type checking of product lines. The FFJ_{PL} compiler expects a set of feature modules and a feature model both of which, together, represent the product line. A feature module is represented by a directory. The files found inside a feature module’s directory are assigned to / belong to the enclosing feature. The FFJ_{PL} compiler stores this information for type checking. Each file may contain multiple classes and class refinements. In Figure 21, we show a snapshot of our test environment, which is based on Eclipse

and a Haskell plugin⁸. We use Eclipse to interpret or compile our FFJ and FFJ_{PL} type systems and interpreters. Specifically, the figure shows the directory structure of our email system. The file `EmailClient.features` contains the user’s feature selection and the feature model of the product line.




Fig. 21. Snapshot of the test environment of the Haskell implementation.

The feature model of a product line is represented by a propositional formula, following the approach of Batory [12] and Czarnecki and Pietroszek [22]. Propositional formulas are an effective way of representing the relationships between features (e.g., of specifying which feature implies the presence and absence of other features and of machine checking whether a feature selection is valid). For example, we have implemented predicate *sometimes* as follows:

$$sometimes(FM, \bar{\Omega}, \Psi) = satisfiable(FM \wedge \Omega_1 \wedge \dots \wedge \Omega_n \wedge \Psi)$$

The feature model is an propositional formula; feature are variables; and *satisfiable* is a satisfiability solver. Likewise, we have implemented predicate *always* on the basis of logical reasoning on propositional formulas:

$$always(FM, \bar{\Omega}, \Psi) = \neg(satisfiable(\neg(FM \Rightarrow ((\Omega_1 \wedge \dots \wedge \Omega_n) \Rightarrow \Psi))))$$

For a more detailed explanation of how propositional formulas relate to feature models and feature selections, we refer the interest to the work of Batory [12].

In Figure 22, we show the textual specification of the feature model of our email system, which can be passed directly to the FFJ_{PL} compiler.

⁸ <http://eclipsefp.sourceforge.net/haskell/>

```

1 features:
2   EmailClient IMAP POP3 MIME SSL Text Mozilla Safari
3
4 model:
5   EmailClient implies (IMAP or POP3);
6   IMAP implies EmailClient;
7   POP3 implies EmailClient;
8   MIME implies EmailClient;
9   SSL implies EmailClient;
10  Text implies (IMAP or POP3);
11  Mozilla implies (IMAP or POP3);
12  Safari implies (IMAP or POP3);
13  Mozilla implies (not Safari);
14  Safari implies (not Mozilla);

```

Fig. 22. Feature model of an email client product line.

The first section (**features:**) of the file representing the feature model defines an ordered set of names of the features of the product line and the second section (**model:**) defines constraints on the features' presence in the derived programs. In our example, each email client supports either the protocols IMAP, POP3, or both. Furthermore, every feature requires the presence of the base feature EMAILCLIENT. Feature TEXT requires either the presence of IMAP or POP3 or both – the same for MOZILLA and SAFARI. Finally, feature MOZILLA requires the absence of feature SAFARI and vice versa.

On the basis of the feature modules and the feature model, FFJ_{PL}'s type system checks the entire product line and identifies valid program variants that still contain type errors. A SAT solver is used to check whether elements are never, sometimes, or always reachable. If an error is found, the product line is rejected as ill-formed. If not, a feature-oriented program guaranteed to be well-formed can be derived on the basis of a user's feature selection. This program can be evaluated using the standard evaluation rules of FFJ, which we have also implemented in Haskell.

In contrast to previous work on type checking feature-oriented product lines [64,23], our type system provides detailed error messages. This is possible due to the fine-grained checks at the level of individual term typing and well-formedness rules. For example, if a field access succeeds only in some program variants, this fact can be reported to the user and the error message can point to the erroneous field access. Previously proposed type systems compose all code of all features of a product line and extract a single propositional formula, which is checked for satisfiability. If the formula is not satisfiable (i.e., a type error has occurred), it is not possible to identify the location that has caused the error (at least not without further information). See Section 5, for a detailed discussion of related approaches.

We made several tests and experiments with our Haskell implementation. However, real-world tests were not feasible because of two reasons. First, in previous work it has been already demonstrated that feature-oriented product lines require proper type systems and that type checking entire real-world product lines is feasible and useful [64]. Second, like FJ, FFJ is a core language into which all Java programs can be compiled and which, by its relative simplicity, is suited for the formal definition and proof of language properties – in our case, a type system and its correctness and completeness.

But, a core language is never suited for the development of real-world programs. This is why our examples and test programs are of similar size and complexity as the FJ examples of Pierce [54]. Type checking our test programs required acceptable amounts of time (in the order of magnitude of milliseconds per product line). We do not claim to be able to handle full-sized feature-oriented product lines by hand-coding them in FFJ_{PL} . Rather, this would require an expansion of the type system to full Java (including support for features as provided by AHEAD [13] or FeatureHouse [9]) – an enticing goal, but one for the future (especially, as Java’s informal language specification [28] has 688 pages). Our work lays a foundation for implementing type systems in that it provides evidence that core feature-oriented mechanisms are type sound and type systems of feature-oriented product lines can be implemented correctly and completely.

Still, we would like to make some predictions on the scalability of our approach. The novelty of our type system is that it incorporates alternative features and, consequently, alternative definitions of classes, fields, and methods. This leads to a type derivation tree with possibly multiple branches denoting alternative term types. Hence, performing a type derivation of product line with many alternative features may consume a significant amount of computation time and memory. It seems that this overhead is the price for allowing alternative implementation of program parts.

Nevertheless, our approach minimizes the overhead caused by alternative features compared to the naive approach. In the naive approach, all possible programs are derived and type checked subsequently. In our approach, we type check the entire code base of the product line and branch the type derivation only at terms that really have multiple, alternative types, and not at the level of entire program variants, as done in the naive approach. Our experience with feature-oriented product lines shows that, usually, there are not many alternative features in a product line, but mostly optional features [42,3,64,37,59,11,9,5,6,57,60]. For example, in the Berkeley DB product line (JE edition; 80 000 lines of code) there are 99 feature modules, but only two pairs of them alternative [9,37]; in the Graph Product Line there are 26 feature modules, of which only three pairs are alternative [42,9]. A further observation is that most alternative features that we encountered do not alter types. That is, there are multiple definitions of fields and methods but with equal types. For example, GPL and Berkeley DB contain alternative definitions of a few methods but only with identical signatures. Type checking these product lines with our approach, the type derivation would have almost no branches. In the naive approach, still many program variants exist due to optional features. Hence, our approach is preferable. For example, in a product line with n features and $c * n$ variants (with c being a constant), in our approach, the type system would have to check n feature modules (with some few branches in the type derivation and solving few simple SAT problems; see below) and, in the naive approach, the type system would have to check, at least, $2 * n$ feature modules but, commonly, $2 * n * m$ with $m < n$. For product lines with a higher degree of variability, e.g., with n^2 or even 2^n variants the benefit of our approach becomes even more significant. We believe that this benefit can make a difference in real world product line engineering.

A further point is that almost all typing and well-formedness rules contain calls to the built-in SAT solver. This results in possibly many invocations of the SAT solver at type checking time. Determining the satisfiability of a propositional formula is in

general an \mathcal{NP} -complete problem. However, it has been shown that the structures of propositional formulas occurring in software product lines are simple enough to scale satisfiability solving to thousands of features [47]. Furthermore, in our experiments, we have observed that many calls to the SAT solver are redundant, which is easy to see when thinking about type checking feature-oriented product lines where the presence of single types or members is checked in many type rules. We have implemented a caching mechanism to decrease the number of calls to the SAT solver to a minimum.

Finally, the implementation in Haskell helped us a lot with the evaluation of the correctness of our type rules. It can serve other researchers to reproduce and evaluate our work and to experiment with further (feature-oriented) language mechanisms. The implementations of FFJ and FFJ_{PL}, along with test programs, can be downloaded from the Web.⁹

5 Related Work

We divide our discussions of related work into two parts: the implementation, formal models, and type systems (1) of feature-oriented programs and (2) of feature oriented product lines.

5.1 Feature-Oriented Programs

FFJ has been inspired by several feature-oriented languages and tools, most notably AHEAD/Jak [13], FeatureC++ [10], FeatureHouse [9], and Prehofer’s feature-oriented Java extension [55]. Their key aim is to separate the implementation of software artifacts, e.g., classes and methods, from the definition of features. That is, classes and refinements are not annotated or declared to belong to a feature. There is no statement in the program text that defines explicitly a connection between code and features. Instead, the mapping of software artifacts to features is established via so-called containment hierarchies, which are basically directories containing software artifacts. The advantage of this approach is that a feature’s implementation can include, beside classes in the form of Java files, also other supporting documents, e.g., documentation in the form of HTML files, grammar specifications in the form of JavaCC files, or build scripts and deployment descriptors in the form of XML files [13]. To this end, feature composition merges not only classes with their refinements but also other artifacts, such as HTML or XML files, with their respective refinements [2,9].

Another class of programming languages that provide mechanisms for the definition and extension of classes and class hierarchies includes, e.g., *ContextL* [29], *Scala* [52], and *Classbox/J* [14]. The difference to feature-oriented languages is that they provide explicit language constructs for aggregating the classes that belong to a feature, e.g., family classes, classboxes, or layers. This implies that non-code software artifacts cannot be included in a feature [11]. However, FFJ still models a subset of these languages, in particular, class refinement.

Similarly, related work on a formalization of the key concepts underlying feature-oriented programming has not disassociated the concept of a feature from the level of

⁹ <http://www.fosd.de/ffj>

code. Especially, calculi for mixins [26,16,1,34], traits [41], family polymorphism and virtual classes [33,25,30,18], path-dependent types [52,51], open classes [20], dependent classes [27], and nested inheritance [50] either support only the refinement of single classes or expect the classes that form a semantically coherent unit (i.e., that belong to a feature) to be located in a physical module that is defined in the host programming language. For example, a virtual class is by definition an inner class of the enclosing object, and a classbox is a package that aggregates a set of related classes. Thus, FFJ differs from previous approaches in that it relies on contextual information that has been collected by the compiler, e.g., the features' composition order or the mapping of code to features.

A different line of research aims at the language-independent reasoning about features [13,44,9,39]. The calculus gDeep is most closely related to FFJ since it provides a type system for feature-oriented languages that is language-independent [4]. The idea is that the recursive process of merging software artifacts, when composing hierarchically structured features, is very similar for different host languages, e.g., for Java, C#, and XML. The calculus describes formally how feature composition is performed and what type constraints have to be satisfied. In contrast, FFJ does not aspire to be language-independent, although the key concepts can certainly be used with different languages. The advantage of FFJ is that its type system can be used to check whether terms of the host language (Java or FJ) violate the principles of feature orientation, e.g., whether methods refer to classes that have been added by other features. Due to its language independence, gDeep does not have enough information to perform such checks.

5.2 Feature-Oriented Product Lines

Our work on type checking feature-oriented product lines was motivated by the work of Thaker et al. [64]. They suggested the development of a type system for feature-oriented product lines that does not check all individual programs but the individual feature implementations. They have implemented an (incomplete) type system and, in a number of case studies on real product lines, they found numerous hidden errors using their type rules. Nevertheless, the implementation of their type system is ad-hoc in the sense that it is described only informally, and they do not provide a correctness and completeness proof. Our type system has been inspired by their work and we were able to provide a formalization and a proof of type safety.

In a parallel line of work, Delaware et al. have developed a formal model of a feature-oriented language, called *Lightweight Feature Java (LFJ)*, and a type system for feature-oriented product lines [23]. Their work was also influenced by the practical work of Thaker et al. So, it is not surprising that it is closest to ours. However, there are numerous differences. First, their formal model of a feature-oriented language is based on *Lightweight Java (LJ)* [62] and not on *Featherweight Java (FJ)*. While LJ is more expressive, it is also more complex. We decided for the simpler variant FJ, omitting, e.g., constructors and mutable state. Second, Delaware et al. do not model feature-oriented mechanisms, such as class or method refinements, directly in the semantics and type rules of the language. Instead, they introduce a transformation step in which LFJ code is “compiled down” to LJ code, i.e., they flatten refinement chains to single classes. Proceeding likewise, we would have to generate first an FJ program from an

FFJ product line and type check the FJ program (that consists of some or all possible features of the product line) subsequently. We refrained from such a transformation step in order to model the semantics of feature-oriented mechanisms directly in terms of dedicated field and method lookup mechanisms as well as special well-formed rules for method and class refinements.

Lagorio et al. have shown that a flattening semantics and a direct semantics are equivalent [40]. An advantage of a “direct” semantics is that it allows a type checking and error reporting at a finer grain. In LFJ, all feature modules are composed and a single propositional formula is generated and tested for satisfiability; if the formula is not satisfiable, it is difficult to identify precisely the point of failure. In FFJ_{PL} , the individual type rules consult the feature model and can point directly to the point of failure.

A further advantage of our approach is that it leaves open *when* feature composition is performed. Currently, feature composition is modeled in FFJ/FFJ_{PL} as a static process done before compilation but, with our approach, it becomes possible to model dynamic feature composition at run time [58,53] by making the class and feature tables and the feature model dynamic, i.e., allowing them to change during a computation. With LFJ this is not possible. Hutchins has shown that feature composition can be performed by an interpreter and partial evaluation can be used to pre-evaluate the parts of a composition that are static [31]. However, Delaware et al. have developed a machine-checked model of their type system formalized with the theorem prover Coq [15]. Our proof is hand-written, but we have a Haskell implementation of the FFJ and FFJ_{PL} calculi that we have tested thoroughly.

Even previously to the work of Thaker et al., Czarnecki et al. presented an automatic verification procedure for ensuring that no ill-structured UML model template instances will be generated from a valid feature selection [22]. That is, they type check product lines that consist not of Java programs but of UML models. They use OCL (object constraint language) constraints to express and implement a type system for model composition. In this sense, their aim is very similar to that of FFJ_{PL} , but limited to model artifacts – although they have proposed to generalize their work to programming languages.

Kästner et al. have implemented a tool, called CIDE, that allows a developer to decompose a software system into features via annotations [38]. In contrast to other feature-oriented languages and tools, the link between code and features is established via annotations. If a user selects a set of features, all code that is annotated with features (using background colors) that are not present in the selection is removed. Kästner et al. have developed a formal calculus and a set of type rules that ensure that only well-typed programs can be generated from a valid feature selection [36]. For example, if a method declaration is removed, the remaining code must not contain calls to this method. CIDE’s type rules are related to the type rules of FFJ_{PL} but, so far, mutually exclusive features are not supported in CIDE. In some sense, FFJ_{PL} and CIDE represent two sides of the same coin: the former aims at the composition of feature modules, the latter at the annotation of feature-related code.

6 Conclusion

A feature-oriented product line imposes severe challenges on type checking. The naive approach of checking all individual programs of a product line is not feasible because of the combinatorial explosion of program variants. Hence, the only practical option is to check the entire code base of a product line, including all features, and, based on the information of which feature combinations are valid, to ensure that it is not possible to derive a valid program variant that contains type errors.

We have developed such a type system based on a formal model of a feature-oriented Java-like language, called Feature Featherweight Java (FFJ). A distinguishing property of our work is that we have modeled the semantics and type rules for core feature-oriented mechanisms directly, without compiling feature-oriented code down to a lower-level representation such as object-oriented Java code. The direct semantics allows us to reason about core feature-oriented mechanisms in terms of themselves and not of generated lower-level code. A further advantage is the fine-grained error reporting and that the time of feature composition may vary between compile time and run time.

We have demonstrated and proved that, based on a valid feature selection, our type system ensures that every program of a feature-oriented product line is well-formed and that our type system is complete. Our implementation of FFJ, including the type system for product lines, indicates the feasibility of our approach and can serve as a testbed for experimenting with further feature-oriented mechanisms.

Acknowledgment

This work is being funded in part by the German Research Foundation (DFG), project number AP 206/2-1.

References

1. D. Ancona, G. Lagorio, and E. Zucca. Jam—Designing a Java Extension with Mixins. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 25(5):641–712, 2003.
2. F. Anfurrutia, O. Díaz, and S. Trujillo. On Refining XML Artifacts. In *Proceedings of the International Conference on Web Engineering (ICWE)*, volume 4607 of *LNCS*, pages 473–478. Springer-Verlag, 2007.
3. S. Apel and K. Böhm. Towards the Development of Ubiquitous Middleware Product Lines. In *Software Engineering and Middleware*, volume 3437 of *LNCS*, pages 137–153. Springer-Verlag, 2004.
4. S. Apel and D. Hutchins. An Overview of the gDeep Calculus. Technical Report MIP-0712, Department of Informatics and Mathematics, University of Passau, 2007.
5. S. Apel, F. Janda, S. Trujillo, and C. Kästner. Model Superimposition in Software Product Lines. In *Proceedings of the International Conference on Model Transformation (ICMT)*, volume 5563 of *LNCS*, pages 4–19. Springer-Verlag, 2009.
6. S. Apel, C. Kästner, A. Gröbinger, and C. Lengauer. Feature (De)composition in Functional Programming. In *Proceedings of the International Conference on Software Composition (SC)*, volume 5634 of *LNCS*, pages 9–26. Springer-Verlag, 2009.

7. S. Apel, C. Kästner, and C. Lengauer. An Overview of Feature Featherweight Java. Technical Report MIP-0802, Department of Informatics and Mathematics, University of Passau, 2008.
8. S. Apel, C. Kästner, and C. Lengauer. Feature Featherweight Java: A Calculus for Feature-Oriented Programming and Stepwise Refinement. In *Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE)*, pages 101–112. ACM Press, 2008.
9. S. Apel, C. Kästner, and C. Lengauer. FeatureHouse: Language-Independent, Automated Software Composition. In *Proceedings of the International Conference on Software Engineering (ICSE)*, pages 221–231. IEEE CS Press, 2009.
10. S. Apel, T. Leich, M. Rosenmüller, and G. Saake. FeatureC++: On the Symbiosis of Feature-Oriented and Aspect-Oriented Programming. In *Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE)*, volume 3676 of LNCS, pages 125–140. Springer-Verlag, 2005.
11. S. Apel, T. Leich, and G. Saake. Aspectual Feature Modules. *IEEE Transactions on Software Engineering (TSE)*, 34(2):162–180, 2008.
12. D. Batory. Feature Models, Grammars, and Propositional Formulas. In *Proceedings of the International Software Product Line Conference (SPLC)*, volume 3714 of LNCS, pages 7–20. Springer-Verlag, 2005.
13. D. Batory, J. Sarvela, and A. Rauschmayer. Scaling Step-Wise Refinement. *IEEE Transactions on Software Engineering (TSE)*, 30(6):355–371, 2004.
14. A. Bergel, S. Ducasse, and O. Nierstrasz. Classbox/J: Controlling the Scope of Change in Java. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 177–189. ACM Press, 2005.
15. Y. Bertot and P. Casteran. *Interactive Theorem Proving and Program Development – Coq’Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, 2004.
16. V. Bono, A. Patel, and V. Shmatikov. A Core Calculus of Classes and Mixins. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 1628 of LNCS, pages 43–66. Springer-Verlag, 1999.
17. G. Bracha and W. Cook. Mixin-Based Inheritance. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP) and International Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, pages 303–311. ACM Press, 1990.
18. D. Clarke, S. Drossopoulou, J. Noble, and T. Wrigstad. Tribe: A Simple Virtual Class Calculus. In *Proceedings of the International Conference on Aspect-Oriented Software Development (AOSD)*, pages 121–134. ACM Press, 2007.
19. P. Clements and L. Northrop. *Software Product Lines: Practices and Patterns*. Addison-Wesley, 2002.
20. C. Clifton, T. Millstein, G. Leavens, and C. Chambers. MultiJava: Design Rationale, Compiler Implementation, and Applications. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 28(3):517–575, 2006.
21. K. Czarnecki and U. Eisenecker. *Generative Programming: Methods, Tools, and Applications*. Addison-Wesley, 2000.
22. K. Czarnecki and K. Pietroszek. Verifying Feature-Based Model Templates Against Well-Formedness OCL Constraints. In *Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE)*, pages 211–220. ACM Press, 2006.
23. B. Delaware, W. Cook, and D. Batory. A Machine-Checked Model of Safe Composition. In *Proceedings of the International Workshop on Foundations of Aspect-Oriented Languages (FOAL)*, pages 31–35. ACM Press, 2009.

24. S. Ducasse, O. Nierstrasz, N. Schärli, R. Wuyts, and A. Black. Traits: A Mechanism for Fine-Grained Reuse. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 28(2):331–388, 2006.
25. E. Ernst, K. Ostermann, and W. Cook. A Virtual Class Calculus. In *Proceedings of the International Symposium on Principles of Programming Languages (POPL)*, pages 270–282. ACM Press, 2006.
26. M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and Mixins. In *Proceedings of the International Symposium on Principles of Programming Languages (POPL)*, pages 171–183. ACM Press, 1998.
27. V. Gasiunas, M. Mezini, and K. Ostermann. Dependent Classes. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 133–152. ACM Press, 2007.
28. J. Gosling, B. Joy, G. Steele, and G. Bracha. *The Java Language Specification*. The Java Series. Addison-Wesley, 3rd edition, 2005.
29. R. Hirschfeld, P. Costanza, and O. Nierstrasz. Context-Oriented Programming. *Journal of Object Technology (JOT)*, 7(3):125–151, 2008.
30. D. Hutchins. Eliminating Distinctions of Class: Using Prototypes to Model Virtual Classes. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 1–19. ACM Press, 2006.
31. D. Hutchins. *Pure Subtype Systems: A Type Theory For Extensible Software*. PhD thesis, School of Informatics, University of Edinburgh, 2008.
32. A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A Minimal Core Calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 23(3):396–450, 2001.
33. A. Igarashi, C. Saito, and M. Viroli. Lightweight Family Polymorphism. In *Proceedings of the Asian Symposium on Programming Languages and Systems (APLAS)*, volume 3780 of LNCS, pages 161–177. Springer-Verlag, 2005.
34. T. Kamina and T. Tamai. McJava – A Design and Implementation of Java with Mixin-Types. In *Proceedings of the Asian Symposium on Programming Languages and Systems (APLAS)*, volume 3302 of LNCS, pages 398–414. Springer-Verlag, 2004.
35. K. Kang, S. Cohen, J. Hess, W. Novak, and A. Peterson. Feature-Oriented Domain Analysis (FODA) Feasibility Study. Technical Report CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie Mellon University, 1990.
36. C. Kästner and S. Apel. Type-Checking Software Product Lines – A Formal Approach. In *Proceedings of the International Conference on Automated Software Engineering (ASE)*, pages 258–267. IEEE CS Press, 2008.
37. C. Kästner, S. Apel, and D. Batory. A Case Study Implementing Features using AspectJ. In *Proceedings of the International Software Product Line Conference (SPLC)*, pages 222–232. IEEE CS Press, 2007.
38. C. Kästner, S. Apel, and M. Kuhlemann. Granularity in Software Product Lines. In *Proceedings of the International Conference on Software Engineering (ICSE)*, pages 311–320. ACM Press, 2008.
39. C. Kästner, S. Apel, S. Trujillo, M. Kuhlemann, and D. Batory. Guaranteeing Syntactic Correctness for all Product Line Variants: A Language-Independent Approach. In *Proceedings of the International Conference on Objects, Models, Components, Patterns (TOOLS EUROPE)*, volume 33 of LNBI, pages 174–194. Springer-Verlag, 2009.
40. G. Lagorio, M. Servetto, and E. Zucca. Featherweight Jigsaw – A Minimal Core Calculus for Modular Composition of Classes. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, LNCS. Springer-Verlag, 2009.
41. L. Liquori and A. Spiwack. FeatherTrait: A Modest Extension of Featherweight Java. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):1–32, 2008.

42. R. Lopez-Herrejon and D. Batory. A Standard Problem for Evaluating Product-Line Methodologies. In *Proceedings of the International Conference on Generative and Component-Based Software Engineering (GCSE)*, volume 2186 of *LNCS*, pages 10–24. Springer-Verlag, 2001.
43. R. Lopez-Herrejon, D. Batory, and W. Cook. Evaluating Support for Features in Advanced Modularization Technologies. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 3586 of *LNCS*, pages 169–194. Springer-Verlag, 2005.
44. R. Lopez-Herrejon, D. Batory, and C. Lengauer. A Disciplined Approach to Aspect Composition. In *Proceedings of the International Symposium Partial Evaluation and Semantics-Based Program Manipulation (PEPM)*, pages 68–77. ACM Press, 2006.
45. O. Madsen and B. Moller-Pedersen. Virtual Classes: A Powerful Mechanism in Object-Oriented Programming. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 397–406. ACM Press, 1989.
46. H. Masuhara and G. Kiczales. Modeling Crosscutting in Aspect-Oriented Mechanisms. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 2743 of *LNCS*, pages 2–28. Springer-Verlag, 2003.
47. M. Mendonca, A. Wasowski, and K. Czarnecki. SAT-based Analysis of Feature Models is Easy. In *Proceedings of the International Software Product Line Conference (SPLC)*. Software Engineering Institute, Carnegie Mellon University, 2009.
48. M. Mezini and K. Ostermann. Variability Management with Feature-Oriented Programming and Aspects. In *Proceedings of the International Symposium on Foundations of Software Engineering (FSE)*, pages 127–136. ACM Press, 2004.
49. G. Murphy, A. Lai, R. Walker, and M. Robillard. Separating Features in Source Code: An Exploratory Study. In *Proceedings of the International Conference on Software Engineering (ICSE)*, pages 275–284. IEEE CS Press, 2001.
50. N. Nystrom, S. Chong, and A. Myers. Scalable Extensibility via Nested Inheritance. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 99–115. ACM Press, 2004.
51. M. Odersky, V. Cremet, C. Röckl, and M. Zenger. A Nominal Theory of Objects with Dependent Types. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 2743 of *LNCS*, pages 201–224. Springer-Verlag, 2003.
52. M. Odersky and M. Zenger. Scalable Component Abstractions. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 41–57. ACM Press, 2005.
53. K. Ostermann. Dynamically Composable Collaborations with Delegation Layers. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 2374 of *LNCS*, pages 89–110. Springer-Verlag, 2002.
54. B. Pierce. *Types and Programming Languages*. MIT Press, 2002.
55. C. Prehofer. Feature-Oriented Programming: A Fresh Look at Objects. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 1241 of *LNCS*, pages 419–443. Springer-Verlag, 1997.
56. T. Reenskaug, E. Andersen, A. Berre, A. Hurlen, A. Landmark, O. Lehne, E. Nordhagen, E. Ness-Ulseth, G. Oftedal, A. Skaar, and P. Stenslet. OORASS: Seamless Support for the Creation and Maintenance of Object-Oriented Systems. *Journal of Object-Oriented Programming (JOOP)*, 5(6):27–41, 1992.
57. M. Rosenmüller, C. Kästner, N. Siegmund, S. Sunkle, S. Apel, T. Leich, and G. Saake. SQL à la Carte – Toward Tailor-made Data Management. In *Datenbanksysteme in Business, Technologie und Web – Fachtagung des GI-Fachbereichs Datenbanken und Informationssysteme*, volume P-144 of *GI-Edition – LNI*, pages 117–136. Gesellschaft für Informatik, 2009.

58. M. Rosenmüller, N. Siegmund, G. Saake, and S. Apel. Code Generation to Support Static and Dynamic Composition of Software Product Lines. In *Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE)*, pages 3–12. ACM Press, 2008.
59. M. Rosenmüller, N. Siegmund, H. Schirmeier, J. Sincero, S. Apel, T. Leich, O. Spinczyk, and G. Saake. FAME-DBMS: Tailor-made Data Management Solutions for Embedded Systems. In *Proceedings of the EDBT Workshop on Software Engineering for Tailor-made Data Management (SETMDM)*, pages 1–6. ACM Press, 2008.
60. N. Siegmund, C. Kästner, M. Rosenmüller, F. Heidenreich, S. Apel, and G. Saake. Bridging the Gap between Variability in Client Application and Database Schema. In *Datenbanksysteme in Business, Technologie und Web – Fachtagung des GI-Fachbereichs Datenbanken und Informationssysteme*, volume P-144 of *GI-Edition – LNI*, pages 297–306. Gesellschaft für Informatik, 2009.
61. Y. Smaragdakis and D. Batory. Mixin Layers: An Object-Oriented Implementation Technique for Refinements and Collaboration-Based Designs. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(2):215–255, 2002.
62. R. Strniša, P. Sewell, and M. Parkinson. The Java Module System: Core Design and Semantic Definition. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 499–514. ACM Press, 2007.
63. P. Tarr, H. Ossher, W. Harrison, and S. Sutton, Jr. N Degrees of Separation: Multi-Dimensional Separation of Concerns. In *Proceedings of the International Conference on Software Engineering (ICSE)*, pages 107–119. IEEE CS Press, 1999.
64. S. Thaker, D. Batory, D. Kitchin, and W. Cook. Safe Composition of Product Lines. In *Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE)*, pages 95–104. ACM Press, 2007.
65. M. VanHilst and D. Notkin. Using Role Components in Implement Collaboration-based Designs. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 359–369. ACM Press, 1996.
66. A. K. Wright and M. Felleisen. A Syntactic Approach to Type Soundness. *Information and Computation*, 115(1):38–94, 1994.

A Type Soundness Proof of FFJ

Before giving the main proof, we state and prove some required lemmas.

LEMMA A.1 If $mtype(m, last(D)) = \bar{C} \rightarrow C_0$, then $mtype(m, last(C)) = \bar{C} \rightarrow C_0$ for all $C <: D$.

Proof Straightforward induction on the derivation of $C <: D$. There are two cases: First, if method m is not defined in the declaration or in any refinement of class C , then $mtype(m, last(C))$ should be the same as $mtype(m, last(E))$ where $CT(\Phi.C) = \text{class } C \text{ extends } E \{ \dots \}$ for some Φ . This follows from the definition of $mtype$ that searches E 's refinement chain from right to left if m is not declared in C 's refinement chain. Second, if m is defined in the declaration or in any refinement of class C , then $mtype(m, last(C))$ should also be the same as $mtype(m, last(E))$ with $CT(\Phi.C) = \text{class } C \text{ extends } E \{ \dots \}$ for some Φ . This case is covered by the well-formedness rules for methods that use the predicate *override* to ensure that m is properly overridden, i.e., the signatures of the overridden and the overriding declaration of m are equal, and that m is not introduced twice, i.e., overloading is not allowed in FFJ. \square

LEMMA A.2 (*Term substitution preserves typing*) If $\Gamma, \overline{x} : \overline{B} \vdash t : D$ and $\Gamma, \overline{s} : \overline{A}$, where $\overline{A} <: \overline{B}$, then $\Gamma \vdash [\overline{x} \mapsto \overline{s}] t : C$ for some $C <: D$.

Proof By induction on the derivation of $\Gamma, \overline{x} : \overline{B} \vdash t : D$.

CASE (T-VAR) $t = x \quad x : D \in \Gamma$

If $x \notin \overline{x}$, then the result is trivial since $[\overline{x} \mapsto \overline{s}] x = x$.¹⁰ On the other hand, if $x = x_i$ and $D = B_i$, then, since $[\overline{x} \mapsto \overline{s}] x = s_i$, letting $C = A_i$ finishes the case.

CASE (T-FIELD) $t = t_0.f_i \quad \Gamma, \overline{x} : \overline{B} \vdash t_0 : D_0 \quad \text{fields}(\text{last}(D_0)) = \overline{C} \overline{f} \quad D = C_i$

By the induction hypothesis, there is some C_0 such that $\Gamma \vdash [\overline{x} \mapsto \overline{s}] t_0 : C_0$ and $C_0 <: D_0$. It is easy to check that $\text{fields}(\text{last}(C_0)) = (\text{fields}(\text{last}(D_0)), \overline{D} \overline{g})$ for some $\overline{D} \overline{g}$. Therefore, by T-FIELD, $\Gamma \vdash ([\overline{x} \mapsto \overline{s}] t_0).f_i : C_i$. The fact that the refinements of a class may add new fields does not cause problems. $\overline{D} \overline{g}$ contains all fields that C_0 , including all of its refinements, add to D_0 .

CASE (T-INVK) $t = t_0.m(\overline{t}) \quad \Gamma, \overline{x} : \overline{B} \vdash t_0 : D_0 \quad \text{mtype}(m, \text{last}(D_0)) = \overline{E} \rightarrow D$
 $\Gamma, \overline{x} : \overline{B} \vdash \overline{t} : \overline{D} \quad \overline{D} <: \overline{E}$

By the induction hypothesis, there are some C_0 and \overline{C} such that:

$$\Gamma \vdash [\overline{x} \mapsto \overline{s}] t_0 : C_0 \quad C_0 <: D_0 \quad \Gamma \vdash [\overline{x} \mapsto \overline{s}] \overline{t} : \overline{C} \quad \overline{C} <: \overline{D}.$$

By Lemma A.1, we have $\text{mtype}(m, \text{last}(C_0)) = \overline{E} \rightarrow D$. Moreover, $\overline{C} <: \overline{E}$ by the transitivity of $<:$. Therefore, by T-INVK, $\Gamma \vdash [\overline{x} \mapsto \overline{s}] t_0.m([\overline{x} \mapsto \overline{s}] \overline{t}) : D$. The key is that subclasses and refinements may override methods but the well-formedness rules of methods ensure that the method's type is not altered, i.e., there is no overloading in FFJ.

CASE (T-NEW) $t = \text{new } D(\overline{t}) \quad \text{fields}(\text{last}(D)) = \overline{D} \overline{f} \quad \Gamma, \overline{x} : \overline{B} \vdash \overline{t} : \overline{C} \quad \overline{C} <: \overline{D}$

By the induction hypothesis, $\Gamma \vdash [\overline{x} \mapsto \overline{s}] \overline{t} : \overline{E}$ for some \overline{E} with $\overline{E} <: \overline{C}$. We have $\overline{E} <: \overline{D}$ by the transitivity of $<:$. Therefore, by rule T-NEW, $\Gamma \vdash \text{new } D([\overline{x} \mapsto \overline{s}] \overline{t}) : D$. Although refinements of class D may add new fields, rule T-NEW ensures that the arguments of the object creation match the overall fields of D , including all refinements, in number and types. That is, the number of arguments (\overline{t}) equals the number of fields (\overline{f}) which function fields returns.

CASE (T-UCAST) $t = (D)t_0 \quad \Gamma, \overline{x} : \overline{B} \vdash t_0 : C \quad C <: D$

By the induction hypothesis, there is some E such that $\Gamma \vdash [\overline{x} \mapsto \overline{s}] t_0 : E$ and $E <: C$. We have $E <: D$ by the transitivity of $<:$, which yields $\Gamma \vdash (D)([\overline{x} \mapsto \overline{s}] t_0) : D$ by T-UCAST.

CASE (T-DCAST) $t = (D)t_0 \quad \Gamma, \overline{x} : \overline{B} \vdash t_0 : C \quad D <: C \quad D \neq C$

¹⁰ Note that $[\overline{x} \mapsto \overline{s}] x$ is an abbreviation for $[x_1 \mapsto s_1, \dots, x_n \mapsto s_n] x$. It means that all occurrences of the variables x_1, \dots, x_n in the term x are substituted with the corresponding terms s_1, \dots, s_n .

By the induction hypothesis, there is some E such that $\Gamma \vdash [\overline{x \mapsto s}] t_0 : E$ and $E <: C$. If $E <: D$ or $D <: E$, then $\Gamma \vdash (D)([\overline{x \mapsto s}] t_0) : D$ by T-UCAST or T-DCAST, respectively. If both $D \not<: E$ and $E \not<: D$, then $\Gamma \vdash (D)([\overline{x \mapsto s}] t_0) : D$ (with a *stupid warning*) by T-SCAST.

CASE (T-SCAST) $t = (D)t_0 \quad \Gamma, \overline{x : B} \vdash t_0 : C \quad D \not<: C \quad C \not<: D$

By the induction hypothesis, there is some E such that $\Gamma \vdash [\overline{x \mapsto s}] t_0 : E$ and $E <: C$. This means that $E \not<: D$ because, in FFJ, each class has just one superclass and, if both $E <: C$ and $E <: D$, then either $C <: D$ or $D <: C$, which contradicts the induction hypothesis. So $\Gamma \vdash (D)([\overline{x \mapsto s}] t_0) : D$ (with a *stupid warning*), by T-SCAST. \square

LEMMA A.3 (*Weakening*) If $\Gamma \vdash t : C$, then $\Gamma, x : D \vdash t : C$

Proof Straightforward induction. The proof for FFJ is similar to the proof for FJ. \square

LEMMA A.4 If $mtype(m, last(C_0)) = \overline{D} \rightarrow D$, and $mbody(m, last(C_0)) = (\overline{x}, t)$, then for some D_0 and some $C <: D$ we have $C_0 <: D_0$ and $\overline{x} : \overline{D}, this : D_0 \vdash t : C$.

Proof By induction on the derivation of $mbody(m, last(C_0))$. The base case (in which m is defined in the most specific refinement of C_0) is easy since m is defined in $CT(last(C_0))$ and the well-formedness of the class table implies that we must have derived $\overline{x} : \overline{D}, this : C_0 \vdash t : C$ by the well-formedness rules of methods. The induction step is also straightforward: if m is not defined in $CT(last(C_0))$, then $mbody$ searches the refinement chain from right to left; if m has not been found, the superclass' refinement chain is searched. There are two subcases: first, m is defined in the declaration or in any refinement of C_0 ; this case is similar to the base case. Second, m is defined in a superclass D_0 of C_0 or in one of D_0 's refinements; in this case, the well-formedness of the class table implies that we must have derived $\overline{x} : \overline{D}, this : D_0 \vdash t : C$ by the well-formedness rules of methods, which finishes the case. \square

Note that this lemma holds because method refinements do not change the types of the arguments and the result of a method, overloading is not allowed, and `this` points always to the class that is introduced or refined.

THEOREM A.1 (*Preservation*) If $\Gamma \vdash t : C$ and $t \longrightarrow t'$, then $\Gamma \vdash t' : C'$ for some $C' <: C$.

Proof By induction on a derivation of $t \longrightarrow t'$, with a case analysis on the final rule.

CASE (E-PROJNEW) $t = \text{new } C_0(\overline{v}).f_i \quad t' = v_i \quad fields(last(C_0)) = \overline{D} f$

From the shape of t , we see that the final rule in the derivation of $\Gamma \vdash t : C$ must be T-FIELD, with premise $\Gamma \vdash \text{new } C_0(\overline{v}) : D_0$, for some D_0 , and that $C = D_i$. Similarly, the last rule in the derivation of $\Gamma \vdash \text{new } C_0(\overline{v}) : D_0$ must be T-NEW, with premises $\Gamma \vdash \overline{v} : C$ and $C <: D$, and with $D_0 = C_0$. In particular, $\Gamma \vdash v_i : C_i$, which finishes the case, since $C_i <: D_i$.

CASE (E-INVKNEW) $t = (\text{new } C_0(\overline{v})).m(\overline{u}) \quad t' = [\overline{x \mapsto \overline{u}}, this \mapsto \text{new } C_0(\overline{v})] t_0$
 $mbody(m, last(C_0)) = (\overline{x}, t_0)$

The final rules in the derivation of $\Gamma \vdash t : C$ must be T-INVK and T-NEW, with premises $\Gamma \vdash \text{new } C_0(\bar{v}) : C_0$, $\Gamma \vdash \bar{u} : \bar{C}$, $\bar{C} <: \bar{D}$, and $mtype(m, last(C_0)) = \bar{D} \rightarrow C$. By Lemma A.4, we have $\bar{x} : \bar{D}, \text{this} : D_0 \vdash t : B$ for some D_0 and B , with $C_0 <: D_0$ and $B <: C$. By Lemma A.3, $\Gamma, \bar{x} : \bar{D}, \text{this} : D_0 \vdash t_0 : B$. Then, by Lemma A.2, we have $\Gamma [\bar{x} \mapsto \bar{u}, \text{this} \mapsto \text{new } C_0(\bar{v})] t_0 : E$ for some $E <: B$. By the transitivity of $<:$, we obtain $E <: C$. Letting $C' = E$ completes the case.

CASE (E-CASTNEW) $t = (D)(\text{new } C_0(\bar{v}))$ $C_0 <: D$ $t' = \text{new } C_0(\bar{v})$

The proof of $\Gamma \vdash (D)(\text{new } C_0(\bar{v})) : C$ must end with T-UCAST since ending with T-SCAST or T-DCAST would contradict the assumption of $C_0 <: D$. The premises of T-UCAST give us $\Gamma \vdash \text{new } C_0(\bar{v}) : C_0$ and $D = C$, finishing the case.

The cases for the congruence rules are easy. We show just the case E-CAST.

CASE (E-CAST) $t = (D)t_0$ $t' = (D)t'_0$ $t_0 \longrightarrow t'_0$

There are three subcases according to the last typing rule used.

SUBCASE (T-UCAST) $\Gamma \vdash t_0 : C_0$ $C_0 <: D$ $D = C$

By the induction hypothesis, $\Gamma \vdash t'_0 : C'_0$ for some $C'_0 <: C_0$. By transitivity of $<:$, $C'_0 <: C$. Therefore, by T-UCAST, $\Gamma \vdash (C)t'_0 : C$ (with no additional *stupid warning*).

SUBCASE (T-DCAST) $\Gamma \vdash t_0 : C_0$ $D <: C_0$ $D = C$

By the induction hypothesis, $\Gamma \vdash t'_0 : C'_0$ for some $C'_0 <: C_0$. If $C'_0 <: C$ or $C <: C'_0$, then $\Gamma \vdash (C)t'_0 : C$ by T-UCAST or T-DCAST (without any additional *stupid warning*). On the other hand, if both $C'_0 \not<: C$ or $C \not<: C'_0$, then $\Gamma \vdash (C)t'_0 : C$ with a *stupid warning* by T-SCAST.

SUBCASE (T-SCAST) $\Gamma \vdash t_0 : C_0$ $D \not<: C_0$ $C_0 \not<: D$ $D = C$

By the induction hypothesis, $\Gamma \vdash t'_0 : C'_0$ for some $C'_0 <: C_0$. Then, also $C'_0 \not<: C$ and $C \not<: C'_0$. Therefore $\Gamma \vdash (C)t'_0 : C$ with a *stupid warning*. If $C'_0 \not<: C$, then $C \not<: C'_0$ since $C \not<: C_0$ and, therefore, $\Gamma \vdash (C)t'_0 : C$ with *stupid warning*. If $C'_0 <: C$, then $\Gamma \vdash (C)t'_0 : C$ by T-UCAST (with no additional *stupid warning*). This subcase is analogous to the case T-SCAST of the proof of Lemma A.2. \square

THEOREM A.2 (*Progress*) Suppose t is a well-typed term.

1. If t includes $\text{new } C_0(\bar{t}).f_i$ as a subterm, then $fields(last(C_0)) = \bar{C} \bar{f}$ for some \bar{C} and \bar{f} .
2. If t includes $\text{new } C_0(\bar{t}).m(\bar{u})$ as a subterm, then $mbody(m, last(C_0)) = (\bar{x}, t_0)$ and $|\bar{x}| = |\bar{u}|$ for some \bar{x} and t_0 .

Proof If t has $\text{new } C_0(\bar{t}).f_i$ as a subterm, then, by well-typedness of the subterm, it is easy to check that $fields(last(C_0))$ is well-defined and f_i appears in it. The fact that refinements may add fields (that have not been defined already) does not invalidate this conclusion. Note that for every field of a class, including its superclasses and all its refinements, there must be a proper argument. Similarly, if t has $\text{new } C_0(\bar{t}).m(\bar{u})$ as a subterm, then it is also easy to show that $mbody(m, last(C_0)) = (\bar{x}, t_0)$ and $|\bar{x}| = |\bar{u}|$ from the fact that $mtype(m, last(C_0)) = \bar{C} \rightarrow D$ where $|\bar{x}| = |\bar{C}|$. This conclusion holds for FFJ since a method refinement must have the same signature than the method refined and overloading is not allowed. \square

THEOREM A.3 (*Type soundness of FFJ*) If $\emptyset \vdash t : C$ and $t \longrightarrow^* t'$ with t' a normal form, then t' is either a value v with $\emptyset \vdash v : D$ and $D <: C$, or a term containing $(D)(\text{new } C(\bar{t}))$ in which $C <: D$.

Proof Immediate from Theorem A.1 and A.2. Nothing changes in the proof of Theorem A.3 for FFJ compared to FJ. \square

B Type Soundness Proof of FFJ_{PL}

In this section, we provide proof sketches of the theorems *Correctness of FFJ_{PL}* and *Completeness of FFJ_{PL}*. A further formalization would be desirable, but we have stopped at this point. As is often the case with formal systems, there is a trade-off between formal precision and legibility. We decided that a semi-formal development of the proof strategies are the best fit for our purposes.

B.1 Correctness

THEOREM B.1 (*Correctness of FFJ_{PL}*) Given a well-typed FFJ_{PL} product line pl (including with a well-typed term t , well-formed class, introduction, and refinement tables CT , IT , and RT , and a feature model FM), every program that can be derived with a valid feature selection fs is a well-typed FFJ program (cf. Figure 10).

$$\frac{pl = (t, CT, IT, RT, FM) \quad pl \text{ is well-typed} \quad fs \text{ is valid in } FM}{\text{derive}(pl, fs) \text{ is well-typed}}$$

The proof strategy is as follows: assuming that the FFJ_{PL} type system ensures that each slice is a valid FFJ type derivation (Lemma B.1) and that each valid feature selection corresponds to a single slice (Lemma B.1), it follows that the corresponding feature-oriented program is well-formed. Before we prove Theorem B.1 we develop two required lemmas that cover the two assumptions of our proof strategy.

LEMMA B.1 Given a well-formed FFJ_{PL} product line, every slice of the product line's type derivation corresponds to a (set of) valid type derivation(s) in FFJ.

Proof (Proof sketch) Given a well-formed FFJ_{PL} product line, the corresponding type derivation consists of possibly multiple slices.

The basic case is easy: there is only a simple derivation without branches due to mutually exclusive features (optional features may be present). In this case, each term has only a single type, which is the one that would also be determined by FFJ. Furthermore, FFJ_{PL} guarantees that referenced types, methods, and fields are present in all valid variants, using the predicate *validref*.

Let us illustrate this with the rule T-FIELD_{PL}; the other rules are analogous:

$$\frac{\forall E \in \bar{E} : \text{validref}(\Phi, E.f) \quad \Gamma \vdash t_0 : \bar{E} \dashv \Phi \quad \text{fields}(\Phi, \overline{\text{last}(\bar{E})}) = \overline{\mathcal{F}, \mathcal{C}f, \mathcal{G}}}{\Gamma \vdash t_0.f : C_{11}, \dots, C_{n1}, \dots, C_{1m}, \dots, C_{nm} \dashv \Phi} \quad (\text{T-FIELD}_{PL})$$

In the basic case there are no branches in the type derivation and thus the term t_0 has only a single type E_1 . For the same reason, *fields* returns only a simple list of fields that contains the declaration of field f . Finally, $T\text{-FIELD}_{PL}$ checks whether the declaration of f is present in all valid variants (using *validref*). Hence, in the basic case, an FFJ_{PL} derivation that ends at the rule $T\text{-FIELD}_{PL}$ is equivalent to a set of corresponding FFJ derivations, which do not contain alternative and optional features and thus t_0 has a single type, *fields* returns a simple list of fields that contains the declaration of f , and the declaration of f is present. The reason that an FFJ_{PL} derivation without mutually exclusive features (i.e., a single slice) corresponds to multiple FFJ derivations is that the FFJ_{PL} derivation may contain optional features whose different combinations correspond to the different FFJ derivations. Using predicate *validref*, all type rules of FFJ_{PL} ensure that all possible combinations of optional features are well-typed.

In the case that there are multiple slices in the FFJ_{PL} derivation, a term t_0 may have multiple types \bar{E} . The type rules of FFJ_{PL} make sure that every possible shape of a given term is well-typed. Each possible type of the term leads to a branch in the derivation tree. The premise of $T\text{-FIELD}_{PL}$ checks whether all possible shapes of a given term are well-typed by taking the conjunction of all branches of the derivation. Hence, if $T\text{-FIELD}_{PL}$ is successful, each individual branch holds, i.e., each slice corresponds to a well-typed FFJ program. Ensuring that, in the presence of optional features, all relevant subterms are well-typed (i.e., all referenced elements are present in all valid variants), a well-typed slice covers a set of well-typed FFJ derivations that correspond to different combinations of optional features, like in the basic case.

For example, in a field projection $t_0.f$, the subterm t_0 has multiple types \bar{E} . For all these types, *fields* yields all possible combinations of fields declared by the variants of the types. It is checked whether, for each type of the subterm t_0 , each combination of fields contains a proper declaration of field f . The different types of f become the possible types of the overall field projection term. Like in the basic case, it is checked whether every possible type of t_0 is present in all valid variants (using *validref*), so that each slice corresponds a valid FFJ derivation, i.e., a whole set of derivations covering different combinations of optional features. \square

LEMMA B.2 Given a well-formed FFJ_{PL} product line, each valid feature selection corresponds to a single slice in the corresponding type derivation.

Proof (Proof sketch) By definition, a valid feature selection does not contain mutually exclusive features. Considering only a single valid feature selection, each term has only a single type. But the type derivation of the overall product line contains branches corresponding to alternative types of the terms. A successive removal of mutually exclusive features removes these branches until only a single branch remains. Consequently, a valid feature selection corresponds to a single slice. \square

Proof (Proof sketch of Theorem B.1 (Correctness of FFJ_{PL})) The fact that the FFJ_{PL} type system ensures that each slice is a valid FFJ type derivation (Lemma B.1) and that each valid feature selection corresponds to a single slice (Lemma B.2), implies that each feature-oriented program that corresponds to a valid feature selection is well-formed. \square

B.2 Completeness

THEOREM B.2 (*Completeness of FFJ_{PL}*) Given an FFJ_{PL} product line pl (including a well-typed term t , well-formed class, introduction, and refinement tables CT , IT , and RT , and a feature model FM), and given that *all* valid feature selections fs yield well-typed FFJ programs, according to Theorem B.1, pl is a well-typed product line according to the rules of FFJ_{PL}.

$$\frac{pl = (t, CT, IT, RT, FM) \quad \forall fs : (fs \text{ is valid in } FM \Rightarrow \text{derive}(pl, fs) \text{ is well-typed})}{pl \text{ is well-typed}}$$

Proof (Proof sketch of Theorem B.2 (Completeness of FFJ_{PL})) There are three basic cases: (1) pl has only mandatory features; (2) pl has only mandatory features except a single optional feature; (3) pl has only mandatory features except two mutually exclusive features. Proving Theorem B.2 for the first basic case is trivial. Since only mandatory features exist, only a single FFJ program can be derived from the product line. If the FFJ program is well-typed, the product line is well-typed, too, because all elements are always reachable and each term has only a single type. In fact, the type rules of FFJ_{PL} and FFJ become equivalent in this case.


In the second basic case, two FFJ programs can be derived from the product line, one including and one excluding the optional feature. The difference between the two programs is the content of the optional feature. The feature can add new classes, refine existing classes by new methods and fields, and refine existing methods by overriding. If the two programs are well-typed, then the overall product line is well-typed as well since the reachability checks succeed in every type rule of FFJ_{PL}. Otherwise, at least one of the two programs would not be well-typed since, in this case, the reachability checks are the only difference between FFJ_{PL}'s and FFJ's type rules (as in the first case, each term has only a single type since there are no mutually exclusive features). The fact that the two FFJ programs are well-typed implies that all elements are reachable in the type derivations of two FFJ programs. Thus, the reachability checks of the FFJ_{PL} derivation succeed in every case, i.e., the product line in question is well-typed.

In the third basic case, two FFJ programs can be derived from the product line, one including the first alternative and the other including the second alternative of the feature in question. The difference between the two programs is, on the one hand, the program elements one feature introduces that are not present in the other and, on the other hand, the alternative definitions of similar elements, like two alternative definitions of a single class. The first kind of difference is already covered by the second basic case. Alternative definitions of a program element (second kind of difference) that are well-typed in the context of their enclosing FFJ programs, are well-typed in FFJ_{PL} because they lead to two new branches in the derivation tree which are handled separately and the conjunction of their premises must hold. Since the corresponding FFJ type rule for the element succeeds in both FFJ programs, their conjunction in the FFJ_{PL} type rule always holds, i.e., the product line in question is well-typed.

Finally, we it remains to show that all other cases, i.e., all other combinations of mandatory, optional, and alternative features, can be reduced to combinations of the three basic cases, which proves Theorem B.2. To this end, we divide the possible relations between features into three disjoint sets: (1) a feature is reachable from another

feature in *all* variants, (2) a feature is reachable from another feature in *some*, but *not in all*, variants, (3) two features are mutually exclusive. From these three possible relations we construct a general case that can be reduced to a combination of the three basic cases.

Assume a feature Φ that is mandatory with respect to a set of features $\overline{\Pi}$, that is optional with respect to a set of features $\overline{\Omega}$, and that is alternative to a set $\overline{\Delta}$ of features. We use arrows to illustrate to which of the three basic cases a pairwise relation between Φ and each element of a list is reduced:



Such an arrow diagram can be created for every feature of a product line. The reason is that the three kinds of relations are orthogonal and there are no further relations relevant for type checking. Hence, the general case covers all possible relations between features and combinations of features. The description of the general case and the reduction finish the proof of Theorem B.2, i.e., FFJ_{PL}'s type system is complete. \square