SFJ: An Implementation of Semantic Featherweight Java

Artem Usov and Ornela Dardha(✉)

School of Computing Science, University of Glasgow, Glasgow, UK
2296905U@student.gla.ac.uk, ornela.dardha@glasgow.ac.uk

Abstract. There are two approaches to defining subtyping relations: the *syntactic* and the *semantic* approach. In semantic subtyping, one defines a model of the language and an interpretation of types as subsets of this model. Subtyping is defined as inclusion of subsets denoting types.

An orthogonal subtyping question, typical of object-oriented languages, is the *nominal* versus the *structural* subtyping. Dardha *et al.* [11,12] defined boolean types and semantic subtyping for Featherweight Java (FJ) and integrated both nominal and structural subtyping, thus exploiting the benefits of both approaches. However, these benefits were illustrated only at a theoretical level, but not exploited practically.

We present SFJ—Semantic Featherweight Java, an implementation of FJ which features boolean types, semantic subtyping and integrates nominal as well as structural subtyping. The benefits of SFJ, illustrated in the paper and the accompanying video (with audio/subtitles) [27], show how static type-checking of boolean types and semantic subtyping gives higher guarantees of program correctness, more flexibility and compactness of program writing.

Keywords: Nominal subtyping · Structural subtyping · Semantic Featherweight Java · Object-oriented languages · Boolean types · Type theory

1 Introduction

There are two approaches to defining subtyping relations: the *syntactic* and the *semantic* approach. Syntactic subtyping [20] is more mainstream in programming languages and is defined by means of a set of formal deductive subtyping rules. Semantic subtyping [1,10] is more recent and less known: one defines a formal model of the language and an interpretation of types as subsets of this model. Then, subtyping is defined as set inclusion of subsets denoting types.

Orthogonally, for object-oriented languages there are two approaches to defining subtyping relations: the *nominal* and the *structural* approach [21,22]. Nominal subtyping is based on declarations by the developer and is *name*-based: "A is a subtype of B if and only if it is declared to be so, that is if the class A extends (or implements) the class (or interface) B". Structural subtyping instead is based on the *structure* of a class, its fields and methods: "a class A is a subtype of a class B if and only if the fields and methods of A are a superset of the fields and methods of B, and their types in A are subtypes of the types in B". For example, the set of inhabitants of a class *Student* is smaller than the set of inhabitants of a class *Person*, as each Student is a Person, but not the other way around. However, the set of fields and methods of *Student* is a superset of that of *Person*. Hence, *Student* is a structural subtype of *Person*, even if it is not declared so.

Dardha *et al.* [11,12] define boolean types—based on set-theoretic operations such as **and**, **not**, **or**—and semantic subtyping for *Featherweight Java* (FJ) [17]. This approach allows for the integration of both nominal and structural subtyping in FJ, bringing in higher guarantees of program correctness, flexibility and compactness in program writing. Unfortunately, these benefits were only presented at a theoretical level and not exploited practically, due to the lack of an implementation of the language, its types and type system.

In this paper, we present SFJ—*Semantic Featherweight Java* Sect. 3, an implementation of FJ with boolean types and semantic subtyping. In SFJ the developer has a larger and more expressive set of types, by using boolean connectives **and**, **not**, **or**, with the expected set-theoretic interpretation. On the other hand, this added expressivity does not add complexity. Rather the opposite is true, as the developer has an easier, more compact and elegant way of programming. SFJ integrates both structural and nominal subtyping, and the developer can choose which one to use. Finally, as discussed in Dardha *et al.* [12, Sect. 8.4], thanks to semantic subtyping, we can easily encode in SFJ standard programming constructs and features of the full Java language, such as lists, or overloading classes via multimethods [5], which are missing in FJ, thus making SFJ a more complete language closer to Java.

Example 1 (Polygons). This will be our running example both in the paper and in the tool video [27] to illustrate the benefits of boolean types and semantic subtyping developed by Dardha *et al.* [11,12] and implemented as SFJ.

Consider the set of polygons, such as triangles, squares and rhombuses given by a class hierarchy. We want to define a method *diagonal* that takes a polygon and returns the length of its longest diagonal. This method makes sense only if the polygon passed to it has at least four sides, hence triangles are excluded. In Java this could be implemented in the following ways:

```
class Polygon {...}
class Triangle extends Polygon {...}
class Other_Polygons extends Polygon {
    double diagonal(Other_Polygons shape) {...}
                    ...
}
class Square extends Other_Polygons {...}
class Rhombus extends Other_Polygons {...}
```

Or by means of an interface *Diagonal*:

```
public interface Diagonal {
    double diagonal(Polygon shape);
}
class Polygon {...}
class Triangle extends Polygon {...}
class Square extends Polygon  implements Diagonal {...}
class Rhombus extends Polygon  implements Diagonal {...}
              // other polygons ...
```

Now, suppose our class hierarchy is such that *Polygon* is the parent class and all other geometric figures extend *Polygon*, which is how one would naturally define the set of polygons. Suppose the class hierarchy is given and is part of legacy code, which cannot be changed. Then again, a natural way to implement this in Java is by defining the method *diagonal* in the class *Polygon* and using an **instanceof**, for example, inside a **try-catch** construct. Then, an exception would be thrown at run time, if the argument passed to the method is a triangle.

We propose a more elegant solution, by combining boolean types and semantic subtyping, where only static type-checking is required and we implement this in SFJ [27]: it is enough to define a method *diagonal* that has an argument of type *Polygon* **and not** *Triangle*, thus allowing the type-checker to check at compile time the restrictions on types:

```
class Polygon {...}
class Triangle extends Polygon {...}
class Square extends Polygon {...}
class Rhombus extends Polygon {...}
              ...
class Diagonal {
              ...
    double diagonal((Polygon and not  Triangle) shape){...}
}
```

We can now call *diagonal* on an argument of type *Polygon*: if the polygon is **not** a *Triangle*, then the method computes and returns the length of its longest diagonal; otherwise, there will be a type error at compile time.

Structure of the Paper: In Sect. 2 we present the types and terms of the SFJ language. In Sect. 3 we present the design and implementation of SFJ; we discuss our two main algorithms, Algorithm 1 in Sect. 3.1 which checks the validity of type definitions, and Algorithm 2 in Sect. 3.2 which generates the semantic subtyping relation. Further, we discuss typing of SFJ in Sect. 3.3; nominal vs. structural subtyping in Sect. 3.5; method types in Sect. 3.6; and code generation in Sect. 3.7. We discuss related work and conclude the paper in Sect. 4.

2   Background

The technical developments behind semantic subtyping and its properties are complex, however, they are completely transparent to the programmer. The framework is detailed and proved correct in the relevant work by Dardha *et al.* [11,12], and SFJ builds on that framework.

In this section we will briefly detail the types and terms of SFJ.

2.1   Types

The syntax of types τ is defined by the following grammar:

$$
\begin{array}{lll}
\tau ::= \alpha \mid \mu & & \textit{Types} \\
\alpha ::= \mathbf{0} \mid \mathbb{B} \mid [\widetilde{l:\tau}] \mid \alpha \textbf{ and } \alpha \mid \textbf{not } \alpha & & \textit{Field types } (\alpha\textit{-types}) \\
\mu ::= \alpha \rightarrow \alpha \mid \mu \textbf{ and } \mu \mid \textbf{not } \mu & & \textit{Method types } (\mu\textit{-types})
\end{array}
$$

The α-types are used to type fields and the μ-types are used to type methods. Type $\mathbf{0}$ is the empty type. Type $\mathbb{B}$ denotes the *basic* types, such as integers, booleans, etc. Record types $[\widetilde{l:\tau}]$, where $\widetilde{l}$ is a sequence of disjoint labels, are used to type objects. Arrow types $\alpha \rightarrow \alpha$ are used to type methods.

The boolean connectives **and** and **not** in the α-types and μ-types have their expected set-theoretic meanings. We let $\alpha \setminus \alpha'$ denote α **and** (**not** α'), and α **or** α' denote **not**(**not** α **and** (**not** α')).

2.2   Terms

The syntax of terms is defined by the following grammar and is based on the standard syntax of terms in FJ [17]:

$$
\begin{array}{ll}
\textit{Class declaration} & L ::= \textbf{class } C \textbf{ extends } C \; \{\widetilde{\alpha\,a}; \; K; \; \widetilde{M} \} \\[4pt]
\textit{Constructor} & K ::= C \; (\widetilde{\alpha\,x}) \; \{ \; \textbf{super}(\widetilde{x}); \; \widetilde{\textbf{this}.a = \widetilde{x};} \} \\[4pt]
\textit{Method declaration} & M ::= \alpha \; m \; (\alpha \; x) \; \{ \; \textbf{return } e; \} \\[4pt]
\textit{Expressions} & e ::= x \mid c \mid e.a \mid e.m(e) \mid \textbf{new } C(\widetilde{e})
\end{array}
$$

We assume an infinite set of names, with some special names: *Object* denotes the root class, **this** denotes the current object and **super** denotes the parent object. We let $A, B, \ldots$ range over classes; $a, b, \ldots$ over fields; $m, n, \ldots$ over methods and $x, y, z, \ldots$ over variables.

A program $(\widetilde{L}, e)$ is a pair of a sequence of class declarations $\widetilde{L}$, giving rise to a class hierarchy as specified by the inheritance relation, and an expression e to be evaluated. A class declaration L specifies the name of the class, the name of the parent class it extends, its typed fields, the constructor K and its method declarations M. The constructor K initialises the fields of the object by assigning values to the fields inherited by the super-class and to the fields declared in the current **this** class. A method declaration M specifies the signature of the method, namely the return type, the method name and the formal parameter as well as the body of the method. Notice that in our theoretical development we use unary methods, without loss of generality: tuples of arguments can be modelled by an object that instantiates a "special" class containing as fields all the needed arguments. Expressions e include variables, constants, field access, method call and object creation.

Following FJ [17], we rule out ill-formed programs, such as declaring a constructor named B within a class named A, or multiple fields or methods having the same name, or fields having the same type as the class they are defined in.

3   The SFJ Language

3.1   On Valid Type Definitions

Since we want to use types τ to program in SFJ, we restrict them to *finite trees* whose leaves are basic types $\mathbb{B}$ Sect. 2.2 with no cycles. For example, a recursive type $\alpha = [a : \alpha]$ denotes an infinite program tree **new** $C(\textbf{new } C(\cdots))$, hence we rule it out as it is not inhabitable. Similarly, the types $\alpha = [b : \beta]$, $\beta = [a = \alpha]$ create a cycle and thus would not be inhabitable. Notice that these types can be defined and inhabited in Java by assigning *null* to all fields in a class, however they are not useful in practice.

SFJ is implemented using ANTLR [24]. We start by defining the grammar of the language in Extended Backus-Naur Form (EBNF), following Sect. 2.1 and by running ANTLR, we can automatically generate a parser for SFJ and extend it in order to implement the required checks for our types and type system.

Running the parser on an SFJ program returns an abstract syntax tree (AST) of that program. When visiting the AST, we check if the program is well-formed, following the intuition at the end of Sect. 2. We mark any classes containing only fields typed with basic types as *resolved*, otherwise we mark them as *unresolved*. Using this information, Algorithm 1 checks if the type definitions in a program are valid, namely if they are finite trees whose leaves are basic types with no

cycles. The algorithm returns $True$ only if all the types in the SFJ program are resolved, otherwise it returns $False$, meaning there is at least one type definition which is invalid and contains a cycle.

Algorithm 1: Validity Check for Type Definitions

Input   : $classes$, the set of classes in an SFJ program marked $resolved$, if their fields contain only basic types, $unresolved$ otherwise.

Output : $True$ if all classes are valid type definitions, $False$ otherwise.

```
 1 begin
 2    do
 3        resolutionOccured ⟵ False
 4        for class that is unresolved in classes do
 5            resolved ⟵ True
 6            for field in class that contains a class type do
 7                if type of field is unresolved then
 8                    resolved ⟵ False
 9                end
10            end
11
12            if resolved = True then
13                class ⟵ resolved
14                resolutionOccured ⟵ True
15            end
16        end
17    while resolutionOccured = True
18
19    if not all classes are resolved then
20        return False
21    else
22        return True
23    end
24 end
```

3.2 Building Semantic Subtyping for SFJ

If Algorithm 1 returns $True$, meaning all type definitions in a program are valid, we can then build the semantic subtyping. Leveraging the interpretation of types as sets of values to define semantic subtyping for FJ [11,12], in SFJ we keep track of the semantic subtyping relation by defining a map from a type to the set of its subtypes, satisfying the property that the set of values of a subtype is included in the set of values of the type.

Algorithm 2: Semantic Subtyping for SFJ—*generateRelation*

Input   : *classes*, the set of classes in an SFJ program.
             relation, the mapping of types to the set of subtypes, initially Map 3.1.

1 begin
2   **Function** *generateRelation(classes: List<Class>)*:
3   *unprocessed : List < Class > ⟵ []*
4   **for** *class in classes* **do**
5   **if** *addClass(class) = False* **then**
6   │ *unprocessed.add(class)*
7   **end**
8   **end**
9   **if** *untyped ≠ []* **then**
10   │ *generateRelation(unprocessed)*
11   **end**
12   **end**
13

    /* algorithm continued on next page... */

We start with basic types and let *Universe* be a supertype of all types. The full mapping for basic types is defined in Map 3.1.

$$
\begin{aligned}
Double &= \{Double, Float, Int, Short, Byte\} \quad Float = \{Float, Short, Byte\} \\
Long &= \{Long, Int, Short, Byte\} \qquad\qquad\; Int = \{Int, Short, Byte\} \\
Short &= \{Short, Byte\} \qquad\qquad\qquad\qquad Byte = \{Byte\} \\
Boolean &= \{Boolean\} \qquad\qquad\qquad\qquad\; Void = \{Void\} \\
Universe &= \{Double, Float, Long, Int, \\
&\qquad Short, Byte, Boolean, Void\}
\end{aligned}
$$

(3.1)

Note that *Int* is not a subtype of *Float* as a 32-bit *float* cannot represent the whole set of 32-bit *integer* values accurately and therefore *Int* is not fully set-contained in *Float*, however this is not the case for *Int* and *Double*. Similarly, *Long* is not a subtype of *Double*.

Algorithm 2 builds the semantic subtyping relation for all class types of an SFJ program by calling the function *generateRelation*. Given that classes are valid type definitions by Algorithm 1, we are guaranteed that Algorithm 2 will terminate.

The semantic subtyping generated by Algorithm 2 is a preorder: it is reflexive and transitive. This is also illustrated by Map 3.1.

Some comments on Algorithm 2 follow. In function *generateRelation* we iterate over the set of classes in an SFJ program. If the class currently being processed contains types in its fields or methods not present in the subtyping relation (lines 5 and 30, 42 in the continuation of the algorithm in the next page), then we add the current class to the list of *unprocessed* classes (line 6) so we can process its fields and methods first and the class itself later after having all required

---

Algorithm 2: Semantic Subtyping for SFJ—*addClass* and *checkSuperSet*

```
     /* ...algorithm continued from previous page */
13
14   Function addClass(class: Class) → boolean:
15       for existing class type in relation do
16           if checkSuperSet(class, existingClass) = False then
17               return False
18           end
19           checkSuperSet(existingClass, class)
20       end
21       relation[class].add(class)
22       relation[Universe].add(class)
23       return True
24   end
25
26   Function checkSuperSet(class: Class, other: Class) → boolean:
27       flag ⟵ True
28       for field in class do
29           if field contains type not in relation then
30               return False
31           end
32           if other does not contain field then
33               flag ⟵ False
34           else
35               if other.field.types does not fully contain field.types then
36                   flag ⟵ False
37               end
38           end
39       end
40       for method in class do
41           if method contains type not in relation then
42               return False
43           end
44           if other does not contain method then
45               flag ⟵ False
46           else
47               if other.method.types does not fully contain method.types then
48                   flag ⟵ False
49               end
50           end
51       end
52       if flag = True then
53           relation[other].add
54       end
55   end
56 end
```

type information. The set of unprocessed classes will then be inspected again in a recursive call (line 10).

The next two functions of the algorithm, *addClass* and *checkSuperSet* given in the next page, check subtyping for the current *class* being processed and update *relation*, which is a mapping from a type to its subtypes and originally only consists of entries from Map 3.1. In function *addClass(class)* we check if the type *class* is a subtype of an existing type in *relation* (lines 15–18), as well as the opposite, meaning if *class* is a supertype of an existing type in *relation* (line 19). In order to do so *checkSuperSet* checks all fields (lines 28–39) and all methods (lines 40–51) in *class* and compares them with an *existingClass* in *relation*. If a subtyping relation is established, then it is added to *relation* (line 53). Finally, upon returning from *checkSuperSet*, we also add *class* to its own relation (line 21) to satisfy reflexivity and to *Universe* (line 22), which is a supertype of all types.

It is worth noticing that the subtyping algorithm finds all nominal and structural subtypes of a given type. This is due to the fact that all pairs of types are inspected. Recall from Sect. 1 that nominal subtyping is name-based and given by the class hierarchy defined by the programmer, whereas structural subtyping is structure-based and given by the set-inclusion of fields and methods. In particular, structural subtyping is contra-variant with respect to this set-inclusion. Algorithm 2 finds all structural subtypes of a given class because it checks that its fields and methods are a superset of existing types in *relation*. For example, all classes are structural subtypes of type $empty = []$. On the other hand, it also finds all nominal subtypes because a class inherits all fields and methods of its superclass and as such its fields and methods are a superset of its superclass. This means that checking for structural subtyping is enough because nominal subtyping will be captured due to inheritance of fields and methods.

Finally, a note on complexity. The complexity of Algorithm 1 is $\mathcal{O}(n)$, and the complexity of Algorithm 2 is $\mathcal{O}(n^2)$, with n being the size of the input. The reason for a quadratic complexity of Algorithm 2 is due to the symmetric check of structural subtyping between a *class* and an *existingClass* in *relation*. Notice that if we were to only work with nominal subtyping, then we would only require traversing the class hierarchy once, which gives an $\mathcal{O}(n)$ complexity.

3.3   Type System for SFJ

The type system for the SFJ language, given in Sect. 2.2, is based on the type system by Dardha *et al.* [11,12] where the formal typing rules and soundness properties are detailed. As these formal developments are beyond the scope of this paper, we discuss typing for SFJ only informally.

A program $(\tilde{L}, e)$ is well typed if both $\tilde{L}$ and e are well typed. Class declaration L and method declaration M are well typed if all their components are well typed. Let us move onto expressions E. Field access $e.a$, method call $e.m(e)$ and object creation **new** $C(\widetilde{e})$ are typed in the same way as in Java: we inspect the type of the field and the type of the method and its arguments to determine the type of the field access and method call, respectively. The type of an object

creation is determined by the type of its class. Regarding constants, in order to respect the set-theoretic interpretation of types as sets of values, we type constants with the most restrictive type, i.e., the type representing the smallest set of values containing the value itself. For example, the type system would assign to the value 42 the type **byte**, which is the smallest in the sequence **byte**, **short**, **int** (see Map 3.1 for details).

Finally, the subtyping relation generated by Algorithm 2 is used in the type system for the SFJ language via a subsumption typing rule:

$$\frac{\Gamma \vdash e : \alpha_1 \qquad \alpha_1 \leq \alpha_2}{\Gamma \vdash e : \alpha_2}$$

We read this typing rule as follows: if an expression e is of type α_1 under a typing context Γ (details of a typing context are irrelevant here) and type α is a subtype $\leq$ of α_2, then expression e can be typed with α_2.

3.4 Polygons: Continued

Let us illustrate the semantic subtyping algorithm on our *Polygons* given in Example 1. Algorithm 2 generates the subtyping relation given in Map 3.2, together with the subtyping relation for basic types, omitted here and defined in Map 3.1. Notice that the mapping for *Universe* is extended with the new types for polygons.

$$
\begin{aligned}
Polygon &= \{Polygon, Triangle, Square, Rhombus\} \quad Triangle = \{Triangle\} \\
Square &= \{Square\} \quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad Rhombus = \{Rhombus\} \\
Universe &= \{Double, Float, Long, Int, Short, Byte \quad Diagonal = \{Diagonal\} \\
&\quad\quad Boolean, Void, Polygon, Square \\
&\quad\quad Square, Rhombus, Diagonal\}
\end{aligned}
$$

$$(3.2)$$

Recall the method *diagonal* in class *Diagonal*, with signature

$$\textbf{double } diagonal((Polygon \textbf{ and not } Triangle)\ shape)$$

The result of the set operation on its parameter type gives the following set of polygons:

$$Polygon \textbf{ and not } Triangle = \{Polygon, Square, Rhombus\}$$

In order to define the **not** *Triangle* type we need the *Universe* type so that we can define it as $Universe \setminus Triangle$. Then, the **and** connective is the intersection of sets of *Polygon* with **not** *Triangle*.

If we write in our SFJ program the following expression:

$$(\textbf{new } Diagonal()).diagonal(\textbf{new } Square())$$

the argument **new** *Square*() of the *diagonal* method is of type *Square*, by the type system in Sect. 3.3 and *Square* is contained in the set of the parameter type of the method, so this expression will successfully type-checks.

However, if we write the following SFJ expression:

$$(\textbf{new } Diagonal()).diagonal(\textbf{new } Triangle())$$

Type $Triangle$ is not contained in $\{Polygon, Square, Rhombus\}$, therefore this expression will not type-check and will return a type error at compile time.

This is further illustrated in the accompanying video of this paper [27].

3.5   Nominal vs. Structural Subtyping

In this section we will comment on pros and cons of nominal vs. structural subtyping.

Structural subtyping allows for more flexibility in defining this relation and the user does not need to explicitly definite it, as would do with nominal subtyping. However, for this flexibility we might need to pay in meaning. For example, consider the following two structurally equivalent classes, hence record types $coordinate = [x : int, y : int, z : int]$ and $colour = [x : int, y : int, z : int]$. While they can be used interchangeably in a type system using structural subtyping, their "meaning" is different and we might want to prohibit it, because intuitively speaking we do not want to use a $colour$ where a $coordinate$ is expected.

On the other hand, while nominal subtyping can avoid the above problem, it can introduce others and in particular, a developer can define an overridden method to perform the opposite logic to what the super class is expecting, as illustrated by the following classes in Java:

```
class A extends Object {          class B extends A {
        . . .                             . . .
    int n;                            int length(){ return −n; }
    int length(){ return n; }     }
}
```

Both approaches have their pros and cons, and they leave an expectation on the developer to use the logic behind subtyping correctly when writing code. The integration of both subtyping approaches in SFJ gives the developer the freedom to choose the most suitable subtyping relation to use for a given task.

3.6   Methods in SFJ

On Multimethods. Since FJ is a core language, some features of the full Java are removed, such as overloading methods. In our framework, by leveraging the expressivity of boolean connectives and semantic subtyping, we are able to restore overloading, among other features [12, §8.4]. We can thus model *multimethods*, [5], which according to the authors is "*very clean and easy to understand [...] it would be the best solution for a brand new language*". As an example, taken from Dardha *et al.* [11,12], consider the following class declarations:

```
class A extends Object {              class B extends A {
      int length (string s){ ... }          int length (int n){ ... }
}                                     }
```

Method *length* has type **string** → **int** in class *A*. However, because class *B* extends class *A*, *length* has type (**string** → **int**) **and** (**int** → **int**) in class *B*, which can be simplified to (**string or int**) → **int**.

Method Types. Let us illustrate the method types given in Sect. 2.1 via an alternative implementation of the class *Diagonal* at the end of Example 1.

```
class Diagonal {
          ...
      double diagonal((diagonal : void → double) shape)
             { return  shape.diagonal(); }
}
```

We define the type of the (outside) *diagonal* method as accepting any type and its subtypes implementing the (inside) *diagonal* method with type signature **void** to **double**.

In order to type check an argument passed to the (outside) *diagonal* method, at compile time we build a collection of types $\{type_1,\ type_2, \ldots\}$ which are class types where the (inside) *diagonal* method is defined. As such, we iterate over the list of *classes* in an SFJ program (as we did in Algorithm 2) to check for the required method. The resulting collection of types is the union of all classes where *diagonal* is defined together with their subtypes ($[\![type_1]\!] \cup [\![type_2]\!] \cup \ldots$), where each $[\![type_i]\!]$ denotes a mapping of $type_i$ to the set of its subtypes, similar to Map 3.2.

However, calculating this collection of types for each method of every class would be computationally inefficient and most importantly unnecessary as only few methods would in turn be used as method types. Therefore we only compute them on demand during type-checking when we come across such a type.

We can therefore use method types to statically include or exclude a portion of our class hierarchy. However, unlike with interfaces as in Example 1, the values that can be accepted by a method type do not have to be related to each other in any way in the class hierarchy. This indeed is useful if we are dealing with legacy code as we can still accept all classes where *diagonal* is defined, without having to go back and add interface implementations.

3.7   Code Generation

SFJ only includes the typechecking component of the language. In this section, we provide a sketch of the code generation algorithm, which is work in progress.

Given the similarity of SFJ to Java, our approach is to translate an SFJ program into Java bytecode and then run it on the Java Virtual Machine (JVM) [19]. This is a standard approach also used by other object-oriented languages, for example, Kotlin[1].

The main challenge in translating SFJ into bytecode is translating types using boolean connectives. For example, a field f_1 of type **int or bool**, will be translated as two Java fields, one of type **int** and one of type **bool**, and only one of the two types will be inhabited by a value. In order to achieve this, we first analyse our program and reduce the boolean types by keeping only the alternatives which actually get used in the program. For example, if the field of type **int or bool** only ever gets initialised with a boolean value, we can reduce it and make it a field of a single type **bool**. After this reduction phase, we then consider the remaining types which use boolean connectives and could not be reduced. On the example above, consider again field f_1 of type **int or bool**. This will be translated as two fields **int**_f_1 and **bool**_f_1 with the corresponding types. In order to initialise these fields, we use the constructor overloading capabilities of the JVM to generate an overloaded version for each alternative type. In each constructor, only the field that matches the type of the parameter is initialised with all other fields set to *null*. To access the field of an object, we generate code that checks each alternative of the field if it is non-null and includes the rest of the code generation for the expression for each branch. At run time, only one branch will be true, and this is the branch of generated code which will be executed.

For methods, we also use the overloading capabilities of the JVM to define an overloaded method for each type in the expanded method parameter, all with the same method body. Depending on which alternative the argument inhabits at run time, a different method will be dispatched to. Like methods with boolean types, we similarly implement methods with method types, as we discussed in Sect. 3.6, by defining an overloaded alternative for each type that implements the specified method. This concludes the code generation phase for all expressions in SFJ given in Sect. 2.2.

4   Related Work and Conclusion

Semantic subtyping goes back to more than two decades ago [1,10]. Hosoya and Pierce [14–16] define XDuce, an XML-oriented language designed specifically to transform XML documents in other XML documents satisfying certain properties. Frisch *et al.* [13] extend XDuce by introducing less XML specific types such as records, boolean connectives and arrow types, and implement it as ℂDuce. Their work is similar to ours in that our class-based semantic type system is a combination of the ℂDuce record types with arrow types. Castagna *et al.* define ℂπ [7], a variant of the asynchronous π-calculus, where channel types are augmented with boolean connectives; semantic subtyping for ML-like languages [8] and semantic subtyping in a gradual typing framework [6]. Ancona

---

[1] https://kotlinlang.org/

and Lagorio [3] define subtyping for infinite types by using union and object type constructors, where types are interpreted as sets of values of the language. Bonsangue *et al.* [4] study a coalgebraic approach to coinductive types and define a set-theoretic interpretation of coinductive types with union types. Pearce [26] defines semantic subtyping for rewriting rules in the Whiley Rewrite Language and for a flow-typing calculus [25].

Regarding implementations of semantic subtyping, to the best of our knowledge, there are only a few works in the literature. Muehlboeck and Tate [23] define a syntactic framework with boolean connectives which has been implemented in the Ceylon programming language [18]. Ancona and Corradi [2] define semantic subtyping for an imperative object-oriented language with mutable fields. In our framework we are considering only the functional fragment of Java, which is FJ, and as a result the semantic subtyping framework is simpler. The authors also propose a prototype implementation of their subtyping algorithm. Chaudhuri *et al.* [9] present the design and implementation of FLOW, which is a type checker for JavaScript. They use boolean connectives **and**, **not**, **or** for their predicates, however they do not define semantic subtyping for their language.

In this paper we presented the design and implementation of SFJ—Semantic Featherweight Java, an extension of Featherweight Java featuring boolean types, semantic subtyping, and integrating both nominal and structural subtyping. Due to the expressivity of semantic subtyping, in SFJ we are able to restore standard Java constructs and features for example, lists and overloading meathods, which were not present in FJ, thus making SFJ a more complete language. We presented Algorithm 1 on validity of type definitions and Algorithm 2 on semantic subtyping, which finds all nominal and structural subtypes for all types in an SFJ program. We also described typing of terms in SFJ, which follows that of Java and builds upon relevant work [11, 12]. As future work, we aim to finalise the code generation phase, which is sketched in Sect. 3.7.

References

1. Aiken, A., Wimmers, E.L.: Type inclusion constraints and type inference. In: Proceedings of the Conference on Functional Programming Languages and Computer Architecture, FPCA, pp. 31–41. ACM, New York (1993). https://doi.org/10.1145/165180.165188
2. Ancona, D., Corradi, A.: Semantic subtyping for imperative object-oriented languages. In: Visser, E., Smaragdakis, Y. (eds.) Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA, pp. 568–587. ACM (2016). https://doi.org/10.1145/2983990.2983992
3. Ancona, D., Lagorio, G.: Coinductive subtyping for abstract compilation of object-oriented languages into horn formulas. In: Proceedings of the Symposium on Games, Automata, Logic, and Formal Verification, GANDALF, EPTCS, vol. 25, pp. 214–230 (2010). https://doi.org/10.4204/EPTCS.25.20
4. Bonsangue, M., Rot, J., Ancona, D., de Boer, F., Rutten, J.: A coalgebraic foundation for coinductive union types. In: Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.) ICALP 2014. LNCS, vol. 8573, pp. 62–73. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-43951-7_6

5. Boyland, J., Castagna, G.: Parasitic methods: an implementation of multi-methods for Java. In: Proceedings of the Conference on Object-Oriented Programming Systems, Languages & Applications OOPSLA, pp. 66–76. ACM (1997). https://doi.org/10.1145/263698.263721

6. Castagna, G., Lanvin, V.: Gradual typing with union and intersection types. Proc. ACM Program. Lang. **1**(ICFP), 41:1–41:28 (2017). https://doi.org/10.1145/3110285

7. Castagna, G., Nicola, R.D., Varacca, D.: Semantic subtyping for the pi-calculus. Theor. Comput. Sci. **398**(1–3), 217–242 (2008). https://doi.org/10.1016/j.tcs.2008.01.049

8. Castagna, G., Petrucciani, T., Nguyen, K.: Set-theoretic types for polymorphic variants. In: Garrigue, J., Keller, G., Sumii, E. (eds.) Proceedings of the International Conference on Functional Programming, ICFP, pp. 378–391. ACM (2016). https://doi.org/10.1145/2951913.2951928

9. Chaudhuri, A., Vekris, P., Goldman, S., Roch, M., Levi, G.: Fast and precise type checking for Javascript. In: Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA, vol. 1, pp. 481–4830 (2017). https://doi.org/10.1145/3133872

10. Damm, F.M.: Subtyping with union types, intersection types and recursive types. In: Hagiya, M., Mitchell, J.C. (eds.) TACS 1994. LNCS, vol. 789, pp. 687–706. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-57887-0_121

11. Dardha, O., Gorla, D., Varacca, D.: Semantic subtyping for objects and classes. In: Beyer, D., Boreale, M. (eds.) FMOODS/FORTE -2013. LNCS, vol. 7892, pp. 66–82. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38592-6_6

12. Dardha, O., Gorla, D., Varacca, D.: Semantic subtyping for objects and classes. Comput. J. **60**(5), 636–656 (2017). https://doi.org/10.1093/comjnl/bxw080

13. Frisch, A., Castagna, G., Benzaken, V.: Semantic subtyping: dealing set-theoretically with function, union, intersection, and negation types. J. ACM **55**(4), 1–64 (2008). https://doi.org/10.1145/1391289.1391293

14. Hosoya, H., Pierce, B.C.: Regular expression pattern matching for XML. SIGPLAN Not. **36**(3), 67–80 (2001). https://doi.org/10.1145/373243.360209

15. Hosoya, H., Pierce, B.C.: XDuce: a statically typed XML processing language. ACM Trans. Internet Technol. **3**(2), 117–148 (2003). https://doi.org/10.1145/767193.767195

16. Hosoya, H., Vouillon, J., Pierce, B.C.: Regular expression types for XML. ACM Trans. Program. Lang. Syst. **27**(1), 46–90 (2005). https://doi.org/10.1145/1053468.1053470

17. Igarashi, A., Pierce, B.C., Wadler, P.: Featherweight Java: a minimal core calculus for Java and GJ. ACM Trans. Program. Lang. Syst. **23**(3), 396–450 (2001). https://doi.org/10.1145/503502.503505

18. King, G.: The Ceylon Language Specification, Version 1.3 (2016). https://ceylon-lang.org/documentation/1.3/spec/

19. Lindholm, T., Yellin, F., Bracha, G., Buckley, A.: Java Virtual Machine Specification, Java SE 7 Edition: Java Virt Mach Spec Java_3. Addison-Wesley (2013)

20. Liskov, B.H., Wing, J.M.: A behavioral notion of subtyping. ACM Trans. Program. Lang. Syst. **16**(6), 1811–1841 (1994). https://doi.org/10.1145/197320.197383

21. Malayeri, D., Aldrich, J.: Integrating nominal and structural subtyping. In: Vitek, J. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 260–284. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70592-5_12

22. Malayeri, D., Aldrich, J.: Is structural subtyping useful? An empirical study. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 95–111. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00590-9_8
23. Muehlboeck, F., Tate, R.: Empowering union and intersection types with integrated subtyping. Proc. Conf. Object-Oriented Program. Syst. Lang. Appl. OOPSLA **2**, 1–29 (2018). https://doi.org/10.1145/3276482
24. Parr, T.: The Definitive ANTLR 4 Reference. Pragmatic Bookshelf (2013)
25. Pearce, D.J.: Sound and complete flow typing with unions, intersections and negations. In: Giacobazzi, R., Berdine, J., Mastroeni, I. (eds.) VMCAI 2013. LNCS, vol. 7737, pp. 335–354. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35873-9_21
26. Pearce, D.J.: On declarative rewriting for sound and complete union, intersection and negation types. J. Comput. Lang. **50**, 84–101 (2019). https://doi.org/10.1016/j.jvlc.2018.10.004
27. Usov, A., Dardha, O.: SFJ: An implementation of Semantic Featherweight Java, On YouTube and on Dardha's website (2020). https://youtu.be/oTFIjm0A2O8, http://www.dcs.gla.ac.uk/~ornela/publications/SFJ.mp4