


Metaprogramming in Cyan

José de Oliveira Guimarães

(Federal University of São Carlos, Sorocaba, SP, Brazil)

 <https://orcid.org/0000-0003-4128-0419>, jose@ufscar.br

Abstract: Certain languages allow a metaprogram to act as a compiler plugin and thus alter the compilation process. The metaprogram interacts with low-level details of the compiler, making its construction difficult and potentially leading to errors. Different parts of the metaprogram may have conflicting interactions, thus producing unintended outcomes. This article introduces metaprogramming in the prototype-based object-oriented language Cyan. This language provides the same core functionality as other metaprogramming systems while introducing features that improve interactions between the compiler and different components of the metaprogram. Furthermore, Cyan incorporates security measures designed to circumvent typical issues encountered in metaprogramming.

Keywords: object-oriented languages, metaprogramming, metaobject, computational reflection, prototype-based languages, compilers

Categories: H.3.1, H.3.2, H.3.3, H.3.7, H.5.1

DOI: 10.3897/jucs.141599

1 Introduction

Metaprogramming is the coding of programs, called *metaprograms*, that treat code as data [Lilis, Y. and Savidis 2019]. The *base program*, or simply the *program*, is the program that is treated as *data*. A *metaprogram* can generate new code, change existing programs, or do checks on them. Metaprogramming provides more advanced code reuse mechanisms than traditional software libraries. It can generate families of related code, as in the case of C++ templates [Stroustrup 2013]; separate functional and nonfunctional concerns, as in AspectJ [Kiczales et al. 2001]; generate code based on specifications, as ANTLR 4 does [Parr 2013]; support new syntax, such as Scala macros [Burmako 2013]; detect program bugs through static analyzers¹; implement new type systems using a pluggable-type system [Bracha 2004]; run a program in multiple stages [Taha 2007], each stage generating and running a new program; change the program at runtime [Redmond and Cahill 2002, Kamin et al.]; and support embedded Domain-Specific Languages [Rompf 2012, Biboudis et al. 2016].

In this paper, the focus is *language support* for Compile-Time Metaprogramming (CTMP), which is the handling of a program by a metaprogram at compile-time. To discuss specific characteristics of compile-time metaprogramming supported by programming languages, we will define some terms. The *program*, or *base code*, is the code that implements the desired functionality for the application. A *metacode* is each of the pieces of code that compose the *metaprogram*. The *metacode* is loaded at the compilation time of the *base code* and works as an integral part of the compiler. Data is exchanged between the compiler and the metacode. The compiler calls metacode at

¹ <https://github.com/google/error-prone>, <https://spotbugs.readthedocs.io/en/stable/>

specific points of compilation. Therefore, metacode can interact with the type checker, code generator, parser, and any other algorithm used by the compiler. They can also add, delete, or replace code in the *program*. In practice, languages restrict what metacode can do to a few things. The metaprogram is designed to help the program achieve the desired functionality.

Metacode interacts closely with the compiler. This interaction is precisely defined through a *protocol*, which specifies which part of each metacode is called in a given compilation phase and how data is passed from and to the compiler. For example, the protocol may specify that during the compilation phase *parsing*, a function² or method *duringParsing* of the metacode may be called. The function or method may perform checks or add code to the program. The protocol would also specify which parameters the function takes, what it returns, and how the return changes the compilation.

Although every language supporting CTMP has a protocol for metacode-compiler interactions, some of them, particularly older languages, are explicitly designed with a Metaobject Protocol (MOP). This paper discusses their common characteristics. Let us evaluate languages supporting CTMP with or without a MOP to show their deficiencies and opportunities for improvement. A typical Metaobject Protocol associates a metaclass to every class and, in general, a meta-*S* to every kind of declaration *S*. Developers can define their own metaclass for a class, which directs the compilation of the class. Using a metaclass, one can add code to the class, do additional checks on it, and change the semantics of the language for that class. For example, the metaclass can check if all class fields are read-only, intercept object creation, and change the meaning of inheritance. However, there are limitations. Usually, only one metaclass is associated with a class and the composition of behavior is difficult. Arguments cannot be passed to the metaclass, which makes its configuration impossible. There may be limitations on what can be changed, as in OpenC++ [Chiba 1995] or OJ [Tatsubori et al. 2000]. Or there may be very few limitations, as in CLOS, but the AST of the class associated with the metaclass is not available. This makes building the metaclass more difficult.

A language may support CTMP without a MOP such as Groovy [König 2007], Scala³, or Java⁴. In general, support for CTMP is at a much lower level than when using a MOP. The metacode developer has to interact with low-level compiler algorithms and data structures. It is easy for the metacode to invalidate compiler invariants or damage internal compiler structures, making the compiler crash.

There are some problems that occur with metaprogramming languages with or without a MOP, although mainly with the latter. Not all problems that follow occur in all languages, although most of them do occur in languages without a MOP. The following list is not exhaustive.

- (a) A metacode associated with a source file can change the base code of another file. Hence, to know the final version of a file, the developer has to know the behavior of every metacode in the metaprogram.
- (b) The compiler triggers the recompilation of a class if any of the classes it depends on changes. Without metacode, the compiler always knows the dependencies of a class *A* because it handles the data structures needed to do checking or generate code for this class. However, a metacode may produce code for *A* depending on a class *B* and the compiler is not aware of this linking. The compiler cannot detect that a change to *B*

² Function as in language C

³ <https://docs.scala-lang.org/scala3/reference/changed-features/compiler-plugins.html>

⁴ <https://docs.oracle.com/en/java/javase/11/docs/api/jdk.compiler/com/sun/source/util/Plugin.html>

requires recompiling A.

(c) The developer has to know many details about the compiler since the metacode interacts with low-level compiler algorithms and data structures. Therefore, metaprogramming is not only harder, but the metacode can bypass compiler checks, invalidate invariants, damage data structures, and crash the compiler.

As an example, the metacode may change the AST after the compiler type-checked the code. The changed AST may have type errors, which may crash the compiler during the following compilation phase.

(d) Metacode can insert new code into the base code. If the new code has errors, the compiler is unable to point out which metacode is responsible for adding the new code. This is because the metacode itself changes the compiler AST to add new AST objects. No trace is left of who did what.

As an example, suppose that, during the compilation, the compiler calls the metacode M that has the following statement.

```
stats.add( new MessagePassing(...) );
```

stats refer to the list of statements of a method named search (stats was created by the compiler). MessagePassing is the AST object that represents a message passing. Hence, the above statement adds a message passing to method search. The compiler does not record that the statement MessagePassing was inserted by M. If it has errors, the compiler does not point out this metacode as its source.

(e) The order in which the metacode is called is not specified by the language. It can even change between compilations (with the same base code). By changing the metacode calling order, their view of the program can change if a metacode can view the code added by a metacode executed previously.

As an example, the metacode may be specified by a compiler option of a fictitious language Tlip:

```
compiler -processor P -processor Q A.tlip
```

Metacode P adds a field count to the class A and Q creates get and set methods for all A fields. This works fine because P is executed before Q (assume this). If we reverse the -processor options, the get and set methods for count will not be created.

(f) A verification performed by one metacode may become obsolete due to modifications introduced by another metacode. As an example, a metacode P checks that all fields of a class A have get and set methods. After that, another metacode Q adds a field ghost but not get and set methods. The check that P made was invalidated by Q.

(g) Metacode may generate metacode that may generate metacode ad infinitum. As an example, metacode M produces code containing references to M (as a macro M producing calls to M). If the compiler processes the newly created code, there will be an infinite creation of code.

(h) Metacode may not terminate its computation. The compiler will not terminate either.

(i) Metacode may unintentionally generate code that uses identifiers already in use. This is the hygiene problem of macros [Kohlbecker et al. 1986].

The Cyan MOP has two main goals. The first one is to build general-use metaobjects such as property to create get and set methods for a field, annot to associate information to declarations, checkStyle to verify the style of code, several metaobjects that give information on the source code (as the line number of the annotation), eval to evaluate code at compile-time, deprecated to issue a warning if a deprecated declaration is used, and many others.

The second goal of the Cyan MOP is to build metaobjects associated with developers'

packages (libraries). These metaobjects would generate code and check the use of the packages. As examples of code generation, package `cyan.lang` uses metaobjects to generate code for several generic prototypes, such as `Tuple` and `Array`. Based on the real argument, such as `Int` in `Array<Int>`, metaobjects add some methods to this prototype.

The semantics of methods, which are usually written in the documentation, can be put in metaobjects that do checks at compile-time. An example of that is method `printf` of prototype `Out`. Similar to the function with the same name in language C, this method accepts a format string as the first argument. A metaobject checks at compile-time if the arguments match the first argument. The semantics of inheritance can also be checked by metaobjects. Annotation `@overrideToo("hashCode")` attached to method `==` of the top Cyan prototype, `Any`, demands that, whenever `==` is overridden in a subprototype, method `hashCode` should be overridden too. A metaobject whose annotation is attached to a method may request that this method be called in every method that overrides it in a subprototype. These examples show that packages' semantics can be checked by metaobjects.

The paper's organization is as follow. Section 2 is a brief introduction to the Cyan language. The Metaobject Protocol of Cyan is explained in Section 3. Section 4 compares the metaprogramming systems of other languages with the Cyan MOP. The last section concludes. Additional information on the Cyan compiler and the language is available at `cyan-lang.org`.

2 The Cyan Language

Cyan is a statically typed prototype-based object-oriented language. A *prototype* is a template from which other objects may be created, the same role is played by *classes* other object-oriented languages. The difference is that the prototype itself is an object like any other.⁵

2.1 Basic Elements of the Language

The look and feel of Cyan is that of a class-based language. The compiler translates Cyan into non-legible Java code. Thus, many language constructs are directly translated into Java, such as packages, inheritance, method overriding, message passing, assignment, and prototype declaration (each prototype is translated to a Java class). Cyan code can import Java packages and classes. The compiler does all the necessary conversions between names and values of the basic types.

Listing 1 shows the declaration of prototype `Student` of the package `university`. Cyan employs a syntax for method declaration and message passing that is in some ways similar to Smalltalk. A *unary message passing* is made up of a *receiver* and an identifier, which should be the name of the *unary method*:

```
aStudent getName
```

`aStudent` is the *message receiver*. A *keyword message passing* is composed of a *receiver* and one or more *message keywords*, or just *keywords* with their parameters:

```
// creates an object
var Array<String> as = Array<String> new;
// 'add:' is a message keyword, a method 'add:' is
called
```

⁵ There are exceptions to this rule, but this is not important to this paper.

```

package university

object Student
  // fields name and number
  var String name
  var Int    number
  // this is a constructor. Use 'Student new: name,
  number'
  // or 'Student(name, number)' to create an object of
  // Student
  func init: String name, Int number {
    self.name = name;
    self.number = number;
  }
  func getName -> String    = name;
  func setName: String name { self.name = name }
  func getNumber -> Int    { return number }
  func setNumber: Int number { self.number = number }
end

```

Listing 1: Prototype Student

```

as add: "first";
  // 'at:' and 'put:' are keywords, method 'at:put:' is
  called
as at: 0 put: "zero";

```

Both *method keywords* and *message keywords* are called *keywords*. To avoid confusion, *Cyan keyword* is used for reserved words in the language.

2.2 Examples of Using Annotations

This Subsection gives a general view of how annotations are used in Cyan code and what metaobjects can do. An *annotation*, or *metaobject annotation*, is the syntax element that links the program to a metacode. There are several kinds of annotations in Cyan. This Subsection will only give examples of the most general of them all, that which starts with “@”, as shown in Listing 2. At compile-time, each annotation is linked to a single *metaobject* (and vice versa). The metaobject is able to do checks, create new prototypes, and add code to the current prototype. In this example, the metaobject associated with `init` will add a constructor to the prototype `Student` to initialize the newly created object with the student’s name and number. The two metaobjects associated with the two property annotations will add get and set methods to `Student`. The resulting prototype is identical to the one in Listing 1. We say that the two property annotations are *attached* to the declarations of name and number. The annotation `init` is *attached* to the prototype `Student`. Basic literals (3, 3.14, 'A'), identifiers, literal arrays, literal tuples, literal maps, and any combination of these can be parameters for annotations.

```

package university

@init(name, number)
object Student
  @property var String name
  @property var Int number
end

```

Listing 2: Prototype Student with annotations

```

1 object Program
2   func run { ... }
3   @doc{* returns the double of the argument. The method call
4     is replaced by the expression 2*n *}
5   @replaceCallBy(once){* 2*n *}
6   func twice: Int n -> Int = n + n;
7 end

```

Listing 3: Annotations with attached text between delimiters

Annotation `init` takes two identifiers as arguments, `name` and `number`, which are treated as strings by the metaobject.

The *metaobjects associated with a prototype P*, or *metaobjects of P*, are the metaobjects associated with annotations of prototype P. There are three metaobjects associated with `Student` of Listing 2. We will use “*metaobject init*” when no confusion may arise. If there are two `init` annotations in a code, “*metaobject init*” will be ambiguous because it may refer to metaobjects associated with both annotations.

Annotations with an Attached DSL

An annotation may have a text or DSL code after its parameters (or the annotation, if it does not take arguments) between *delimiters*. The usual *delimiters* are `{* and *}`, as shown in lines 3-5 of Listing 3. The text between delimiters will be called *attached text* or *attached Domain-Specific Language (DSL) code*. It works like a last literal string argument to the annotation.

The DSL of annotation `replaceCallBy` (line 5) is a subset of Cyan in which only expressions are legal. Any text is valid for the `doc` annotation (line 3). It is intended as the documentation for the declaration to which the annotation is attached (`twice`: in this case).

Annotations can be used inside expressions if they were designed with this goal, as shown in the following example.

```

var factorial_of_10 = @eval("cyan.lang", "Int"){*
  var r = 2;
  for n in 3..10 { r = r*n }
  return r
}

```

```

@onOverride{*
  if (method getStatementList: env) getStatementList
                                     size < 10 {
    metaobject addError: "method should have >= 10 stats"
  }
*} func test { }

```

Listing 4: Annotation `onOverride` whose code is interpreted whenever the attached method is overridden

```

*});

```

The DSL code attached to `eval` is in language *Myan*, which is a dynamically typed simplified version of Cyan. The metaobject `eval` interprets the Myan code and returns 10! at compile-time.

Java packages can be imported and used within Cyan and Myan code. Therefore, one could read a file or open a web connection and create, based on it, a bunch of methods or fields in the current prototype.

Myan code has access to several compiler objects through the variables `method`, `metaobject`, and `env`. Listing 4 shows an example where the Myan code attached to the annotation `onOverride` is interpreted whenever the attached method (`test`) is overridden in a subprototype. It issues an error if the method has fewer than 10 statements.

Generic Prototypes and Metaobjects

Concepts [Stroustrup 2003] are constraints on real arguments for generic classes, functions, and prototypes, an idea that originated in CLU [Liskov et al. 1977]. Concepts in Cyan are implemented using metaobject concept without any help from the language itself. The DSL code attached to the annotation specifies the restrictions that the generic parameters should obey. That means developers can implement their own variation of this metaobject by copying and changing its source code. This is more flexible than hardwire concepts in the language.

3 The Cyan Metaobject Protocol

The Cyan Metaobject Protocol (MOP) describes the *interactions* between the Cyan code being compiled, the compiler, the MOP library, the metaprogram, and annotations in the Cyan code. The metaprogram in Cyan is composed of Java classes, Cyan prototypes, or a mixture of both. The compiler is implemented in Java making it convenient to use Java classes as the metaprogram. But since the compiler translates each Cyan prototype into a Java class, Cyan can also be used as the metaprogramming language.

Before presenting the MOP in detail, Subsection 3.1 shows how to build a metaobject class and how to use metaobject annotations. Although the description is tutorial-like, very concrete, this is necessary in order to understand many details of the rest of the paper. The interactions between the metacode and the compiler are presented in Subsection 3.2. Subsection 3.3 explains the phases of the Cyan compiler. The important features of the Cyan MOP are summarized in Subsection 3.4.

```

1 package art
2 import cyanPack
3 object Painting
4   func init: String name { self.name = name }
5   @get String name
6 end

```

Listing 5: Annotation get attached to field name

3.1 An Example of a Metaobject

This Subsection explains how to build a metaobject get whose annotation should be attached to a prototype field. The metaobject adds to the prototype a method `get_field` that just returns the attached field. Listing 5 shows annotation `get` attached to field `name` in line 5 and, therefore, the following method is created and added to `Painting`.

```
func get_name -> String = name;
```

We will give step-by-step instructions on how to build metaobject `get`. First, create a Java file that declares a class `CyanMetaobjectGet` that inherits from

```
CyanMetaobjectAtAnnot
```

Next, select an interface to implement based on the goals of the metaobject. Use Table 1 (end of the article). This table relates goals and the interfaces the metaobject class should implement to achieve them. As will be explained in detail later, the interfaces implemented by a metaobject class direct the compilation. They tell the compiler when (compilation phase) and which metaobject methods should be called.

The goal of metaobject `get` is to add a method. Hence, according to Table 1, class `CyanMetaobjectGet` should implement interface `IAction_afterResTypes`. Based on the documentation of this interface, we choose to override the method

```
afterResTypes_codeToAdd
```

that is used for adding fields and methods to the current prototype. The complete metaobject class is in Listing 6 (without the imports), which will be explained.

The annotation name is "get", it does not take parameters, and it should be attached to a field. This is all communicated in the call to `super` in lines 5-6. To code method `afterResTypes_codeToAdd` starting at line 9, we again remember the goals of the metaobject, which is to create a method whose name is based on the attached field. So, we need the name of the field attached to the annotation. Method `getAnnotation` in line 13 returns an object with information on the annotation. From it, the code retrieves the attached declaration using method `getDeclaration` (a field in this case). Since, by line 6, we know that the attached declaration is always a field, the code can cast the object returned to class `WrFieldDec`. This is the AST class for Cyan prototype fields. If the annotation `get` took arguments, these could be retrieved using another method of the object returned by `getAnnotation`. The attached field name is got in line 14 and the method source code, as a string, is created in lines 15-16. The return type of `afterResTypes_codeToAdd` is a tuple consisting of the source code of the fields and methods to be inserted in the current prototype (`Painting` in the example) and the signatures of these declarations. The signature of a method is the method without its

```

1 public class CyanMetaobjectGet extends CyanMetaobjectAtAnnot
2     implements IAction_afterResTypes {
3
4     public CyanMetaobjectGet() {
5         super("get", AnnotationArgumentsKind.ZeroParameters,
6             new AttachedDeclarationKind[] { AttachedDeclarationKind
7                 .FIELD_DEC });
8     }
9     @Override
10    public Tuple2<StringBuffer, String> afterResTypes_codeToAdd(
11        ICompiler_afterResTypes compiler,
12        List<Tuple2<WrAnnotation, List<ISlotSignature>>> infoList) {
13
14        final WrFieldDec field = (WrFieldDec) this.getAnnotation().getDeclaration();
15        final String name = field.getName();
16        String methodSig = "func get_" + name + " ->" + field.getType()
17            .getFullName();
18        String method = methodSig + " = " + name + ";;";
19        return new Tuple2<StringBuffer, String>(new StringBuffer(method),
20            methodSig);
21    }
22 }

```

Listing 6: The `CyanMetaobjectGet` class in Java

body, and the signature of a field is its declaration without any initialization (`= expr`). If the compiler were smarter, the return value could be just the code to be inserted.

Class `CyanMetaobjectGet` should be compiled with the file `saci.jar` (available on the website www.cyan-lang.org). The resulting `.class` file should be put in a directory `"cyanPack\--meta\javaPack"` if class `CyanMetaobjectGet` is in the Java package `javaPack`. Now, any Cyan code that imports package `cyanPack` can use annotation `get`.

Let us study what happens in the compilation of `Painting` of Listing 5. At the start of the compilation, the compiler imports package `cyan.lang` (automatically) and `cyanPack` (line 2 of this example) and loads dynamically the classes of directories `--meta` of these packages. It then creates one object for each class. Using method `getName`, it has access to the annotation name of each metaobject, which was given in the call to `super` (line 5 of Listing 6). During parsing, when the compiler finds an annotation `get`, it looks for a metaobject called "get" in the imported packages. `CyanMetaobjectGet` is found, and the compiler creates an object of this class, which is associated with the annotation of line 5 of Listing 5.

In the compilation phase `afterResTypes` (to be seen later), the compiler calls the method `afterResTypes_codeToAdd`. If the returned value is not `null`, the returned code is inserted in the source code of the prototype, which needs to be compiled again. Only the text of the prototype in the compiler memory is changed. In the following

compilation phases, the compiler will find a `get_name` method in prototype `Painting`.

A metaobject class seems difficult to build because it demands the knowledge of a lot of cyan MOP classes, which includes the AST. However, this is only partially true because the MOP and the IDE guide the developer throughout the process. Let us see why.

A metaobject class has a standard superclass, and its constructor should call `super` with the annotation name, information on the number of annotation arguments, and a list of allowed attached declarations. The Java compiler will demand that the developer pass this information to the superclass. Based on the goals of the metaobject, the developer uses Table 1 to discover which interfaces the metaobject class should implement. The documentation of each interface tells which methods should be overridden. The coding of a metaobject class method should be based on information on the annotation or the source code (the current prototype or method). Information on the annotation (attached declaration, arguments, or attached text between `{*` and `*`}) is obtained from the method `getAnnotation`. Information on the source code is obtained from an interface method parameter called `compiler` or `env`. For example, the parameter `compiler` of method `afterResTypes_codeToAdd` (line 10 of Listing 6). As an example, suppose we want to check if there is a method in the prototype with the same name as the method the metaobject will add. This can be done with the following code, which should be put after line 14.

```

    for (WrMethodDec m : compiler.getMethodDecList()) {
        if ( m.getName().equals("`get_" + name) ) {
            return null; // no method will be added
        }
    }
}

```

“`compiler.getMethodDecList()`” returns a list of AST objects of class `WrMethodDec`, one for each method of the prototype in which the annotation is. There is no need to learn all MOP classes and methods. Using the “code completion” feature of the IDE (after typing “`paramName.`”), the developer has access to methods and, from them, to the classes she or he needs.

3.2 Interactions with the Compiler

The *metacode of a Cyan program* (code) is composed of all Java classes or Cyan prototypes of the metaobjects associated with the annotations used by the program. To support the metacode, there are two versions of a *MOP library* in the `saci.jar` file: one in Java and another in Cyan (one is a mirror of the other). The Cyan compiler knows both MOP libraries, and some of their classes and prototypes are superclasses or superprototypes of the classes and prototypes of the metacode. The class `CyanMetaobjectGet` of Listing 6 is part of the metacode of the Cyan program that has class `Painting`. This Java class imports its superclass from the MOP library that is in the `saci.jar` file, which is also part of the metacode.

When parsing source code, the Cyan compiler creates, for each annotation, an object of the AST internal to the compiler, which is represented by the round rectangle “annot. AST object” on the left side of Figure 1. This object represents the annotation as a syntactical element. Based on the annotation name, the compiler discovers the metaobject class associated with it, `CyanMetaobjectGet` in this case, and creates a metaobject, which is represented by the round rectangle on the left-bottom side of the figure. The metaobject gets the annotation data, as well as its arguments and attached declaration,

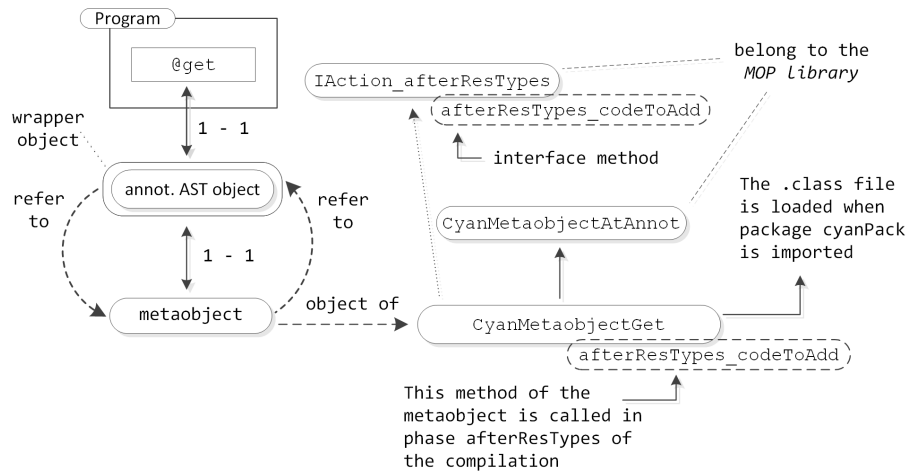


Figure 1: Relationship between metaobjects, annotations, metaobject classes, MOP interfaces, and compilation phases

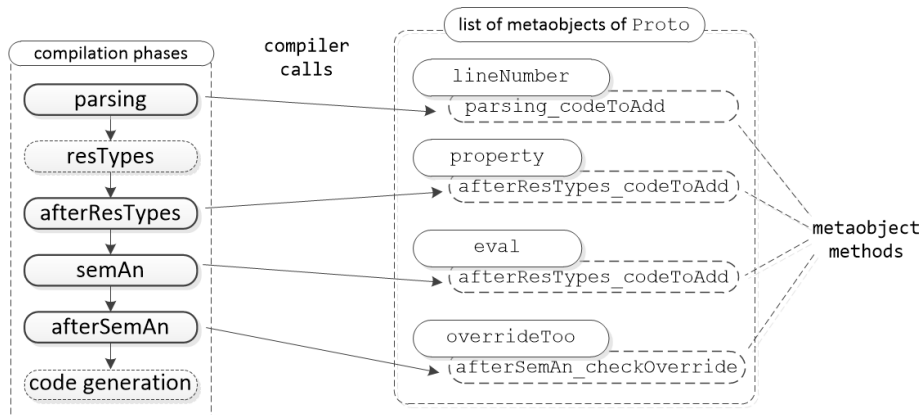


Figure 2: The compilation phases and their links to methods of metaobjects at compile time

from a wrapper AST object. In the figure, this wrapper AST object is represented by a round rectangle enveloping the rectangle “annot. AST object”. Both objects will be called *AST objects*. Figure 1 also shows that there are one-to-one relationships between metaobjects, wrapper AST objects, annotation AST objects, and annotations (the syntactic element).

3.3 Phases of Compilation

The Cyan compiler goes through six compilation phases for each source file, as shown inside the rectangle with dashed lines on the left of Figure 2. The flow of control is from top to bottom. Phase **parsing** does the syntactical analysis and builds the Abstract

Syntax Tree (AST) of the source file. Some AST objects are associated with a type and have a `type` field that is initially set to `null`. A `type` field, for example, exists in AST objects representing method parameters, prototype fields, implemented interfaces, the superprototype, message passings, and expressions.

There are two kinds of AST objects associated with types: those representing expressions and local variable declarations, which are always inside method bodies, and those outside method bodies. The `type` field of the later AST objects is set in phase **resTypes** (resolving types). Thus, the field name of `Student` in Listing 1 is represented by an AST object whose field `type` is `null` at the beginning of phase **resTypes**. During this phase, the compiler sets the `type` field to the AST object representing the prototype `String`. Phase **resTypes**, therefore, is included in the *semantic analysis* of the source code. The compiler goes through phase **resTypes** on a source file only after parsing all source files referenced in this file or loading the jar file with the referenced types.

Phases **afterResTypes** and **afterSemAn** only exist because of the MOP; they could be eliminated if Cyan did not support metaprogramming. Phase **afterResTypes** means *after resolving types*. Phase **afterSemAn** means *after semantic analysis*. In phase **semAn** (the remaining part of the semantic analysis), the compiler sets the `type` field of AST objects representing expressions and local variable declarations. In this phase, the compiler also does other checks, as demanded by the language. The last compilation phase is code generation. Currently, no metaobject method is called in this phase.

During the compilation of a prototype `P`, for each phase `X` in the set `{ parsing, afterResTypes, semAn, afterSemAn }`, the compiler runs the following algorithm:

```
for every metaobject metaObj associated with P {
  if the class/prototype of metaObj implements any interface
    associated with phase X
  then
    for each method 'met' of each interface associated
      with phase X {
        metaObj.met(...); // call method 'met'
      }
    endif
  }
}
```

The right-hand side of Figure 2 represents metaobjects associated with annotations of a hypothetical prototype `Proto` (not shown). According to the above algorithm, during phase **parsing**, the compiler calls the method `parsing_codeToAdd` of the metaobject associated with `lineNumber`. During phase **afterResTypes**, the compiler sends the message `afterResTypes_codeToAdd` to the metaobject property, and so on.

The Cyan compiler calls specific metaobject methods in each compilation phase, as exemplified in Figure 2. It also calls some metaobject methods when some events happen, such as “message passing”, “method missing”, “field access”, “field missing”, “inheritance”, and “method overriding”. That is, some interfaces are associated with *triggers*. For example, methods of interface

`IActionMethodMissing_semAn`

are called whenever the compiler is unable to find a method that matches a message passing. The methods of this interface can return an expression that replaces the message passing. Hence, this interface allows the implementation of *virtual* methods. The annotation associated with the metaobject should be attached to a prototype or a method. In light of what was presented in this Subsection, the reader is invited to study Table 1 (end of the article). It shows the goals and the interfaces that the metaobject class should

implement to achieve them. For example, to intercept “message passing”, the metaobject class should implement the method

```
IActionMessageSend_semAn
```

Note that the interface name ends with the compilation phase.

3.4 Important MOP Features

The following list complements the information on the MOP given in the previous Subsection.

- (a) Metaobjects always generate code as strings. The code is added to a copy in memory of the prototype source code; the original file is not changed.
- (b) The class or prototype of a metaobject may implement any number of interfaces for any number of phases. As a result, a metaobject can act in multiple compilation phases.
- (c) In each phase, all metaobjects associated with a prototype have the same view of the code, which is how the code was in the last compilation phase. In phase **afterResTypes**, however, each of a prototype’s metaobjects knows what fields and methods the other metaobjects generated. This information is passed through the parameter `infoList` (see line 11 of Listing 6). The methods are called in rounds. In the first round, `infoList` is `null`. In the second round, `infoList` contains information on the code generated in the previous round. And so on until a maximum of 5 rounds. Each metaobject can then change the code it generates based on the code generated in the previous round. In phase **semAn**, no metaobject is aware of the code added by other metaobjects in that phase. A consequence of the above is that the order in which metaobject methods are called is not important. It would be relevant if a metaobject could view the code just added by other metaobjects in the same phase.
- (d) Fields and methods can only be added in phase **afterResTypes**, and statements and expressions can only be added in phase **semAn**. Hence, metacode developers can be sure that, in phase **semAn**, the fields and methods of prototypes will not change.
- (e) Base code can be changed by metaobjects in only one way: method statements and expressions can be replaced by others. Metaobjects cannot change inheritance (add or remove), method parameter types, implemented interfaces, and so on.
- (f) A metaobject may add code in several phases. The added code may have annotations. However, the associated metaobjects will only be active in the next phase. Accordingly, in compilation phase X, a metaobject may generate code with annotations. But the metaobjects associated with these annotations will only be used in the compilation phases following X.

For example, suppose an annotation `@addField(0)` adds the following code to a prototype in phase **afterResTypes**.

```
@addField(1)
var Int field0;
```

The annotation `@addField(1)` should also add field `field1` to the prototype. However, this does not happen: the metaobject associated with this annotation will only be active in the next compilation phase, **semAn**, and in this phase, metaobjects cannot add fields to prototypes.

- (g) The compiler keeps track of which code was inserted by which metaobject. If there is a compilation error, it points out exactly the annotation that inserted the code with

errors. That is only possible because: (i) metaobject methods return code as strings; and (ii) metaobjects do not change the AST directly, it is the compiler that inserts the code.

(h) The Cyan compiler is the active part of the metaprogramming system. It calls the metaobject methods, which may return code to be inserted. However, it is the compiler that inserts the code.

(i) Metaobject methods only have access to a simplified and read-only version of the compiler AST. For example, the class `WrFieldDec` used in line 13 of Listing 6 is the read-only version of the AST class `FieldDec` of the compiler (that represents a field of a prototype). Internal compiler details of the latter are not revealed in the former.

(j) A metaobject method can add fields and methods to the prototype `P` in which the metaobject annotation is. In this case, the *natural prototype context* of the metaobject is `P`. A metaobject method whose annotation is in prototype `P` can be called when `P` is inherited by `Q`. In this case, the metaobject method will be called in the context of `Q`, which will be the *natural prototype context* of the metaobject (*natural context*, for short). The AST of `Q` will be visible to the metaobject method.

A prototype cannot view the fields (all are private) and private methods of other prototypes.⁶ A metaobject method cannot view the private fields and methods of a prototype different than its *natural context* prototype. The visibility of a metaobject method is the same as its *natural context*.

We will explain that using the example of Subsection 3.1. Suppose the expression `((WrPrototype) field.getType()).getFieldList(compiler.getEnv())` is inside the `afterResTypes_codeToAdd` method of the `CyanMetaobjectGet` class of Listing 6. This expression tries to get private information about the type of the field, which is the list of fields of the type. There would be a compilation error in the example of Listing 5 because a metaobject of `Painting` would be trying to access private information of `String`, the type of the field name. However, there would be no error if the type of name were `Painting` because a prototype has the right to know its own private parts.

(k) When the compiler calls a metaobject method, it passes as an argument an object describing the compiler itself. The method parameter has names such as `compiler` (line 10 of Listing 6) or `env`. The available information varies according to the interface phase in which the metaobject method was declared. During **parsing**, metaobjects do not know method statements, other prototypes, or any information on code that comes textually after the annotation. In phase **resTypes**, metaobjects have access to AST objects that describe everything outside method statements. In phase **semAn**, metaobjects have access to all information except the types of variables and expressions that come after the annotation. Finally, all the information is available for metaobjects in phase **afterSemAn** (the code and the AST cannot be changed in this phase).

(l) Metaobjects cannot add code in phase **afterSemAn**. Therefore, the metaobjects that act in this phase view the final code, which cannot be changed any longer.

(m) Listing 2 shows examples of annotations starting with `@`, the most common kind. There are other annotations with different syntax or having other roles: annotations to types, macros, literal numbers and strings, and *Codegs*. Annotations can be attached to types for *additional* type checking. Hence, Cyan supports a *limited* pluggable-type system [Bracha 2004]. For example, annotation `range` limits the values that can be assigned to a variable of that annotated type. Hence, if a variable has type `Int@range(1, 12)`, only integers in this range can be assigned to it. This is checked at compile-time (if

⁶ For a complete discussion of the visibility rules of Cyan, refer to the Cyan manual.

possible) or at runtime. Macros work similarly to other languages although they are full metaobjects. A literal number ending with an identifier such as `110bin` or `0FFF_Hex` is an annotation. Literal strings starting with an identifier are annotations, as `r"x[A-Z]*"` and `xml"" code in XML""`. The metaobject associated to a literal number or string replace the literal with any expression and it has the full power of a metaobject. *Codegs* are metaobjects that depend on an IDE plugin to work (a single plugin for all metaobjects). The plugin allows the metaobject to get data at editing time through a graphical interface made specially for that metaobject (it is part of it). This data can be used by the metaobject at compile time for code checking and generation.

4 Comparison with Related Work

This section presents some metaprogramming systems and how they are related to Cyan. The first subsection describes mechanisms for code generation and the benefits and drawbacks of each of them. Subsection 4.2 compares Cyan with runtime metaprogramming. The comparison related to embedded external DSLs are made in Subsection 4.3. The last Subsection compares Cyan with languages and systems supporting compile-time metaprogramming.

4.1 How Code is Generated and Represented

Metaprograms generate code in many representations using several mechanisms [Smaragdakis 2015], described next.

(a) *As text*. Code is generated in string format. There is no guarantee that the generated code is error-free.

(b) *Handling of the program Abstract Syntax Tree*. Code is generated by creating objects of the AST representing it, considering that the compiler is implemented in an object-oriented language.⁷ Therefore, the developer has to know a great number of classes (it would be more than one hundred in Cyan). Code generation is difficult because it demands the mapping, by the metaprogrammer, of human-legible source code into the creation of AST objects. AST handling has the advantage that the metaprogram compiler usually catches all syntactic errors in the generated code. However, the AST objects created by metaprogramming may have semantic errors. If this happens, the compiler will not be able to identify the metacode that produced the AST objects that caused the errors.

(c) *Quoting* A special language syntax is used to transform text into AST objects. Therefore, the metaprogram handles text that is converted into AST objects. As a short example, in the Python-based language Converge [Tratt 2008], the AST of `1 + 2` is obtained by the *quasi-quoted* expression `[| 1 + 2 |]`.

(d) *Specialized languages* Domain-Specific Languages are used to generate code. There are many languages that fit in this category: Genoupe [Draheim et al. 2005A] [Draheim et al. 2005B], SafeGen [Huang et al. 2005], and PTFJ [Miao and Siek 2012], CTR [Fähndrich 2006], MorphJ [Huang and Smaragdakis 2011], MTJ [Reppy and Turon 2007], and PTFJ [Miao and Siek 2012]. They share the common characteristic that the modifications they can make are limited. Either the code generation is inflexible (based on a pattern) or limited to certain tasks (as adding fields and methods to a class and only

⁷ The reasoning does not change if the AST is implemented in a language that is not object-oriented.

this). These languages are not further discussed because they are not comparable with the Cyan MOP, which is more general.

Cyan generates code as strings, and any errors in them are not caught either at compile-time or runtime of the metaprogram (metaobject code). However, errors are discovered in the compilation of the base program. The compiler will give precise error messages because it associates every code added to the source file with an annotation, which is the annotation associated with the metaobject that generated the code. The compiler will indicate precisely which annotation is linked to the code with errors.

4.2 Runtime Metaprogramming

The Smalltalk MOP [Goldberg and Robson 1983] [Nierstrasz et al. 2009] is fundamentally different from that of Cyan because it cannot change the program. However, a Smalltalk program can change itself at runtime using methods inherited from fundamental classes such as `Behavior`, which are outside the MOP. There are methods that can, for example, add a new method to a class.

In Python 3 [Ramalho 2015], metaclasses are used to change classes, including adding code to them. Each class has a single metaclass, a limitation that drastically reduces the complexity of metaprogramming in Python because there would be no interactions between metacode. Many of the problems with CTMP cited in Section 1 do not exist in Python. But the limit of just one metaclass per class severely damages the usability of metaprogramming. A metaclass cannot intercept class inheritance or method overrides in a subclass. There are no compile-time guarantees in relation to metaclasses because classes are created only at runtime. Metaobjects in Cyan have access to the AST of the current prototype. In Python, the AST is not readily available. In Cyan, there are four compiler phases in which metaobjects may act. In Python, metaclasses act only when the class is created.

The prime example of a Metaobject Protocol is that of CLOS [Kiczales et al. 1991] [Kiczales et al. 1993] [Paepcke 1993] [Bobrow 1993] [DeMichiel and Gabriel 1987], an extension of Common Lisp [Steele 1990] with features for object-oriented programming. The CLOS MOP acts at runtime, allowing the intercepting of several operations: object creation, memory allocation, the calculus of superclass precedence,⁸ method calls, field access, and many more. The MOP of this language uses *metaclasses*, which are classes of classes and methods.⁹ Metaclasses are objects too. By using a user-made metaclass for a class, we change its expected behavior. For example, a metaclass can introduce a field into a class that keeps track of how many objects were created. The method that creates instances of the class may increment this field every time it is called. The Cyan MOP has a great number of the features of the CLOS MOP, such as the intercepting of method calls and the addition of code to classes/prototypes. However, there are many differences:

- (a) metaprogramming occurs at runtime in CLOS and at compile time in Cyan. In Cyan, errors are detected earlier, and DSL code attached to annotations can be checked at compile-time;
- (b) in CLOS, a class or other element has a single metaclass. In Cyan, any number of metaobjects can be attached to a declaration, and each annotation may have parameters and an attached DSL code, allowing for easy metaobject configuration;

⁸ The superclasses have to be ordered because the language supports multiple inheritance.

⁹ CLOS has both *methods* and *generic methods*. For our purposes, it is not necessary to distinguish between them.

- (c) the AST is not available for the metaclass. This limits what metaclasses can do;
- (d) Cyan has a lot of security measures that CLOS does not support. In Cyan, a metaobject has limited view privileges, which are usually the same as the view of the prototype associated with it.

4.3 DSL Code Embedded Within Base Code

Cyan supports external DSLs embedded in its code. The DSL code is put either between `{*` and `*}` attached to the annotation (see Listing 3) or inside a string preceded by an identifier (as in `xml" . . ."`). The metaobject associated with the annotation is responsible for parsing and analyzing the DSL code at compile time. Converge [Tratt 2008], a Python-based language, supports a similar mechanism:

```
$<<compilerFunc>>
  dslCode
```

Function `compilerFunc` is called at compile time with `dslCode` as an argument (a string) and returns an AST object that replaces the *DSL block* (this whole example). Unlike metaobjects, this function cannot add fields, methods, or classes to the program. This Converge language mechanism is the closest we have found in the literature to Cyan annotations with attached DSL code.

4.4 Compile-Time Metaprogramming

This subsection is divided into parts. First, the text presents some compile-time metaprogramming systems. They are then compared with Cyan.

Metaprogramming Systems with a MOP

OpenC++ [Chiba 1995] is a C++ extension in which metaclasses for classes and methods are given the opportunity of changing the AST after parsing. A metaclass for a class `C` may intercept method calls whose receivers have type `C`. The method call may, after the interception, be changed or replaced. The MOP of OpenC++ also allows interception of variable declarations, creation of objects, and reading and writing in fields. OJ [Tatsubori et al. 2000] [Tatsubori 1999] is a Java extension in which a class may be associated with a user-defined metaclass. Methods of the metaclass have the opportunity of changing the AST. For example, a method called `translateDefinition` of a metaclass may add methods to the class. `expandFieldRead` can change the read of a class field. The user-defined metaclass can also define methods for intercepting object creation, array allocation, writing to fields, method calls, and casts to the class.

Metaprogramming Systems without a MOP

Languages Xtend¹⁰, Groovy [König 2007], and Nemerle¹¹ [Skalski 2005] support *compile-time metaprogramming* without a Metaobject Protocol. These languages support *metaprogramming features*. They share many similar characteristics, as described below, and therefore will be considered together.

- (a) Annotations are attached to classes, methods, and other declarations;

¹⁰ <https://www.eclipse.org/xtend>

¹¹ <http://nemerle.org>

- (b) An annotation is linked to a *Processor Class* (PC) that can implement interfaces and define methods that change the compilation;
- (c) Methods of the PC are invoked in several phases of the compilation, like before parsing, after parsing, before typing members (similar to **afterResTypes** of the Cyan compiler), after semantic analysis, during code generation, etc.;
- (d) Methods of the *Processor Class* have parameters that represent language elements that can be changed at compile time. For example, the AST object of the annotated class or method is passed as an argument. Methods of the PC can, using these AST objects, add methods to an annotated class, change inheritance, add statements to an annotated method, change method statements, and so on. Any AST object reachable from the method arguments can be changed. Therefore, a method can be added to a class that is not annotated or is not directly related to the annotated class. The class may be, for example, just the type of a method parameter accessible to the PC method;
- (e) A method of the *Processor Class* that overrides an interface method is used in the compilation phase associated with that interface (much like Cyan).

Compiler Plugins

A *compiler plugin* is composed of metacode that interacts at low-level with the compiler, changing the compilation process. The difference between the terms *compiler plugins* and *metaprogramming systems without a MOP* is that the former emphasizes implementation aspects (low-level details), whereas the latter emphasizes conceptual aspects (high-level specifications). *Compiler plugins* are supported by the languages Scala¹², Java¹³, X10 [Nystrom and Saraswat 2007], Kotlin,¹⁴ TypeScript¹⁵, and Rust¹⁶. Java annotation processors [Darcy 2006] are compiler plugins for Java that allow checks but not code modifications. They are used, for example, for implementing pluggable types [Bracha 2004]. Project Lombok [Kimberlin 2010] is a Java annotation processor whose supported annotations can add code to classes because it uses non-supported downcasts. Compiler plugins will not be discussed in depth in this paper because there is a shortage of good documentation about them. However, languages whose compilers accept plugins have all the main characteristics of languages supporting *compile-time metaprogramming* without a Metaobject Protocol, as discussed above.

Comparison with the Cyan MOP

The Cyan MOP was built on the metaprogramming system of modern languages, such as Groovy, but with some characteristics of MOPs. From Metaobject Protocols, Cyan took the interception of message passings, inheritance, field access, and errors such as method or field missing. From the other systems, it took all the rest.

The following comparison will be guided by the list of problems with metaprogramming given in the introduction. The problem letter precedes each discussion.

¹² <https://docs.scala-lang.org/scala3/reference/changed-features/compiler-plugins.html>

¹³ <https://docs.oracle.com/en/java/javase/11/docs/api/jdk.compiler/com/sun/source/util/Plugin.html>

¹⁴ <https://kotlinlang.org/docs/all-open-plugin.html>

¹⁵ <https://github.com/microsoft/TypeScript/wiki/Using-the-Compiler-API>

¹⁶ Compiler plugins are an unstable feature of the language. See <https://doc.rust-lang.org/beta/unstable-book/language-features/plugin.html>

(a) *A metacode associated with a source file can change the base code of another file. Hence, to know the final version of a file, the developer has to know the behavior of every metacode in the metaprogram.*

Languages OJ, Xtend, Groovy, and Nemerle allow non-local changes through AST handling. Hence, a metacode of one class can change another class. In particular, metacode associated with annotations of a source file can change another source file, which is called *obliviousness* [Filman and Friedman 2000]. Generally, any metaprogramming system that supplies the AST to metacode has this problem [Clifton and Leavens 2003]. CLOS, OpenC++, and BJS limit the changes to the scope of the metaclass or metacode. In Cyan, metaobjects have access to a read-only AST and it is the compiler that changes the base code. The compiler only allows changes in the current prototype or the prototype that the metaobject was intended to change, which is its *natural context prototype*, as explained in item j of Subsection 3.4.

(b) *A class may not have access to the private fields and methods of another class. But the metacode can because it has access to the compiler data structures.*

As explained in item j of Subsection 3.4, there is a compilation error when a metaobject attempts to access private data of a prototype that is not its *natural context prototype*.

(c) *The developer has to know many details about the compiler since the metacode interacts with low-level compiler algorithms and data structures. Therefore, metaprogramming is not only harder, but the metacode can bypass compiler checks, invalidate invariants, damage data structures, and crash the compiler.*

In Cyan, metaobjects have access to a read-only and *simplified* AST of the base code. In most metaprogramming systems, the metacode has direct access to the compiler's AST. Thus, metacode can easily damage the AST. If the compiler discovers the damage, it will not be able to point out the guilty metacode. The metacode may break compiler invariants after the compiler does the final checks, leading it to generate incorrect code or crash. That is impossible in Cyan. Besides that, Cyan metaobjects generate code as strings. Hence, metaobject developers do not need to have deep knowledge of the AST or the compiler.

(d) *Metacode can insert new code into the base code. If the new code has errors, the compiler is unable to point out which metacode (or associated annotation) is responsible for adding the new code. This is because the metacode itself changes the compiler AST to add new AST objects. No trace is left of who did what.*

This problem occurs in all metaprogramming systems that allow metacode to handle the AST. In Cyan, metaobject methods return code as strings that are inserted in the base code by the compiler. The compiler associates the inserted code with the annotation that asked for its insertion. If there is an error in the code produced by the metaobject, the compiler will point out exactly which annotation is responsible for it.

The Converge programming language [Tratt 2005] [Tratt 2008] supports compile-time metaprogramming. The compiler tracks who produced which code to issue precise error messages. It goes beyond Cyan in two aspects: (a) every bytecode¹⁷ knows its origin, which can be used in runtime error messages, and (b) an AST node can be associated with more than one location (an error may be associated with more than one source).

(e) *The order in which the metacode is called is not specified by the language. It can even change between compilations (with the same base code). By changing the metacode calling order, their view of the program can change if a metacode can view the code added by a metacode executed previously.*

¹⁷ The source code is translated into bytecodes of a Converge VM.

To our knowledge, the only metaprogramming languages that guarantee the execution order of metacode are BSJ [Palmer and Smith 2011] and Cyan (as explained in item c of Subsection 3.4). The order is important if there are two or more metacode associated with the same class, prototype, or source file. What one metacode does may impact the following metacode the compiler calls.

(f) *A verification performed by one metacode may become obsolete due to modifications introduced by another metacode.*

In Cyan, metaobjects should do checks in phase `afterSemAn`, in which the code cannot be changed anymore. In languages that allow unrestricted changes in the AST, checks made by a metacode may be invalidated by code added later by another metacode.

(g) *Metacode may generate metacode that may generate metacode ad infinitum.*

This problem usually does not occur in metaprogramming languages because the base code is processed just once for each metacode. In Cyan, metaobjects may generate code with annotations, but these will only be used in the next compilation phase. As the number of phases is finite, so will be the compilation.

(h) *Metacode may not terminate its computation. The compiler will not terminate either.*

In Cyan, each metaobject method has a time limit for execution, a feature that can be turned on and off by a compiler option. Since enforcing this limit requires the compiler to spawn a dedicated thread, the operation incurs significant computational overhead and should therefore be enabled only during the testing phase. Developers can specify distinct time limits for the program as a whole and for individual packages. To the best of our knowledge, no other metaprogramming language with or without a MOP employs this solution.

(i) *Metacode may generate code that uses identifiers already in use. This is the hygiene problem of macros.*

This problem does occur in Cyan and, to our knowledge, in all metaprogramming systems with or without a MOP. In Cyan, it can be easily avoided. There is a MOP method that returns an identifier name that is guaranteed not to crash with any other identifier in the program.

4.5 Comparison of Features

The Cyan MOP has some features that are not found in other metaprogramming systems without a MOP, generally. These are the systems directly related to Cyan, both in terms of implementation and functionality. These features are described next. It is not difficult to add them to any language.

In Cyan, a source file that imports a package has access to all metaobjects whose “.class” bytecode files were put in directory `--meta` of the package. In other languages, the path (directory and file name) of the metacode should be specified through some configuration file or through the command line.

In Cyan, annotations can be expressions such as `eval` presented in Subsection 2.2 or `@compilationInfo("currentmethodname")` that return the current method name as a string. We are unaware of any other metaprogramming language that allows this.

Cyan has a project file that specifies which packages compose the program. Metaobjects can be used in the project file to specify compiler options, create and associate variables with values (that can be used by metaobjects), and apply an annotation to a bunch of prototypes. For example, instead of attaching metaobject `checkStyle` to all

prototypes of the package `cyanPack`, we can just attach the annotation to the package in the project file.

Metaobjects can transfer information from a prototype `P` to a generic prototype that takes `P` as an argument. As an example, a metaobject associated with prototype `Int` supplies the code of a method `sum` that is added to `Array<Int>` by a metaobject associated with this last prototype.

In phases `afterResTypes`, `semAn`, and `afterSemAn`, metaobjects of the same prototype can communicate with each other through the exchange of data. This may be used for coordinating their actions.

Some metaobjects take interpreted Cyan (Myan) code as the attached DSL text. Using this feature, metacode can be specified inside regular Cyan code. As an example, see Listing 4 of Subsection 2.2. The attached text of the annotation is metacode in Myan, which can be put in standard directories of a package and loaded at compile time.

When there is a compilation error in a source file, the compiler produces a file with the final source code, which includes the code added by metaobjects. The developer can view the code generated by metaobjects and discover possible errors in them.

To the best of our knowledge, no comprehensive studies have investigated the impact of metaprogramming on compilation time in other programming languages. In Cyan, executing metaobject methods incurs computational overhead in addition to the inherent cost of running the methods themselves. Enabling the time limit for these methods requires the compiler to create a dedicated thread, a computationally expensive operation. Conversely, when this feature is disabled, the additional computational cost during most compilation phases — except for semantic analysis — is limited to a type check and a method invocation. This results in only a negligible increase in compilation time, which is barely noticeable. During semantic analysis, however, the compiler's performance is further affected by pre- and post-processing algorithms executed after each metaobject method call. This additional overhead is approximately half the time required to spawn a new thread.

5 Conclusion

No sufficiently powerful metaprogramming system is easy to use, which includes the Cyan MOP. However, we consider that this MOP is not too difficult to use in relation to its power because of the following reasons. The goals direct the implementation of metaobject classes (which interfaces to implement; see Table 1 at the end of this article and Subsection 3.1). Therefore, the metaprogrammer, guided by the goals, makes the most important decisions *before* starting to code. In each compilation phase, metaobject methods *ask* the compiler to add code. As a result, the metaprogram acts passively in relation to the compiler, who is in control of the execution flow of the metaprogram. The developer has access to simplified read-only compiler data structures (which include the AST of the base code), the developer does not need to know the compiler internal details, there cannot be a compiler crash caused by metacode. When there is a compilation error in a source file, the compiler points out exactly the annotation that caused the error and produces a new file with the code added by metaobjects. There are many security mechanisms that are not present in other systems.

Metaobjects have been extensively used in the Cyan libraries. There are many general metaobjects, such as `property`, `init`, `immutable`, `eval`, and `doc`, but also those for specific prototypes and methods. Among the latter, there are metaobjects for generating code, such as `createTuple` and `createFunction` (they create all methods of

prototypes `Tuple<T>` and `Function<T+>`). And there are ones for doing checks based on the semantics of methods, such as `checkIsA` (check if an argument is a prototype), `checkMethodEqualEqual` (check if the result of “`e1 == e2`” is known at compile time and, therefore, the comparison is unnecessary), and `overrideToo` (demand that, if `==` is overridden in a subprototype, method `hashCode` is too). The use of metaobjects in the Cyan libraries is a good justification of the main goals of the Cyan MOP, which are to build general metacode and to make libraries easier to build (because of code generation) and safer to use (because of compile-time checks).

We have plans to further improve the Cyan MOP. One of the planned features is to support variable ownership, like in language Rust [Klabnik and Nichols 2024]. The Cyan compiler is available for download at `cyan-lang.org`. There one can find the language manual, a complete description of the Cyan MOP, extended versions of this article (with all the source code), and a list of around one hundred metaobjects with examples.

Acknowledgements

This project was partially financed by FAPESP (São Paulo - Brazil) under Process number 2014/01817-3.

References

- [Biboudis et al. 2016] Biboudis, A., Inostroza, P., Storm, T.: “Recap: Java Dialects as Libraries”; Proceedings Of The 2016 ACM SIGPLAN International Conference On Generative Programming: Concepts And Experiences (2016), 2-13.
- [Bobrow 1993] Bobrow, D., Gabriel, R., White, J.: “CLOS in Context: The Shape of the Design Space”; Chapter of the book *Object-oriented Programming: The CLOS Perspective*, Cambridge, MA, USA (1993), 29-61.
- [Bracha 2004] Bracha, G.: “Pluggable type systems”; In *OOPSLA’04 Workshop On Revival Of Dynamic Languages* (2004).
- [Burmako 2013] Burmako, E.: “Scala Macros: Let Our Powers Combine! On How Rich Syntax and Static Types Work with Metaprogramming”; *Proc. Of the 4th Workshop On Scala* (2013).
- [Chiba 1995] Chiba, S.: “A Metaobject Protocol for C++”; *Proceedings Of The Tenth Annual Conference On Object-oriented Programming Systems, Languages, and Applications* (1995), 285-299.
- [Clifton and Leavens 2003] Clifton, C., Leavens, G.: “Obliviousness, Modular Reasoning, and the Behavioral Subtyping Analogy”; *Computer Science Technical Reports — Iowa State University* (2003);
- [Darcy 2006] Darcy, J.: “Java Specification Request 269: Pluggable annotation processing API” (2006) <http://jcp.org/en/jsr/detail?id=269>
- [DeMichiel and Gabriel 1987] DeMichiel, L., Gabriel, R.: “The Common Lisp Object System: An Overview”; *European Conference On Object-oriented Programming ECOOP ’87* (1987), 151-170.
- [Draheim et al. 2005A] Draheim, D., Lutteroth, C., Weber, G.: “A Type System for Reflective Program Generators”; *Proceedings Of The 4th International Conference On Generative Programming And Component Engineering* (2005), 327-341.
- [Draheim et al. 2005B] Draheim, D., Lutteroth, C., Weber, G.: “Generative programming for C#”; *ACM SIGPLAN Notices*, 40, 8 (2005), 29-33.

- [Fähndrich 2006] Fähndrich, M., Carbin, M., Larus, J.: “Reflective Program Generation with Patterns”; Proceedings Of The 5th International Conference On Generative Programming And Component Engineering” (2006), 275-284.
- [Filman and Friedman 2000] Filman, R., Friedman, D.: “Aspect-Oriented Programming is Quantification and Obliviousness”; OOPSLA 2000 Workshop On Advanced Separation Of Concerns (2000).
- [Flanagan and Matsumoto 2008] Flanagan, D., Matsumoto, Y.: “The Ruby Programming Language”; O’Reilly (2008).
- [Fowler 2010] Fowler, M.: “Domain Specific Languages”; Addison-Wesley Professional (2010).
- [Goldberg and Robson 1983] Goldberg, A., Robson, D.: “Smalltalk-80: the language and its implementation”; Addison-Wesley Longman Publishing Co. (1983).
- [Guimarães 2020] Guimarães, J.: “The Cyan Language” (2024) <http://cyan-lang.org/docs>
- [Guimarães 2024] Guimarães, J.: “The Cyan Language Metaobject Protocol” (2024) <http://cyan-lang.org/docs>
- [Huang et al. 2005] Huang, S., Zook, D., Smaragdakis, Y.: “Statically Safe Program Generation with Safegen”; Proceedings Of The 4th International Conference On Generative Programming And Component Engineering (2005), 309-326.
- [Huang and Smaragdakis 2011] Huang, S., Smaragdakis, Y.: “Morphing: Structurally Shaping a Class by Reflecting on Others”; ACM Trans. Program. Lang. Syst., 33, 2 (2011).
- [Kamin et al.] Kamin, S., Clausen, L., Jarvis, A.: “Jumbo: Run-Time Code Generation for Java and Its Applications”; Proceedings Of The International Symposium On Code Generation And Optimization: Feedback-Directed And Runtime Optimization (2003), 48-56.
- [Kiczales et al. 1991] Kiczales, G., Rivières, J., Bobrow, D.: “The Art of Metaobject Protocol”; MIT Press (1991).
- [Kiczales et al. 1993] Kiczales, G., Ashley, J., Rodriguez, L., Vahdat, A., Bobrow, D.: “Metaobject protocols: Why we want them and what else they can do”; Chapter of the book Object-oriented Programming: The CLOS Perspective, Cambridge, MA, USA (1993), 101-118.
- [Kiczales et al. 2001] Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., Griswold, W.: “An Overview of AspectJ”; ECOOP, (2001), 327-353.
- [Kimberlin 2010] Kimberlin, M.: “Reducing boilerplate code with project Lombok” (2010) <http://jnb.ocicweb.com/jnb/jnbJan2010.html>
- [Klabnik and Nichols 2024] Klabnik, S., Nichols, C.: “The Rust Programming Language”; No Starch Press (2024) <https://doc.rust-lang.org/book>
- [Kohlbecker et al. 1986] Kohlbecker, E., Friedman, D., Felleisen, M., Duba, B.: “Hygienic Macro Expansion”; Proceedings Of The 1986 ACM Conference On LISP And Functional Programming (1986), 151-161.
- [Liskov et al. 1977] Liskov, B., Snyder, A., Atkinson, R., Schaffert, C.: “Abstraction Mechanisms in CLU”; Communications Of The ACM, 20, 8 (1977) 564-576.
- [König 2007] König, D.: “Groovy in Action”; Manning (2007).
- [Lilis, Y. and Savidis 2019] Lilis, Y., Savidis, A.: “A Survey of Metaprogramming Languages”; ACM Comput. Surv., 52, 6 (2019).
- [Miao and Siek 2012] Miao, W., Siek, J.: “Pattern-Based Traits”; Proceedings Of The 27th Annual ACM Symposium On Applied Computing” (2012), 1729-1736.
- [Nierstrasz et al. 2009] Nierstrasz, O., Ducasse, S., Pollet, D.: “Squeak by Example”; Square Bracket Associates (2009).

- [Nystrom and Saraswat 2007] Nystrom, N., Saraswat, V.: “An annotation and compiler plugin system for X10”; Technical Report RC24198, IBM TJ Watson Research Center (2007).
- [Paepcke 1993] Paepcke, A.: “User-level Language Crafting: Introducing the CLOS Metaobject Protocol”; Chapter of the book *Object-oriented Programming: The CLOS Perspective*, Cambridge, MA, USA (1993), 65-99.
- [Palmer and Smith 2011] Palmer, Z., Smith, S.: “Backstage Java: Making a Difference in Metaprogramming”; *SIGPLAN Not.*, 46, 10 (2011), 939-958.
- [Papi 2008] Papi, M., Ali, M., Correa, T., Perkins, J., Ernst, M.: “Practical Pluggable Types for Java”; *Proceedings Of The 2008 International Symposium On Software Testing And Analysis (2008)*, 201-212.
- [Parr 2013] Parr, T.: “The Definitive ANTLR 4 Reference”; *Pragmatic Bookshelf* (2013).
- [Ramalho 2015] Ramalho, L. “*Fluent Python: Clear, Concise, and Effective Programming*”; *O’Reilly Media* (2015).
- [Redmond and Cahill 2002] Redmond, B., Cahill, V.: “Supporting Unanticipated Dynamic Adaptation of Application Behaviour”; *Proceedings of the 16th European Conference On Object-Oriented Programming (2002)*, 205-230.
- [Reppy and Turon 2007] Reppy, J., Turon, A. “Metaprogramming with Traits”; *Proceedings Of The 21st European Conference On Object-Oriented Programming (2007)*, 373-398.
- [Rompf 2012] Rompf, T., Amin, N., Moors, A., Haller, P., Odersky, M.: “Scala-Virtualized: Linguistic Reuse for Deep Embeddings”; *Higher Order Symbol. Comput.*, 25, 1 (2012), 165-207.
- [Skalski 2005] Skalski, K.: “Syntax-extending and type-reflecting macros in an object-oriented language”; *Master Thesis, University of Wroclaw* (2005).
- [Smaragdakis 2015] Smaragdakis, Y., Biboudis, A., Fourtounis, G.: “Structured Program Generation Techniques”; *Grand Timely Topics In Software Engineering - International Summer School GTTSE (2015)*, 154-178.
- [Steele 1990] Steele, G.: “*Common LISP: The Language*”; 2nd Ed., *Digital Press* (1990).
- [Stroustrup 2003] Stroustrup, B.: “Concept Checking - A More Abstract Complement to Type Checking”; *C++ Standards Committee Papers. ISO/IEC JTC1/SC22/WG21 (2003)* <http://www.stroustrup.com/n1510-concept-checking.pdf>
- [Stroustrup 2013] Stroustrup, B. “*The C++ Programming Language*”; *Addison-Wesley Professional* (2013).
- [Subramaniam 2013] Subramaniam, V.: “*Programming Groovy 2: Dynamic Productivity for the Java Developer*”; *Pragmatic Bookshelf* (2013).
- [Taha 2007] Taha, W.: “A Gentle Introduction to Multi-stage Programming, Part II”; In *Generative And Transformational Techniques In Software Engineering II, International Summer School, GTTSE 2007 (2007)*.
- [Tatsubori 1999] Tatsubori, M.: “An Extension Mechanism for the Java Language”; *University of Tsukuba* (1999).
- [Tatsubori et al. 2000] Tatsubori, M., Chiba, S., Itano, K., Killijian, M.: “OpenJava: A Class-Based Macro System for Java”; *Proceedings Of The 1st OOPSLA Workshop On Reflection And Software Engineering: Reflection And Software Engineering, Papers From OORaSE 1999 (2000)*, 117-133.
- [Tratt 2005] Tratt, L.: “The Converge programming language”; *Technical Report TR-05-01, Department of Computer Science, King’s College London* (2005).
- [Tratt 2008] Tratt, L.: “Domain Specific Language Implementation via Compile-time Metaprogramming”; *ACM Trans. Program. Lang. Syst.*, 30, 10 (2008), 1-40.

Goal	Metaobject interface
Create a new prototype in the package of P	IActionNewPrototypes_parsing IAction..._afterResTypes IActionNewPrototypes_semAn
Associate information with a declaration that can be later retrieved by other metaobjects	ICompilerInfo_Parsing
Add new fields and methods to P, add code before the first statement of a method of P, or check the signatures of fields and methods	IAction_afterResTypes
Add code after the annotation myAnnot, know the fields and methods of P, know the AST of statements before the annotation, or replace a statement of P by another statement	IAction_semAn
Add code after the annotation, which should be attached to a local variable declaration	IActionVariableDeclaration_semAn
Replace a message passing by any statement. Attach myAnnot to a method	IActionMessageSend_semAn
Replace a message passing for which there is no associated method (it would cause a compilation error).	IActionMethodMissing_semAn
Replace assignments to a field or the <i>get</i> of a field. myAnnot should be attached to a field	IActionFieldAccess_semAn
Replace the <i>get</i> of a non-existing field by an expression; replace the assignment to a non-existing field by an expression. myAnnot should be attached to P	IActionFieldMissing_semAn
Implement pluggable-type systems	IActionAttachedType_semAn
Check the subprototypes of P. The metaobject methods are executed in the context of the subprototype	ICheckSubprototype_afterSemAn
Check overridden subprototype methods. If the annotation myAnnot is attached to method m of P, the metaobject can check methods of subprototypes that override m.	ICheckOverride_afterSemAn
Check a declaration, which may be a prototype, method, or field. The metaobject has access to the whole AST of the declaration	ICheckDeclaration_afterSemAn
Check every message passing that may call the method attached to myAnnot	ICheckMessageSend_afterSemAn

Table 1: Table “goal” to “which interface to implement” for a metaobject myAnnot whose annotation is attached to prototype P or is inside P