

Compile-time Reflection and Metaprogramming for Java

Weiyu Miao Jeremy Siek
University of Colorado at Boulder
{weiyu.miao, jeremy.siek}@colorado.edu

Abstract

Java reflection enables us to write reusable programs that are independent of certain classes. However, when runtime performance is a big concern, we propose to use compile-time reflection for writing metaprograms that generate non-reflective class/type specific code, which has lower runtime cost.

We proposed both compile-time reflection and metaprogramming for Java, and extended our previous work on pattern-based traits. Pattern-based traits integrate pattern-based reflection with flexible composition of traits. They are capable of pattern-matching over class members and generating expressive code. We add and formalize pattern-based reflection at the statement-level, which enables a meta program to generate statements. We add reified generics for pattern-based traits, which enables a pattern to iterate over any class when traits are instantiated. We implemented an ORM tool (called PtjORM) using pattern-based traits. PtjORM uses compile-time reflection and our benchmark tests show that it has competitive runtime performance compared to the mainstream Java ORM tools.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definition and Theory; D.3.3 [Programming languages]: Language Constructs and Features

General Terms Languages

Keywords compile-time reflection, metaprogramming, object-oriented programming, traits, object-relational mapping

1. Introduction

Java reflection [9] has the ability to inspect and analyze class members (i.e. metadata of a class) at run time. It enables programmers to write reusable programs that are independent of certain classes. For example, we are given a great number of classes and asked to implement the function `hasGetters`. The function checks if all the fields in a class have getters (i.e. field get methods) and prints out those fields without getters. This function is useful when class objects need to be serialized (e.g. converting objects into XML data) or object-relational mapped: many ORM tools require fields to have getters for access. It is impractical to manually write a `hasGetters` for each class, because each time we change the fields of a class,

we also likely need to change the `hasGetters` for that class, therefore such implementation is not adaptable to the change of a class.

With runtime reflection, we can write a `hasGetters` that is applicable to any class. Following shows the partial implementation of the `hasGetters`.

```
1 public boolean hasGetters(Class cls) {  
2     Field[] fds = cls.getFields();  
3     for(int i = 0; i < fds.length; ++i) {  
4         Field fd = fds[i];  
5         String mn = "get"+capitalize(fd.getName());  
6         Class fTy = fd.getType();  
7         try {  
8             Method _m = cls.getMethod(mn, new Class[0]);  
9             if (! _m.getReturnType().equals(fTy)) {  
10                ... /* print the message that field fd has no getter */  
11            }  
12        } catch (Exception e){  
13            ... /* print the message that field fd has no getter */  
14        }  
15    }  
16    ...  
17 }
```

The `hasGetters` iterates over all the fields of a class (line 2-3). For each field, it creates the name of the getter (line 5, function `capitalize` converts the first character to upper case) and searches in the class for a method that has the same name and with no parameters (line 8). If such method is found and its return type is equal to the field's type (line 9), then the method is a getter method.

Though Java runtime reflection enables programmers to write general programs, it has three drawbacks.

First, runtime reflection is not efficient. Suppose a program uses the `hasGetters` method to check a class that is statically known. Each time we run the program, the function needs to iterate over all the fields and the methods in the class. We quote the following words from the Java's documentation¹:

"Reflective operations have slower performance than their non-reflective counterparts, and should be avoided in sections of code which are called frequently in performance-sensitive applications."

That is why many Java ORM (Object Relational Mapping) tools (such as Hibernate, Ebean, etc) which use runtime reflection, raise the concern about the runtime cost of reflection.

Second, Java has no specific language that is dedicated to reflection. It uses objects to store metadata. For ease of use and for expressiveness, we would like to have a meta language, which enables users to inspect and manipulate metadata more easily. For example, the meta language of Garcia's calculus [10] for type-reflective metaprogramming has abundant type-level operators and uses types as terms.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PEPM '14, January 20–21, 2014, San Diego, CA, USA.
Copyright © 2014 ACM 978-1-4503-2619-3/14/01...\$15.00.
<http://dx.doi.org/10.1145/2543728.2543739>

¹see <http://docs.oracle.com/javase/tutorial/reflect/>

Third, Java reflection has little support for code generation. Forman’s book [9] mentioned code generation via reflection, but the object-level code is written in string, which has no type safety guarantees. We would like Java to have an object-level language that the compiler can type-check before code generation.

To overcome the drawbacks of runtime reflection, researchers have proposed compile-time reflection for its lower run-time cost and improved safety. Draheim [3] developed reflective program generators with compile-time loops that can iterate over the members of one class to create the members for another class. Fährndrich [6] introduced *pattern-based reflection*, which provides high-level patterns and templates. For better code reuse, researchers integrated pattern-based reflection with different composition mechanisms. Huang [11, 12] introduced MorphJ, which increased the expressiveness and safety guarantees of pattern-based reflection. In MorphJ, reflective patterns reside in mixin-like structures [8]: a reflective pattern can match the members of a parametrized superclass and generate members for its subclass. Mixins support linear composition, which imposes some restriction on code reuse. Borrowing the concept of the reflective patterns from MorphJ, we proposed pattern-based traits [20]. They are the combination of pattern-based reflection with traits [23], which offer more expressive forms of composition.

In this paper, we present pattern-based reflection at the statement level. Pattern-based reflection is not new, but to our knowledge, its use at the statement level has never been fully discussed and formalized. We find out that statement-level pattern-based reflection supports the writing of reusable metaprograms, thus it is necessary to formalize the language and prove that it always generates well-typed statements.

Following shows the `hasGetters` implemented in statement-level pattern-based reflection.

```

1 trait getterChecker<class X>
2 provides {
3   public boolean hasGetters() {
4     boolean has = false; boolean ret = true;
5     pattern <F, name f> F $f in X {
6       pattern <> public F get#$(f) in X {
7         has = true;
8       }
9     }
10    if (! has) {
11      ret = false;
12      println("Field " + $f::getName() + " has no getter!");
13    } else { has = false; }
14  }
15  return ret;
16 }

```

Trait `getterChecker` is parameterized over the class that is to be reflected over. Inside the body of function `hasGetters`, the outer pattern matches each field in `X` (line 5) and passes the field’s name and the type to the inner pattern, which searches the field’s getter method in `X` (line 6). The symbol `#` is the name concatenation operator. In line 11, expression `$f::getName()` is a meta function call that, at compile time, returns the string of field `f`’s name. The `hasGetters` in the trait can be specialized for any class and also adaptable to the change of a class.

For example, suppose we are given the following class `Person`:

```

1 class Person {
2   private String full_name;
3   private int age;
4   public int getAge() { return this.age; }
5 }

```

In class `Person`, field `full_name` has no getter method. When trait `getterChecker` is applied to class `Person`, the trait generates the following `hasGetters` function:

```

1 public boolean hasGetters() {
2   boolean has = false; boolean ret = true;
3   if (! has) {
4     ret = false;
5     println("Field " + full_name + " has no getter!");
6   } else { has = false; }
7   has = true;
8   if (! has) {
9     ret = false;
10    println("Field " + age + " has no getter!");
11  } else { has = false; }
12  return ret;
13 }

```

The generated `hasGetters` function is specified for class `Person` and it has no runtime reflection in the body. Please note that in line 7, the value of variable `has` is changed into `true`. So, the second `println` is not executed.

Besides statement-level pattern-based reflection, we also introduce reified generics, which enables a pattern to iterate over any specific class when type parameters are instantiated. Previously, the set of classes that a pattern-based trait can reflect over includes the class that uses the trait and the class’s superclasses. So, we increased expressiveness of pattern-based reflection.

This paper makes the following contributions:

- We introduce pattern-based reflection at the statement level, which supports statement-level code generation and reuse (Section 2.2).
- We add generics to traits and extend ranges with class types, which enable a pattern to reflect over a parametrized class. (Section 2.2).
- We present reified generics, which enable generated code to access specialized types (Section 2.3). We introduce meta-level functions, which enable users to inspect metadata at the compile time (Section 2.3).
- We formalize our language, both the object language (Section 3) and the meta language, and present their type systems, which are implementable (Section 4).
- We implemented an ORM tool with compile time reflection and benchmark tested the tool to seek the improvement of runtime performance by using compile-time reflection (Section 5).

We implemented our proposed programming language features based on Polyglot². The source code and the ORM tool is available at <http://code.google.com/p/pattern-based-traits/>.

2. Language Features

In this section, we introduce the language features. We start with an overview of traits and pattern-based reflection. Even though they were introduced in our previous paper [20], it is worthwhile to review their syntax not only for readers’ convenience but also for the modification we have made to their syntax.

2.1 Traits

A trait [5, 23] is a unit for code reuse. It provides a set of method implementations, which may depend on class members (i.e. fields or methods) that are not given by the trait. A trait can import/use methods from other traits, which means a trait can be composed of methods imported/used from other traits. When methods are

² <http://www.cs.cornell.edu/projects/polyglot/>

merged during trait composition, we can manipulate those methods (such as method renaming, method exclusion, method aliasing, etc) to avoid name conflicts.

We use the Java trait syntax proposed by Reppy and Turon [22]. A basic trait without parameters has the following syntax:

```
trait-declaration ::=
  trait trait-name
  requires { requirements }
  provides { members }
```

The optional **requires** clause contains the signatures of dependent class members. The **provides** clause contains members which include provided methods and the **use** clauses for importing methods from other traits. In detail, the **use** clause contains a list of expressions of trait names and method manipulations.

2.2 Pattern-Based Reflection

A reflective pattern at the statement level resides in a method body. It performs both reflection and code generation. The header part can iterate over a sequence of class members (i.e. fields, methods, and constructors), and pattern-match against each of them. The body contains template code that is parameterized over names and/or types. A reflective pattern generates different code instantiations from different names and types obtained using pattern-matching.

In the following, we give the syntax for reflective patterns at the statement level.

```
pattern-declaration ::=
  pattern
  <parameters> [modifier-pattern] member-pattern
  in range
  { statements }
```

A pattern declaration may have a name, but it is not necessary at the statement level. Pattern parameters include name parameters and (constrained) type parameters.

A member pattern is in the form of a field signature, a method signature, or a constructor signature. It may be prefixed with an access-level modifier pattern for pattern-matching a group of class members with certain access level(s). An access-level modifier pattern can be **public**, **private**, **none** for package-private, **nonprivate** for the access levels that are not private, etc.

A range represents a sequence of class members that a pattern can iterate over. It has the following syntax:

```
range ::= identifier | range{identifiers} | range\{identifiers}
```

which can refer to a class type, a class type variable (including the reserved type variable: **thisType**), a member-selection operation (i.e. selecting specified members from a range), or a member-exclusion operation (i.e. removing specified members from a range). Our previous paper [20] also proposed the sum of two ranges, which is not allowed here.

The body of a statement-level pattern is a sequence of statements. We do not allow a return statement to appear inside a pattern's body, otherwise a pattern can generate statements containing unreachable code.

As we mentioned in the introduction, statement-level reflective patterns enable us to write general methods that are applicable to different classes and also adaptable to class change. For example, when we design an extensible programming language framework, we have to use the visitor pattern for separating the syntax from its behaviors. Following is the syntax of a tiny language for integers and immutable arrays:

```
E ::= Integer | E + E | {E} | E[E]
```

In the syntax, expression $\{\overline{E}\}$ represents immutable array definition, which is like a sequence of expressions; and $E[E]$ means array access. The below shows the Java code of the language's abstract syntax tree nodes:

```
1 interface Node { void visit(NodeVisitor v); }
2
3 class IntLit extends Node {
4   private int; public void visit(NodeVisitor v) { }
5 }
6 class Addition extends Node {
7   private Node left; private Node right;
8   public void visit(NodeVisitor v) {
9     this.left = v.visit(this.left); this.right = v.visit(this.right);
10  }
11 }
12 class Array extends Node {
13   private List fds;
14   public void visit(NodeVisitor v) {
15     List tmp = new ArrayList();
16     for(Iterator i = this.fds.iterator(); i.hasNext(); ) {
17       Object fd = i.next();
18       if (fd instanceof Node)
19         tmp.add(v.visit((Node) fd));
20     }
21     this.fds = tmp;
22   }
23 }
24 class ArrAccess extends Node {
25   private Node arr; private Node idx;
26   public void visit(NodeVisitor v) {
27     this.arr = v.visit(this.arr); this.idx = v.visit(this.idx);
28   }
29 }
```

In the above code, we assume that abstract class `NodeVisitor` gives the method `Node visit(Node n) { n.visit(this); return n; }`. For each syntax tree node, we implement the `visit` method that visits its sub-nodes. With this visitor pattern, we can easily write a type-checker, an evaluator, and a syntax tree printer for this language (detailed implementations are omitted).

If the above language is fully extended, it will be tedious to manually write a `visit` method for each syntax node. Therefore, we write the following metaprogram to generate the `visit` methods.

```
1 trait VisitGen<>
2 provides {
3   public void visit(NodeVisitor v) {
4     pattern<F extends Node, name f> F $f in thisType {
5       this.$f = (F) v.visit(this.$f);
6     }
7     pattern<F extends Collection, name f> F $f in thisType
8     {
9       List tmp = new ArrayList();
10      for(Iterator i = this.$f.iterator(); i.hasNext(); ) {
11        Object e = i.next();
12        if (e instanceof Node)
13          tmp.add(v.visit((Node) e));
14      }
15      this.$f.clear(); this.$f.addAll(tmp);
16    }
17    // insert a reflective pattern for the fields of array types
18    ...
19  }
20 }
```

In trait `VisitGen` and inside `visit`, the first pattern (line 4-6) matches the fields which are syntax tree nodes and generates the statements for them. The second pattern (line 7-16) matches and generates statements for the fields of type `Collection`. In the body of the second pattern, the code traverses the collection and applies the

visit method to an element if it is a syntax tree node. We restrict the range to the type variable **thisType**, which will be automatically substituted for the name of the class that uses VisitGen. In the patterns, we have the code (in lines 5, 10, and 15) that may access the private fields via the **this** variable, and the **thisType** range guarantees the code is well-typed. For instance, when class Addition uses VisitGen, the **thisType** variable is substituted for Addition; a visit method is generated for Addition; and inside the visit method, it is safe to access Addition's private fields. In section 2.4, we give more detailed discussion about type safety.

2.3 Reified Generics

To support code compatibility, Java implements generics using type erasure³ so that the specialized types are not available at runtime. For instance, the following expressions are not accepted in Java (we assume X is a well-defined type variable): `X.class`, `obj instanceof X`, `new X()`, `new X[10]`, etc.

Using metaprogramming, we generate specialized code for the instantiation of generics, like C++ templates [1], so that specialized types are preserved in generated code.

For example, in the following, we give a generic equals function at the meta level, which can be instantiated into a specialized equals function for any class.

```

1  trait EqualGen<>
2  provides {
3      public boolean equals(Object obj) {
4          if (obj instanceof thisType)
5              return equals_k(this, (thisType) obj);
6          return false;
7      }
8      private boolean equals_k(thisType obj1, thisType obj2) {
9          boolean is_equal = true;
10         pattern <primitive T, name f> T $f in thisType {
11             if (obj1.$f != obj2.$f)
12                 is_equal = false;
13         }
14         pattern <T extends Object, name f> T $f in thisType {
15             // code for comparing two objects
16             ...
17         }
18         return is_equal;
19     }
20 }

```

In line 4, the `instanceof` operator is applied to the **thisType** variable. Inside function `equals_k`, the first pattern matches and compares the fields of primitive types, while the second pattern matches and compares the fields of reference types. When the trait `EqualGen` is used by some concrete class C, the variable **thisType** is specialized into type C.

Following is the example for another use of reified generics: generating instance creation functions for the support of the factory method pattern⁴.

```

1  trait InstanceGen<class X, class S>
2  provides {
3      public S createInstance() throws InstantiationException {
4          X obj = null;
5          pattern <> public constructor() in X {
6              obj = new X();
7          }
8          if (obj instanceof S) {
9              println("An instance of "+X::getName()+" is created.");
10             return (S) obj;
11         }

```

³ see <http://docs.oracle.com/javase/tutorial/java/generics/erasure.html>

⁴ see http://en.wikipedia.org/wiki/Factory_method_pattern

```

12     throw new InstantiationException("...");
13 }
14 }

```

The trait receives two arguments: the type of a class whose instance can be created and the interface that the class implements. In line 5-6, the pattern matches a public nullary (or default) constructor for class X and creates an instance via the `new` operator if the constructor exists. In line 8, we check if the created object is an instance of type variable S. In line 9, we have the meta function call `X::getName()` that returns the name of X at compile time. Suppose we have class `Student`, which has the nullary constructor and implements interface `Person`. When trait `InstanceGen` is applied to `Student` and `Person`, it generates the following result:

```

1  public Person createInstance() throws InstantiationException {
2      Student obj = null;
3      obj = new Student();
4      if (obj instanceof Person) {
5          println("An instance of "+Student+" is created.");
6          return (Person) obj;
7      }
8      throw new InstantiationException("...");
9  }

```

Besides the meta function `getName`, we have other predefined meta functions that enable users to inspect metadata at compile time. For instance, expression `X::getSimpleName()` returns the simple name of type X; `X::equals(Y)` checks if X is equal to Y; `X::isSubType(Y)` checks if X is a subtype of Y; `X::superClass()` returns the direct superclass of X; `X::isPrimitive()` checks if X is a primitive type; and so on.

2.4 Member Accessibility

A pattern-based trait uses the pattern structure to reflect over the members of some class A and uses the composition power of traits to extend some class B. We have to discuss the relation between A and B because it has influence over the accessibility of a member. In this section, we discuss the conditions when the members of a class can be accessed via the **this** variable.

Consider the following metaprogram, which generates a function to backup the fields in a superclass.

```

1  trait FieldBackupGen<this-class S>
2  provides {
3      public void backup() {
4          pattern <T, name f> nonprivate T $f in S {
5              pattern <> T backup#$f in thisType {
6                  this.backup#$f = this.$f;
7              }
8          }
9      }
10 }

```

The type variable S in the metaprogram cannot be instantiated with an arbitrary class. For instance, trait `FieldBackupGen` used as follows generates ill-typed code.

```

1  class Account { protected int balance; }
2  class AccBackup {
3      private int backupBalance;
4      use FieldBackupGen<Account>;
5  }

```

For the above, the body of the instantiated function `backup` is

```
this.backupBalance = this.balance
```

but field `balance` cannot be accessed via the **this** variable inside class `AccBackup`. The correct implementation is to let class `AccBackup`

inherit from `Account`, so we can access `Account`'s non-private members via the `this` variable.

In the definition of trait `FieldBackupGen`, we use keyword `this-class` to restrict the set of classes that a trait can be applied to. It means that an accepted class must be a super type of the class that uses the trait. In detail, if the trait is used by class `A`, then an accepted class must be a super class of `A` or class `A` itself. Suppose some type variable `X` is prefixed with `this-class`; the subtype relationship between `thisType` and type variable `X` is that `thisType <: X`.

In the above code, trait `FieldBackupGen`'s type parameter `S` is restricted by the `this-class`. Trait `FieldBackupGen` is used by class `AccBackup` and is applied to class `Account`. However, class `Account` is not a super type of class `AccBackup`, thus the trait application is rejected by the type system.

For an unrestricted type variable `X`, we allow users to only access its public members via an instance of `X`.

3. Calculus for the Object Language

Start from this section, we formalize our language. First, we present the calculus for our object language, which is the modest extension of FJ (short for Featherweight Java) [15] with mutable variables and three basic kinds of statements: variable declaration, assignment, and return statement.

x, y	term variables
m	method name
f	field name
C	type (i.e. class name)
class decl.	$L ::= \text{class } C \text{ extends } C \{ \bar{C} \bar{f}; K \bar{M} \}$
constructor	$K ::= C(\bar{C} \bar{f}) \{ \text{super}(\bar{f}); \text{this}.\bar{f}=\bar{f}; \}$
method	$M ::= C m(\bar{C} \bar{x}) \{ s; \}$
statements	$s ::= \text{return } e \mid C x = e; s \mid x = e; s$
expressions	$e ::= x \mid e.f \mid e.m(\bar{e}) \mid \text{new } C(\bar{e}) \mid (C)e \mid \mathbb{E}[s]$

Figure 1. FJ syntax extended with statements.

Figure 1 shows the abstract syntax of the object language. In the figure, the syntax of statements is the addition to FJ. A statement can be a return statement or a sequence of statements with a return statement at the end. The variable declaration $C x = e$ declares local variable x of type C and initializes x with the value of expression e . The assignment $x = e$ assigns the value of e to x . In Java, an assignment is also an expression that gives the value of e , but in our object language, we treat an assignment only as a statement. The rest of the syntax in the figure is the same as FJ except for methods and the expression $\mathbb{E}[s]$. In the object language, the body of a method is a sequence. The expression $\mathbb{E}[s]$ does not appear in a concrete program. It wraps a statement and splices it into an expression.

Statement typing:	$\Gamma \vdash s : C$
	$\frac{\Gamma \vdash x : C_1 \quad \Gamma \vdash e : C_0 \quad C_0 <: C_1 \quad \Gamma \vdash s : C}{\Gamma \vdash x = e; s : C}$
	$\frac{\Gamma \vdash e : C_0 \quad C_0 <: C \quad \Gamma, x : C \vdash s : C_1}{\Gamma \vdash C x = e; s : C_1} \quad \frac{}{\Gamma \vdash \text{return } e : C}$

Figure 2. Typing rules for statements.

Figure 2 shows the rules for typing statements. We define Γ as a typing environment, which is a finite mapping from variables to types. An environment extension has the form $\Gamma, x : C$, which

means extending Γ with variable typing $x : C$ only if x does not appear in Γ . The judgments $C_0 <: C_1$ means class C_0 is a subtype of class C_1 ; and $\Gamma \vdash e : C$ means expression e has type C under the typing environment Γ . We use the subtyping rules and the typing rules for expressions from FJ's type system (see the figures Fig.1 and Fig.2 in [15]). For the additional expression $\mathbb{E}[s]$, its type should be equal to the type of statement s . See the following typing rule:

$$\frac{\Gamma \vdash s : C}{\Gamma \vdash \mathbb{E}[s] : C}$$

We also use the typing rules from FJ to type-check methods and class declarations. For FJ's method typing rule, we need to override the rule of type-checking a method body: we use the statement typing rule to type-check the method body. Due to the space limitation, we do not present those rules in this paper.

Statement reduction: $s \longrightarrow s'$

$$\frac{C x = v; s \longrightarrow [v/x]s}{x = v; s \longrightarrow [v/x]s} \quad \frac{e \longrightarrow e'}{C x = e; s \longrightarrow C x = e'; s} \quad \frac{e \longrightarrow e'}{x = e; s \longrightarrow x = e'; s} \quad \frac{e \longrightarrow e'}{\text{return } e \longrightarrow \text{return } e'}$$

Figure 3. Statement reduction rules.

$$\begin{aligned} [e'/x](\text{return } e) &= \text{return } [e'/x]e \\ [e'/x](C y = e; S) &= (C y = [e'/x]e; [e'/x]S) \\ [e'/x](x = e; S) &= (x = [e'/x]e; S) \\ [e'/x](y = e; S) &= (y = [e'/x]e; [e'/x]S) \text{ where } x \neq y \end{aligned}$$

Figure 4. Substitution in statements.

Figure 3 shows the reduction rules for statements. We define v as the value of an expression, which is $\text{new } C(\bar{v})$. Similar to $[e'/x]e$, the operation $[e'/x]s$ means the substitution of x in statement s for e . Figure 4 shows its definition. The substitution $[e'/x](C x = e; S)$ does not need to be defined. The substitution implies that x is declared twice in the same scope, which is precluded by the type system. In the figure, we use the reduction rules for expressions (the form $e \longrightarrow e'$) from FJ (see Figure Fig.3 in [15]). But we need to revise the computation rule for method invocation and provide the reduction rules for $\mathbb{E}[s]$. Those rules are shown in Figure 5.

Expression reduction: $e \longrightarrow e'$

$$\frac{\text{mbody}(m, C) = \bar{x}.s}{(\text{new } C(\bar{e})) . m(\bar{e}_0) \longrightarrow \mathbb{E}[[\bar{e}_0/\bar{x}, \text{new } C(\bar{e})/\text{this}]s]}$$

$$\frac{s \longrightarrow s'}{\mathbb{E}[s] \longrightarrow \mathbb{E}[s']} \quad \frac{}{\mathbb{E}[\text{return } v] \longrightarrow v}$$

... (The rest of the reduction rules are the same as those for FJ.)

Figure 5. Selected reduction rules for expressions. (Like FJ's *mbody* function, the *mbody* function in the figure returns the parameter(s) and the body of C 's method m .)

The calculus of our object language is type safe. It inherits the properties of FJ: type preservation and progress for expressions. Besides, the properties of type preservation and progress also apply to statements. See the following properties for statements.

Lemma 1 (Substitution Preserves Typing). *If $\Gamma, x : C_0 \vdash s : C$ and $\Gamma \vdash e : C_1$ where $C_1 <: C_0$, then $\Gamma, x : C_0 \vdash [e/x]s : C'$ for some C' such that $C' <: C$.*

Please note that, in Lemma 1, when type-checking $[e/x]s$, we do not remove the typing of x from typing context Γ , because the substitution of x in s does not always substitute all the x s in s (see the third rule in Figure 4). So, variable x might be still in $[e/x]s$.

Theorem 1 (Type Preservation). *If $\Gamma \vdash s : C$ and $s \longrightarrow s'$, then $\Gamma \vdash s' : C'$ for some C' such that $C' <: C$.*

Theorem 2 (Progress). *If $\Gamma \vdash s : C$, then s is either a statement value (i.e. `return v`) or there is some s' with $s \longrightarrow s'$.*

4. Calculus for the Meta Language

Our calculus for the meta language captures the new language features and presents the kernel part of the language, thus we omit some features that were fully discussed in the previous papers, such as the manipulation of trait members (e.g. name exclusion, member aliasing, etc). We also omit type bounds, modifier patterns, and meta-level functions, which are considered as advanced features.

4.1 Kernel Syntax of the Meta Language

Figure 6 shows the syntax of the core meta language. In the figure, a type τ is a class name associated with the signatures of the class's members. It gives both the nominal and the structural representation of a class type, and it is the value of a range. A range variable X can be used as a type variable. When the range variable X is substituted for some τ in type T , X should be substituted for the nominal representation of τ , that is the substitution $[C \diamond \{\bar{F}; Q; \bar{H}\} / X]T$ is reduced into $[C/X]T$. But it performs a normal substitution when the range variable X is substituted for some τ in range R .

In the expressions, we can use the `new` operator to create an instance of a parameterized class, for instance, we can write the expression `new X(\bar{e})`.

A trait TR is parameterized over ranges. We attach the symbol $+$ or the symbol $-$ to each range parameter. The parameter X^+ means the members (excluding constructors) of X can be accessed via the `this` variable. The parameter X^- means the members of X can be only accessed via an instance of X .

Ranges include member selection ($R[\{\bar{I}\}]$) and member exclusion ($R \setminus \{\bar{I}\}$). A range is evaluated into a range value τ , which does not appear in a concrete program. For the range of some class C , we compute its value by collecting the signatures of the members in C , including C 's inherited members.

4.2 Type System

In this section, we discuss the type system of the meta language calculus. First, we give some preliminary definitions (see Figure 7) which are used by the type system.

In the figure, member types present the types for different kinds of class members: T for fields, $\bar{T} \rightarrow T$ for methods, and $\bar{T} \rightarrow \cdot$ for constructors. A member name ℓ extends 1 with the reserved name C as the unified name for all constructors. A member typing context Δ is the mapping from names to member types. The notation $\Delta[\ell \mapsto \phi]$ means to extend Δ with the mapping from ℓ to ϕ . The notation $\Delta_1 \oplus \Delta_2$ means to merge two member typing contexts if the domains of Δ_1 and Δ_2 are disjoint; otherwise it generates a type error. A structural typing context Θ is the mapping from structure names to member typing contexts. The notation $\Theta \uplus [\kappa \mapsto \Delta]$ means to extend Θ with the new mapping from κ to Δ if $\kappa \notin \Theta$, or means $(\Theta \setminus \kappa) \uplus [\kappa \mapsto \Delta \oplus \Theta(\kappa)]$ if $\kappa \in \Theta$. A variable binding context Γ can bind a term variable to a nominal type ($x : T$), a range variable X^o , the `thisType` variable, or a name variable η .

x, y	term variables
X, Y	type/range variables
η	member name variable
m	method name
f	field name
C	class name
\mathbb{C}	reserved name constant
\mathcal{T}	trait name
L	object class declaration (see Figure 1)

(1) types, names, and member signatures:

type	$\tau ::= C \diamond \{\bar{F}; Q; \bar{H}\}$
nonvar. types	$N ::= C \mid \text{thisType}$
nominal types	$U, T ::= X \mid N$
member names	$l ::= \eta \mid f \mid m$
field sig.	$F ::= T \ 1$
constructor sig.	$Q ::= C(\bar{T})$
method sig.	$H ::= T \ 1(\bar{T})$

(2) meta classes:

meta class	$MC ::= L \ \text{uses} \ \bar{E}$
meta meth.	$M ::= T \ m(\bar{T} \ \bar{x})\{s;\}$
single stmts.	$ss ::= T \ x = e \mid x = e \mid ps$
statements	$s ::= \text{return } e \mid ss; s$
expressions	$e ::= x \mid e.1 \mid e.1(\bar{e}) \mid \text{new } T(\bar{e}) \mid (T)e$

(3) traits:

trait	$TR ::= \text{trait } \mathcal{T} \langle \bar{X}^o \rangle \ \text{req} \ \{ \bar{F}; \bar{H} \} \ \text{prov} \ D$
access ctrl.	$o ::= + \mid -$
trait body	$D ::= \{ \bar{M}; \ \text{use} \ \bar{E} \}$
trait app.	$E ::= \mathcal{T} \langle \bar{R} \rangle$

(4) pattern statements:

pattern stmt.	$ps ::= \text{pattern} \langle \bar{X}, \bar{\eta} \rangle \ P \ \text{in} \ R \ \{ ns; \}$
patterns	$P ::= F \mid Q \mid H$
ranges	$R ::= X \mid N \mid R[\{\bar{I}\}] \mid R \setminus \{\bar{I}\} \mid \tau$
non-return stmts.	$ns ::= ss \mid ss; ns$

Figure 6. Syntax of the calculus for the meta language.

member types	$\phi ::= T \mid \bar{T} \rightarrow T \mid \bar{T} \rightarrow \cdot$
member names	$\ell ::= 1 \mid C$
structure names	$\kappa ::= T \mid \mathcal{T}$
member typings	$\Delta ::= \cdot \mid \Delta[\ell \mapsto \phi] \mid \Delta \oplus \Delta$
structural typings	$\Theta ::= \cdot \mid \Theta \uplus [\kappa \mapsto \Delta]$
variable bindings	$\Gamma ::= \emptyset \mid \Gamma, x : T \mid \Gamma, X^o \mid \Gamma, \text{thisType} \mid \Gamma, \eta$

Figure 7. Preliminary definitions for the type system.

We define the function \mathbb{N} , which computes the nominal representation of a range: $\mathbb{N}(X) = X$, $\mathbb{N}(N) = N$, $\mathbb{N}(R|\{\bar{I}\}) = \mathbb{N}(R)$, $\mathbb{N}(R|\{\bar{I}\}) = \mathbb{N}(R)$, and $\mathbb{N}(C \diamond \{\dots\}) = C$. We define the function δ , which literally translates a member pattern into a single member typing context: $\cdot[\ell \mapsto \phi]$.

Before type-checking a program, the compiler computes the structural type (this is member typings) for each class and each trait in the program, thus the following two structural typing contexts are available if there are no name conflicts. First is the context Θ^{OL} that maps the names of all the classes into their member typings. The members of a class include declared members, members imported from used traits, and members inherited from the super-classes. Second is the context Θ^{TR} that maps the names of all the traits into their member typings. The members of a trait include declared members, members imported from used traits, and members in trait requirement.

Pattern statement typing: $\Gamma; \Theta \vdash ps \text{ ok}$

$$\frac{\mathbb{N}(R) = C \quad \Gamma \vdash R \text{ ok} \quad \Gamma' = \Gamma, \bar{X}^-, \bar{\eta} \quad \Gamma' \vdash P \text{ ok} \quad \Theta' = \Theta \uplus [C \mapsto \delta(P)] \quad \Gamma'; \Theta' \vdash ns \text{ ok}}{\Gamma; \Theta \vdash \text{pattern}\langle \bar{X}, \bar{\eta} \rangle P \text{ in } R \{ ns; \} \text{ ok}}$$

$$\frac{\mathbb{N}(R) = X \quad X^+ \in \Gamma \quad \Gamma' = \Gamma, \bar{X}^-, \bar{\eta} \quad \Gamma' \vdash P \text{ ok} \quad \delta(P) = \Delta \quad (P = F \text{ or } P = H) \text{ implies } \Theta' = \Theta \uplus [X \mapsto \Delta] \uplus [\text{thisType} \mapsto \Delta] \quad P = Q \text{ implies } \Theta' = \Theta \uplus [X \mapsto \Delta] \quad \Gamma'; \Theta' \vdash ns \text{ ok} \quad \Gamma \vdash R \text{ ok}}{\Gamma; \Theta \vdash \text{pattern}\langle \bar{X}, \bar{\eta} \rangle P \text{ in } R \{ ns; \} \text{ ok}}$$

$$\frac{(\mathbb{N}(R) = X \text{ and } X^- \in \Gamma) \text{ or } \mathbb{N}(R) = \text{thisType} \quad \Gamma' = \Gamma, \bar{X}^-, \bar{\eta} \quad \Gamma' \vdash P \text{ ok} \quad \Gamma \vdash R \text{ ok} \quad \Theta' = \Theta \uplus [\mathbb{N}(R) \mapsto \delta(P)] \quad \Gamma'; \Theta' \vdash ns \text{ ok}}{\Gamma; \Theta \vdash \text{pattern}\langle \bar{X}, \bar{\eta} \rangle P \text{ in } R \{ ns; \} \text{ ok}}$$

Statement typing: $\Gamma; \Theta \vdash s : T$ and $\Gamma; \Theta \vdash ns \text{ ok}$

$$\frac{\Gamma; \Theta \vdash ps \text{ ok} \quad \Gamma; \Theta \vdash s : T \quad \Gamma; \Theta \vdash e : T_0 \quad T_0 <: T \quad \Gamma, x : T; \Theta \vdash s : U}{\Gamma; \Theta \vdash ps; s : T} \quad \frac{\Gamma; \Theta \vdash e : T_0 \quad T_0 <: T \quad \Gamma; \Theta \vdash s : U}{\Gamma; \Theta \vdash T x = e; s : U}$$

$$\frac{\Gamma; \Theta \vdash x : T \quad \Gamma; \Theta \vdash e : T_0 \quad T_0 <: T \quad \Gamma; \Theta \vdash s : U}{\Gamma; \Theta \vdash T x = e; s : U}$$

$$\frac{\Gamma; \Theta \vdash e : T \quad \text{for some } e \quad \Gamma; \Theta \vdash (ns; \text{return } e) : T}{\Gamma; \Theta \vdash \text{return } e : T} \quad \Gamma; \Theta \vdash ns \text{ ok}$$

Expression typing: $\Gamma; \Theta \vdash e : T$

$$\frac{x : T \in \Gamma \quad \Gamma; \Theta \vdash e : T \quad \Gamma \vdash 1 \text{ ok} \quad \Delta = \Theta(T)}{\Gamma; \Theta \vdash x : T} \quad \frac{\Gamma; \Theta \vdash e : T \quad \Gamma \vdash 1 \text{ ok} \quad \Delta = \Theta(T)}{\Gamma; \Theta \vdash e.1 : \Delta(1)}$$

$$\frac{\Gamma; \Theta \vdash e : U \quad \Gamma \vdash 1 \text{ ok} \quad \Delta = \Theta(U) \quad \Delta(1) = \bar{T} \rightarrow T \quad \Gamma; \Theta \vdash \bar{e} : \bar{U} \quad \bar{U} <: \bar{T}}{\Gamma; \Theta \vdash e.1(\bar{e}) : T}$$

$$\frac{\Gamma \vdash T \text{ ok} \quad \Delta = \Theta(T) \quad \Delta(C) = \bar{T} \rightarrow \cdot \quad \Gamma; \Theta \vdash \bar{e} : \bar{U} \quad \bar{U} <: \bar{T}}{\Gamma; \Theta \vdash \text{new } T(\bar{e}) : T}$$

$$\frac{\Gamma \vdash T \text{ ok} \quad \Gamma; \Theta \vdash e : U \quad \Gamma \vdash T \text{ ok} \quad \Gamma; \Theta \vdash e : U \quad U <: T \text{ or } T <: U \quad U \not<: T \text{ and } T \not<: U \quad \text{stupid warning}}{\Gamma; \Theta \vdash (T)e : T} \quad \Gamma; \Theta \vdash (T)e : T$$

Figure 8. Statement and expression typing.

Due to the space limitation, we omit the rules for type-checking ranges ($\Gamma \vdash R \text{ ok}$), member patterns ($\Gamma \vdash P \text{ ok}$), names ($\Gamma \vdash 1 \text{ ok}$), and nominal types ($\Gamma \vdash T \text{ ok}$). Those rules are not difficult to write. We also omit the subtyping rules ($T <: T$). Because we do not have

bounded type parameters in the kernel calculus, thus the subtyping rules are similar to those for FJ [15].

In Figure 8, the pattern statement typing has three typing rules. The first rule is for the case when the range R 's nominal part is some class name C . In this case, we extend the member typing context of C with the member typing generated from the pattern P : $\Theta \uplus [C \mapsto \delta(P)]$. The first rule type-checks the following code:

```

1 C obj;
2 pattern <name f> public int $f in C {
3   println("$f+" has value "+ obj.$f);
4 }

```

By type-checking the header of the pattern (line 2), C 's member typing context is extended with $f \mapsto \text{int}$, which means C has field f of type int within the scope of the pattern's body. So in line 3, the expression $\text{obj}.\$f$ is well-typed. The second rule is for the case when R 's nominal part is some range variable X and $X^+ \in \Gamma$, thus we know X is a super type of thisType and the members (excluding constructors) of X should be accessible via the this variable. In this case, if P is a field or a method signature, we extend the member typing contexts of both X and thisType with $\delta(P)$; or if P is a constructor signature, then only the member typing context of X is extended, because a constructor is not inheritable. In the following code:

```

1 trait T<this-class X>
2 provides {
3   public m() {
4     pattern <T, name f> public T $f in X {
5       T fx = this.$f;
6     }
7     pattern <> public constructor() in X {
8       X objx = new X();
9       thisType objths = new thisType(); // error
10    }
11  }
12 }

```

The expression of field access in line 5 is well-typed; the instance creation in line 8 is well-typed, but the instance creation in line 9 is not well-typed. The third rule is for the case when R 's nominal part is some range variable X with $X^- \in \Gamma$, or it is the thisType variable. If it is X , then we extend X 's member typing context with $\delta(P)$, otherwise we extend the thisType 's member typing context with $\delta(P)$. The third rule type-checks the following code:

```

1 trait T<class X> requires { X obj; }
2 provides {
3   public m() {
4     pattern <T, name f> public T $f in X {
5       T fobj1 = obj.$f;
6       T fobj2 = this.$f; // error
7     }
8     pattern <T, name f> public T $f in thisType {
9       T fths1 = this.$f;
10      T fths2 = obj.$f; // error
11    }
12  }
13 }

```

Type variables X and thisType are unrelated (i.e. no subtype relations). In the above code, the first pattern only allows us to access f via an instance of X , so the field access in line 6 is not well-typed. The second pattern only allows us to access f via the this variable, so the field access in line 10 is not well-typed.

In FJ, expression typing uses the additional functions for finding a field/method type. In contrast, because the member typing context of a class changes within different environments, thus the rules for

typing expressions in Figure 8 obtain the type of a member via a member typing context.

Method typing: $\boxed{\Gamma; \Theta \vdash M \text{ ok in } N}$

$$\frac{\Gamma \vdash \bar{T}, T \text{ ok} \quad \Gamma, \bar{x} : \bar{T}, \text{this} : C; \Theta \vdash s : U \quad U <: T \quad \text{class } C \text{ extends } C_0 \{ \dots \} \text{ uses } \dots \quad (\Theta^{CL}(C_0))(m) = \bar{U} \rightarrow U \text{ implies } U = T \text{ and } \bar{U} = \bar{T}}{\Gamma; \Theta \vdash T \text{ m}(\bar{T} \bar{x})\{s; \} \text{ ok in } C}$$

$$\frac{\Gamma \vdash \bar{T}, T \text{ ok} \quad \Gamma, \bar{x} : \bar{T}, \text{this} : \text{thisType}; \Theta \vdash s : U \quad U <: T}{\Gamma; \Theta \vdash T \text{ m}(\bar{T} \bar{x})\{s; \} \text{ ok in } \text{thisType}}$$

Trait application typing: $\boxed{\Gamma; \Theta \vdash E \text{ ok in } N}$

$$\frac{\text{trait } \mathcal{T} \langle \bar{X}^o \rangle \text{ req } \{ \bar{F}; \bar{H} \} \text{ prov } \{ \dots \} \quad \Gamma \vdash \bar{R} \text{ ok} \quad \Gamma \vdash \bar{N}(\bar{R}) \text{ in } N \Rightarrow \bar{o} \quad \Theta(N) = \Delta \quad \text{for each } T \text{ f in } \bar{F}: \Delta(\text{f}) = U \text{ and } U = T \quad \text{for each } T \text{ m}(\bar{T}) \text{ in } \bar{H}: \Delta(m) = \bar{U} \rightarrow U \text{ and } U = T \text{ and } \bar{U} = \bar{T}}{\Gamma; \Theta \vdash \mathcal{T} \langle \bar{R} \rangle \text{ ok in } N}$$

Trait typing: $\boxed{\text{TR ok}}$

$$\frac{\Gamma = \bar{X}^o, \text{thisType} \quad \Gamma \vdash \bar{F} \text{ ok} \quad \Gamma \vdash \bar{H} \text{ ok} \quad \Theta = \Theta^{CL} \uplus [\text{thisType} \mapsto \Theta^{TR}(\mathcal{T})] \quad \Gamma; \Theta \vdash \bar{M} \text{ ok in } \text{thisType} \quad \Gamma; \Theta \vdash \bar{E} \text{ ok in } \text{thisType}}{\text{trait } \mathcal{T} \langle \bar{X}^o \rangle \text{ req } \{ \bar{F}; \bar{H} \} \text{ prov } \{ \bar{M}; \text{use } \bar{E} \} \text{ ok}}$$

Meta class typing: $\boxed{\text{MC ok}}$

$$\frac{K = C(\bar{C}_0 \bar{F}_0, \bar{C} \bar{F})\{\text{super}(\bar{F}_0); \text{this}.\bar{f}=\bar{f}; \} \quad \text{fields}(C_0) = \bar{C}_0 \bar{F}_0 \quad \emptyset; \Theta^{CL} \vdash \bar{M} \text{ ok in } C \quad \emptyset; \Theta^{CL} \vdash \bar{E} \text{ ok in } C}{\text{class } C \text{ extends } C_0 \{ \bar{C} \bar{F}; K \bar{M} \} \text{ uses } \bar{E} \text{ ok}}$$

Figure 9. Method, trait application, trait, and class typing.

Figure 9 shows the rules for typing methods, trait applications, traits, and meta classes.

There are two rules for typing a method. The first rule type-checks a method that resides in some class C , then the type of the `this` variable is C and the validity of method overriding needs to be checked. The second rule type-checks a method that resides in a trait, then the type of the `this` variable is the `thisType` variable.

For the rule of type-checking a trait application, we have function $\Gamma \vdash \bar{N}(\bar{R})$ in $N \Rightarrow \bar{o}$ that computes the `this` access control modifiers for a sequence of ranges. According to the definition of function N , the value of $N(\bar{R})$ can be either a type variable or a class name. The definition of function $\Gamma \vdash \bar{N}(\bar{R})$ in $N \Rightarrow \bar{o}$ is shown as follows:

$$\frac{X^o \in \Gamma}{\Gamma \vdash X \text{ in } N \Rightarrow o} \quad \frac{N <: N_0}{\Gamma \vdash N_0 \text{ in } N \Rightarrow +} \quad \frac{N \not<: N_0}{\Gamma \vdash N_0 \text{ in } N \Rightarrow -}$$

The rule for typing a trait application checks if the modifiers of the parameters are consistent with the modifiers of the trait arguments; and also checks if the trait requirements are satisfied in the use context N .

For the rule of typing a trait declaration, we add the `thisType` variable into the binding context so that the `thisType` can appear inside a trait. The initial structural typing context is

$$\Theta^{CL} \uplus [\text{thisType} \mapsto \Theta^{TR}(\mathcal{T})]$$

so that we can access a member via the `this` variable in a trait.

For a meta class, we do not allow an imported method to override a method in the meta class or in the meta class's superclasses. Such restriction is checked when the compiler computes the structural typing context for the meta class before type-checking. So in

the rule of typing a meta class, we do not need to check if an imported method correctly overrides a method in a superclass.

4.3 Meta Evaluation

The meta evaluation of a metaprogram includes (1) pattern-matching and code generation, and (2) trait flattening (i.e. methods from traits are inlined into classes). Because many previous papers have fully discussed the process of trait flattening [7, 19, 21, 22], we here focus on the former. Because pattern-matching and code generation is performed at the statement level while trait flattening is performed at the member level, there is no side effect if we discuss the evaluate rules for those separately.

A statement-level pattern is evaluated into a sequence of statements. In Figure 10, we present part of the reduction rules for patterns. The rules are for the pattern-matching of fields. Those rules can be applied to the pattern-matching of methods or constructors with some slight modification. In the figure, the first rule is for the case when the pattern matches a field (it can be implemented by the unification algorithm), then it generates a sequence of instantiated statements. If the local variables in a generated statement have name conflicts with other local variables, we perform α renaming to substitute those local variable for fresh ones. The second rule is for the case when the match fails, then the pattern continues to pattern-match the rest fields. The third rule is for the case when there is no field, then the pattern generates an empty statement sequence.

4.4 Soundness

In this section, we give the properties of our meta language. Our purpose is to show that a metaprogram always generates a piece of well-typed code.

Lemma 2 (Staged Type Preservation of Statements). *For some meta-level non-pattern statement s (i.e. statement without patterns), if $\Gamma; \Theta \vdash s : C$ where $\text{TyVars}(\Gamma) = \emptyset$, then when applying object-level type-checking to s , we have $\Gamma' \vdash s : C$ for some object-level typing environment Γ' .*

Definition 1 (Class Name Alias \tilde{N}). *An alias of a class name, written \tilde{N} , is a distinct class name. A class name can have multiple name alias.*

(1) *An alias of a class name preserves the nominal subtyping relations of that class.*

(2) *For any class alias \tilde{N} , $\tilde{N} \notin \text{dom}(\Theta^{CL})$.*

(3) *For some class name alias \tilde{N} and some Θ , if \tilde{N} does not appear in the domain of Θ , then $\Theta(\tilde{N}) = \Theta^{CL}(N)$; otherwise because there exists some Δ such that $\tilde{N} \mapsto \Delta \in \Theta$, $\Theta(\tilde{N}) = \Delta$.*

Lemma 3 (Type Substitution Preserves Typing).

i) *Suppose $\text{this} : N_0 \in \Gamma, \Gamma'$; \tilde{N} is an alias of N ; and $\tilde{N} \notin \text{dom}(\Theta)$. If $\Gamma, X^o, \Gamma'; \Theta \vdash e : U$, $\Gamma \vdash N \text{ ok}$, and $\Gamma \vdash N(N)$ in $N_0 \Rightarrow o$, then $\Gamma, [\tilde{N}/X]\Gamma'; [\tilde{N}/X]\Theta \vdash [\tilde{N}/X]e : [\tilde{N}/X]U$.*

ii) *Suppose $\text{this} : N_0 \in \Gamma, \Gamma'$; \tilde{N} is an alias of N ; and $\tilde{N} \notin \text{dom}(\Theta)$. If $\Gamma, X^o, \Gamma'; \Theta \vdash ns \text{ ok}$, $\Gamma \vdash N \text{ ok}$, and $\Gamma \vdash N(N)$ in $N_0 \Rightarrow o$, then $\Gamma, [\tilde{N}/X]\Gamma'; [\tilde{N}/X]\Theta \vdash [\tilde{N}/X]ns \text{ ok}$.*

Lemma 4 (Name Substitution Preserves Typing).

i) *If $\Gamma, \eta, \Gamma'; \Theta \vdash e : U$, and $\Gamma \vdash 1 \text{ ok}$, then $\Gamma, \Gamma'; [1/\eta]\Theta \vdash [1/\eta]e : U$.*

ii) *If $\Gamma, \eta, \Gamma'; \Theta \vdash ns \text{ ok}$, and $\Gamma \vdash 1 \text{ ok}$, then $\Gamma, \Gamma'; [1/\eta]\Theta \vdash [1/\eta]ns \text{ ok}$.*

Lemma 5 (Statement Concatenation Preserves Typing). *Suppose function $lvar$ collects all the local variables in a sequence of statements. If $\Gamma; \Theta \vdash ns_1 \text{ ok}$, $\Gamma; \Theta \vdash ns_2 \text{ ok}$, and $lvars(ns_1) \cap lvars(ns_2) = \emptyset$, then $\Gamma; \Theta \vdash ns_1; ns_2 \text{ ok}$.*

$$\begin{aligned}
\text{pattern}\langle\bar{X},\bar{\eta}\rangle T f \text{ in } C \diamond \{(T_0 f_0),\bar{F};Q;\bar{H}\} \{ ns; \} &\longrightarrow \frac{[\bar{1}/\bar{\eta}][\bar{T}/\bar{X}]ns; (\text{pattern}\langle\bar{X},\bar{\eta}\rangle T f \text{ in } C \diamond \{\bar{F};Q;\bar{H}\} \{ ns; \})}{\text{where } \exists\bar{T}.\exists\bar{1} \text{ such that } [\bar{T}/\bar{X}]T = T_0 \text{ and } [\bar{1}/\bar{\eta}]f = f_0} \\
\text{pattern}\langle\bar{X},\bar{\eta}\rangle T f \text{ in } C \diamond \{(T_0 f_0),\bar{F};Q;\bar{H}\} \{ ns; \} &\longrightarrow \text{pattern}\langle\bar{X},\bar{\eta}\rangle T f \text{ in } C \diamond \{\bar{F};Q;\bar{H}\} \{ ns; \} \\
&\text{where } \exists\bar{T}.\exists\bar{1} \text{ such that } [\bar{T}/\bar{X}]T = T_0 \text{ and } [\bar{1}/\bar{\eta}]f = f_0 \\
\text{pattern}\langle\bar{X},\bar{\eta}\rangle T f \text{ in } C \diamond \{[];Q;\bar{H}\} \{ ns; \} &\longrightarrow []
\end{aligned}$$

Figure 10. Pattern-matching and code generation using a field pattern.

Theorem 3 (Type Preservation for Patterns). *If $\Gamma; \Theta \vdash ps \text{ ok}$ and $ps \longrightarrow ns$, then $\Gamma; \Theta \vdash ns \text{ ok}$*

Theorem 4 (Progress for Patterns). *If $\emptyset; \Theta \vdash ps \text{ ok}$ and ps is in the form of $\text{pattern}\langle\bar{X},\bar{\eta}\rangle P \text{ in } R \{ ns_0 \}$, then there is some ns with $ps \longrightarrow ns$ (suppose ns includes an empty list of statements).*

5. ORM via Compile-Time Reflection

Object-relational mapping (ORM for short) is the ability of mapping objects from/to records in a relational database. Some Java ORM tools, such as Hibernate⁵, use run-time reflection, which may cause some runtime performance overhead.

We present an ORM tool, called PtjORM, which is a real-world application of pattern-based traits. PtjORM is not as powerful as Hibernate, but it supports compile-time reflection and avoids the overhead of runtime reflection. Like many other ORM tools, PtjORM supports property mapping, association mapping, and inheritance mapping. Current version of PtjORM supports two kinds of inheritance mapping strategies: table per subclass and table per concrete class.

Besides the type of reflection, the performance of an ORM tool can be influenced by other factors, such as object saving/fetching mode and the number of executed SQL statements. So, instead of evaluating compile-time reflection solely, we evaluate the overall performance of PtjORM.

We evaluated PtjORM by testing it with a benchmark based on the 007 benchmark [2]. The 007 benchmark was originally implemented in C++. It tests the performance of persistence for a CAD application. Ibrahim and Cook provided its Java implementation [14]. Our modified 007 benchmark generates the databases of three different sizes: tiny, small, and medium. The tiny one has 996 records; the small one has 11,465 records, and the medium one has 75,779 records. We compare the runtime performance of saving and fetching objects for the following four ORM tools: PtjORM (version 2.0), Hibernate (version 4.1), Ebean (version 2.7.7), and EclipseLink (version 2.4.1). In default, they all use some Java bytecode generation and manipulation (runtime) tool for reflection optimization.

We ran the 007 benchmark on a laptop with a 2.2 GHz, Intel[®] Core[™] i7 processor (4 cores) and 8GB memory. The MySQL database was installed on a desktop with a 1.80GHz Intel[®] Pentium[™] E2160 dual processor and 1GB memory. The desktop and the laptop were connected in the same local network.

5.1 Object Saving Evaluation

First, we give Table 1, which shows the number of the SQL statements executed during the object saving tests. While the others used the table per hierarchy strategy, PtjORM used the table per subclass strategy for inheritance mapping and it generated the largest number of SQL statements.

The ORM tools were tested with MySQL database. The test results about the object-saving performance is show in Table 2. The test results show that even with the largest number of executed SQL

statements, PtjORM achieved the best object-saving performance. Approximately, PtjORM is 20% faster than Hibernate, 24% faster than EclipseLink, and 15% faster than Ebean.

ORM tools	DB size		
	tiny	small	medium
Ebean	1111	11.7K	75.9K
EclipseLink	1071	11.5K	75.7K
Hibernate	1005	11.6K	75.6K
PtjORM	1125	12.6K	76.7K

Table 1. The number of SQL statements executed by each of the ORM tools for object saving.

ORM tools	DB size		
	tiny	small	medium
Ebean	955ms	9610ms	62476ms
EclipseLink	1120ms	10084ms	64357ms
Hibernate	1169ms	9817ms	60178ms
PtjORM	823ms	8258ms	50365ms

Table 2. Time used for object-saving. Unit ms is short for millisecond.

5.2 Object Fetching Evaluation

We use the batch fetching mode for object fetching. For fetching a tiny database, we use batch sizes: 4, 8, 12, 40, 70, and 100. For fetching a small database, we use batch sizes: 40, 100, 160, 400, 700, and 1000. For fetching a medium database, we use batch sizes: 200, 600, 1000, 2000, 3500, and 5000. The tables in Table 3 show the number of SQL statements executed by each of the ORM tools for fetching the databases of three different sizes. From the tables, we note that PtjORM generates the largest number of SQL statements when the batch size is small. With the increase of batch size, the number of SQL statements generated by PtjORM drops dramatically. When fetching a large number of objects, we suggest PtjORM users to assign a large batch size for efficiency. For EclipseLink and Hibernate, increasing the batch size does not have a big influence on the number of SQL statements.

The tables in Table 4 show the time used by each of the ORM tools for fetching the databases of three different sizes. We learned that Ebean has the best performance and one reason is that it generates the least number of SQL statements. For fetching a tiny database, PtjORM is 2%–70% faster than Hibernate, -2%–70% faster than EclipseLink. When fetching a small database, except for the batch size of 40, PtjORM is 8%–58% faster than Hibernate; and PtjORM is 60% faster than EclipseLink in average. When fetching a medium database, PtjORM does not perform well with small batch sizes, but with larger batch sizes, its performance exceeds the performance of Hibernate and EclipseLink. With a small batch size, because PtjORM generates a greater number of SQL statements, the cost of executing the SQL statements and commuting with the server outweighs the saving from compile-time reflection.

⁵ <http://www.hibernate.org>

tion. But when the batch size is equal or larger than 3500, its performance exceeds the performance of Hibernate and EclipseLink. In summary, PtjORM has the second best performance for fetching objects. Of course, from the benchmark test results, we feel the need to improve PtjORM by reducing the number of SQL statements that PtjORM generates.

tools \ batch size	4	8	12	40	70	100
Ebean	203	118	91	50	41	37
EclipseLink	159	130	125	149	161	160
Hibernate	228	185	153	142	142	140
PtjORM	458	284	208	79	63	54

tools \ batch size	40	100	160	400	700	1000
Ebean	171	83	65	48	44	41
EclipseLink	431	314	303	290	288	286
Hibernate	428	395	368	335	358	374
PtjORM	976	469	287	141	109	74

tools \ batch size	200	600	1000	2000	3500	5000
Ebean	160	75	58	47	43	41
EclipseLink	454	404	395	389	394	393
Hibernate	615	485	518	527	526	523
PtjORM	1673	648	404	295	178	126

Table 3. The number of SQL statements executed by each of the ORM tools for fetching the databases of three sizes

6. Related Work

For Composition FeatherTrait Java (FTJ) [19] extends Featherweight Java with traits. The FTJ does not support the compile-time reflection but its type system is modular. Chai [25] also extends Java with traits. Chai allows traits to be more generally used: a trait not only performs a building block for a class but also creates a subtype relation with the class. Moreover, when used, a trait can be dynamically substituted by another trait with the same interface. Featherweight Jigsaw [18] applies the concept of traits into classes and therefore, classes are also building blocks with a set of composition operators. Jigsaw uses the symmetric sum for composition. Besides the member-level operations, Jigsaw introduces member modifiers to solve name conflicts. Reppy and Turon [22] introduced the traits that can be parametrized over names, types, and values. These traits offer the same ability of pattern-based reflection and can be used to generate fields and methods for a class. MixML [4] extends ML modules with imports and outputs (like require and provide in a trait) for flexible composition. MixML does not flatten a composed module. Instead, it keeps module hierarchy. So, access to a member in a module may need to refer to its namespace.

For Reflection Programming languages such as Genoupe [3], SafeGen [13], CTR [6], and MorphJ [11] support reflection. Genoupe introduces a type system for reflective program generators, which offers an particular high degree of static safety for reflection. However, the type system does not guarantee that generated code is always well-typed. SafeGen uses first-order logic formulae to express patterns and those formulae are used by a theorem prover to check the safety of generated code. Therefore, the type system of SafeGen is undecidable. CTR introduces transforms, which support pattern-based compile-time reflection and static type safety. MorphJ refines the type system of CTR with a modular type system. Considering its power for static reflection, we

believe that MorphJ is the calculus closest relative to ours. However, one reflective power of MorphJ that we do not have is negative nested patterns, which can be used to prevent name conflicts. MorphJ cannot extend a class in place while our pattern-based traits can.

For Code Generation Some programming languages, such as MetaML [26], MetaOCaml, and Template Haskell [24], enable metaprogramming by providing multiple stages of computation, where earlier stages can manipulate code for late stages. These languages use explicit stage annotations for the support of code manipulation at expression level. Our programming language supports (two-staged) metaprogramming: pattern-matching, code generation, and trait flattening are performed at compile time while generated program is evaluated at run time. Another difference is that these metaprogramming languages do not support compile-time type reflection. In paper [10], Garcia presented a metaprogramming language that has type-reflection at the meta level, but it does not support generics.

Comparison with Aspect-Oriented Programming Aspect-oriented programming (AOP) [17] intends to separate cross-cutting concerns. Though not dedicated for AOP, our pattern-based traits can be used to describe some cross-cutting concerns. Compared with AspectJ [16], which weaves advice code into original application code at the Java bytecode level, pattern-based traits generates source code, which enables a programmer to inspect the effect of a metaprogram. The current stable version of AspectJ, that is AspectJ 5, supports Java generics, but it does not allow the use of type variables and name variables. AspectJ inserts code flexibly, allowing users to define different point cuts for insertion, but pattern-based traits enable us to merely add wrapper code for members.

References

- [1] David Abrahams and Aleksey Gurtovoy. *C++ Template Metaprogramming: Concepts, Tools, and Techniques from Boost and Beyond (C++ in Depth Series)*. Addison-Wesley Professional, 2004.
- [2] Michael J. Carey, David J. DeWitt, and Jeffrey F. Naughton. The 007 benchmark. In *Proceedings of the 1993 ACM SIGMOD international conference on Management of data, SIGMOD '93*, pages 12–21, New York, NY, USA, 1993. ACM.
- [3] Dirk Draheim, Christof Lutteroth, and Gerald Weber. A type system for reflective program generators. In *Proc. of the Conf. on Generative Programming and Component Engineering*. Springer, 2005.
- [4] Derek Dreyer and Andreas Rossberg. Mixin' up the ml module system. In *Proc. of the 13th ACM SIGPLAN International Conference on Functional Programming*. ACM, 2008.
- [5] Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, and Andrew P. Black. Traits: A mechanism for fine-grained reuse. *ACM Trans. Program. Lang. Syst.*, 2006.
- [6] Manuel Fähndrich, Michael Carbin, and James R. Larus. Reflective program generation with patterns. In *Proc. of the Conf. on Generative Programming and Component Engineering*. ACM, 2006.
- [7] Kathleen Fisher. A typed calculus of traits. In *In Workshop on Foundations of Object-oriented Programming*, 2004.
- [8] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In *Proc. of the Symposium on Principles of Programming Languages*. ACM, 1998.
- [9] Ira R. Forman and Nate Forman. *Java Reflection in Action (In Action series)*. Manning Publications Co., Greenwich, CT, USA, 2004.
- [10] Ronald Garcia and Andrew Lumsdaine. Toward foundations for type-reflective metaprogramming. In *Proc. of the Conf. on Generative Programming and Component Engineering*. ACM, 2009.
- [11] Shan Shan Huang and Yannis Smaragdakis. Expressive and safe static reflection with morphj. In *Proc. of the Conf. on Programming Language Design and Implementation*. ACM, 2008.

batch size \ tools	4	8	12	40	70	100
Ebean	345ms	254ms	217ms	159ms	148ms	140ms
Eclipselink	514ms	484ms	478ms	564ms	632ms	622ms
Hibernate	639ms	599ms	608ms	591ms	625ms	586ms
PtjORM	627ms	458ms	353ms	224ms	195ms	186ms

batch size \ tools	40	100	160	400	700	1000
Ebean	677ms	541ms	509ms	458ms	461ms	456ms
Eclipselink	7964ms	2465ms	2479ms	2390ms	2440ms	2391ms
Hibernate	1505ms	1463ms	1423ms	1365ms	1405ms	1421ms
PtjORM	2167ms	1352ms	1103ms	826ms	697ms	601ms

batch size \ tools	200	600	1000	2000	3500	5000
Ebean	2473ms	2287ms	2299ms	2241ms	2281ms	2285ms
Eclipselink	8337ms	8287ms	8165ms	7704ms	7746ms	7706ms
Hibernate	4970ms	4648ms	4743ms	4952ms	4952ms	5139ms
PtjORM	18365ms	9869ms	6802ms	5310ms	4453ms	4253ms

Table 4. Time used for fetching objects from the databases of three sizes

- [12] Shan Shan Huang, David Zook, and Yannis Smaragdakis. Morphing: Safely shaping a class in the image of others. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer-Verlag, 2007.
- [13] Shan Shan Huang, David Zook, and Yannis Smaragdakis. Statically safe program generation with safegen. *Sci. Comput. Program.*, 2011.
- [14] Ali Ibrahim and William R. Cook. Automatic prefetching by traversal profiling in object persistence architectures. In *Proceedings of the 20th European conference on Object-Oriented Programming, ECOOP'06*, pages 50–73, Berlin, Heidelberg, 2006. Springer-Verlag.
- [15] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.*, 2001.
- [16] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William G. Griswold. An overview of aspectj. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer-Verlag, 2001.
- [17] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean marc Loingtier, and John Irwin. Aspect-oriented programming. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer-Verlag, 1997.
- [18] Giovanni Lagorio, Marco Servetto, and Elena Zucca. Featherweight jigsaw: A minimal core calculus for modular composition of classes. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer-Verlag, 2009.
- [19] Luigi Liquori and Arnaud Spiwack. Feathertrait: A modest extension of featherweight java. *ACM Trans. Program. Lang. Syst.*, 2008.
- [20] Weiyu Miao and Jeremy Siek. Pattern-based traits. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing, SAC '12*, pages 1729–1736. ACM, 2012.
- [21] Oscar Nierstrasz, Stéphane Ducasse, and Nathanael Schärli. Flattening traits. *Journal of Object Technology*, 2006.
- [22] John Reppy and Aaron Turon. Metaprogramming with traits. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer, 2007.
- [23] Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Andrew P. Black. Traits: Composable units of behaviour. In *Proc. European Conf. on Object-Oriented Programming*. Springer, 2003.
- [24] Tim Sheard and Simon Peyton Jones. Template metaprogramming for Haskell. In *Haskell Workshop '02*. ACM, 2002.
- [25] Charles Smith and Sophia Drossopoulou. *Chai: Traits for Java-Like Languages*. In *Proc. of the European Conf. on Object-Oriented Programming*. Springer, 2005.
- [26] Walid Taha and Tim Sheard. Multi-stage programming with explicit annotations. In *Proc. of the 1997 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation*. ACM, 1997.