



Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico


A new modular implementation for stateful traits

Pablo Tesone^{a,b,*}, Stéphane Ducasse^a, Guillermo Polito^c, Luc Fabresse^b,
Noury Bouraqadi^b

^a Inria Lille-Nord Europe, France

^b IMT Lille Douai, Univ. Lille, Unité de Recherche Informatique Automatique, F-59000 Lille, France

^c Univ. Lille, CNRS, Centrale Lille, Inria, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille, F-59000 Lille, France



ARTICLE INFO

Article history:

Received 30 April 2019

Received in revised form 8 April 2020

Accepted 14 April 2020

Available online 29 April 2020

Keywords:

Traits

Dynamic languages

Meta-object protocol

Language extension

Modular languages

Metaclass specialization

ABSTRACT

The term traits is overloaded in the literature. In this work we refer to traits as the stateless model and implementation described in Schaerli et al. articles.

Traits provide a flexible way to support multiple inheritance code reuse in the context of a single inheritance language. The Pharo programming language includes the second implementation of stateless traits based on the original version of Schaerli's one. Even if it is the second iteration of such an implementation, it presents several limitations. First, it does not support state in traits. Second, its implementation is monolithic *i.e.*, it is deeply coupled with the rest of the language kernel: it cannot be loaded nor unloaded. Furthermore, trait support impacts all classes, even classes not using traits. In addition, while the development tools include full support to work with classes, trait support is more limited because classes and traits do not present the same Metaobject Protocol (MOP). Finally, being monolithic and integrated in the language kernel, it is difficult to extend this current implementation.

This article describes a new *modular* and *extensible* implementation of traits: it is easily loadable and unloadable as any other package. In addition, classes not using traits are not impacted. Finally, this new implementation includes a new and carefully designed Metaobject Protocol (MOP) that is compatible with both classes and traits. This allows one to reuse existing tools as they do not require special support for traits. Then, following the semantics proposed for *stateful* traits in [6], we present a new implementation of stateful traits. This implementation is an extension of our new modular implementation.

We implemented modular traits using specialized metaclasses as our main language extension mechanism. By replacing the implementation we reduced the Pharo Language Kernel size by 15%. This model and implementation are used in production since Pharo7.0 (January 2019).

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

The literature contains multiple concepts and implementations named *traits*. Several languages offer traits: Self [22], Scala [34], SEDEL [10], Groovy [23], Fortress [1] and Scheme [18]. Language extensions have been also provided for Java [4,7],

* Corresponding author at: Inria Lille-Nord Europe, France.

E-mail address: pablo.tesone@inria.fr (P. Tesone).

Javascript [47]. However each of these implementations has its own semantics and considerations. Moreover, these implementations introduce conflicting definitions.

In this paper, we use the term *traits* to refer to traits as in Schaerli's et al. articles [44,14,33] and implementation [25].

"A trait is a set of methods, divorced from any class hierarchy. Traits can be composed in arbitrary order. The composite entity (class or trait) has complete control over the composition and can resolve conflicts explicitly, without resorting to linearization." [14]

In addition, since several variations have been made such as Stateful traits [6] and freezable traits [16], we refer to this original model and implementation as *stateless traits*. A variation of such an implementation has been introduced in Perl [40] and PHP [26]. Moreover, SEDEL [10] presents a related toy implementation of traits into a static-typed language that adds first class traits.

Traits provide support for multiple inheritance code reuse in the context of a single inheritance language. They solve the diamond problem [42,45,28]. Duplicated methods are extracted as traits which in turn are shared by many classes. Traits are also composable to form other traits [44,14,43,33]. Reppy and Turon propose a way of reducing boilerplate code by the use of traits [41] without the use of metaprogramming.

The Pharo programming language [5] includes an implementation of *stateless* traits [25]. It is the reference implementation described in the literature [44,14,33]. The Pharo 6.0 implementation of *stateless* traits fully reimplemented the original version [44,25]; however this second version has still several limitations at the implementation and model levels.

Implementation level. The implementation is *monolithic*: it is deeply coupled with the rest of the language kernel (Section 2). This prevents Pharo developers to easily load, unload or update the trait support. It restricts the flexibility of the system and makes the bootstrap process more complex [35]. In particular, classes not using traits have an useless dependency with the Trait implementation (e.g., unused instance variables, conditional code).

Model level. The implemented model presents limitations: Stateless traits do not support state definition in traits. The programmer is then forced to define accessors and initialization in all the classes using a trait that requires this instance state [6]. *Stateful traits* were proposed to extend traits with state support [6]. In this alternative, instance variables, accessors and initialization needed by a trait are defined in the trait itself but at the expense of a change of instance layout (using dictionary-based access instead of index access). However, stateful traits were never fully implemented or adopted because of the complexity of modifying the existent monolithic implementation, and updating all the existing tools.

The impact of the introduction of a new concept in an existing model and system is often not taken into account explicitly by language designers. Based on more than 10 years of experience working with traits, this paper reflects on the introduction of traits in an existing system and its impact on existing tools. Doing so, we analyse the design space of the classes and traits API in terms of the Metaobject protocol¹ exposed to the developer.

This paper describes a modular implementation for introducing traits in a programming language as well as a new implementation of stateful traits based on the semantics defined in [6]. The contributions of this work are:

- An analysis of the limitations of the current Pharo *stateless* traits implementation. This analysis goes beyond missing state support (Section 2). It shows that traits should be polymorphic² to classes and expose a revisited API to support tool builders.
- A new implementation of *stateless* traits that is compatible with classes by design, providing a smooth integration with existing programming tools (Section 3). This implementation solves the key issues of the compatibility between classes and traits from an API perspective. We redesigned a structural Metaobject Protocol [24] for classes and traits that takes into account the origin of the methods.
- A new Metaobject Protocol supporting a system with and without traits, both supporting stateless and stateful traits (Section 4).
- A new implementation of stateful traits. We extend our new implementation with stateful traits following the semantics described in the literature [6]. This new implementation is the foundation of the trait model of Pharo 7.0 and is used in an industrial setup (Section 5).
- A modular implementation. Our new implementation is modular since it is a loadable and extensible library. This new library also ensures that classes not using traits are not impacted because it relies on specialized metaclasses³ that encapsulate the integration with traits. This modular implementation gives developers the freedom to load or unload

¹ A Metaobject protocol is an API allowing to access and often reflectively change language implementation aspects otherwise hidden. Kiczales et al. [24] defines it as: "First, the basic elements of the programming language - classes, methods and generic functions - are made accessible as objects. Because these objects represent fragments of a program, they are given the special name of Metaobjects. Second, individual decisions about the behavior of the language are encoded in a protocol operating on these Metaobjects - a Metaobject protocol. Third, for each kind of Metaobjects, a default class is created, which lays down the behavior of the default language in the form of methods in the protocol."

² An object exposes the same API as another one with which it is polymorphic. In Java terms, two objects implement the same interface, so they are polymorphic.

³ A metaclass is the class of a class, thus a metaclass defines how a class behaves. Using a specialized metaclass allows one to modify the properties of classes [24,13,17]. Metaclasses provide the flexibility to assign different properties to different classes [9,15,39].

this language feature bringing flexibility to the development team [12]. It is even possible to modify the library without affecting the whole environment (Section 6).

Compared with the current stateless trait model and implementation in Pharo 6.0,⁴ our solution improves the modularisation of the existing code base: It simplifies it – it reduces the size of the Pharo Language Kernel by 15%. It speeds up the overall bootstrapping process of 30% [35] for building reduced language kernels [37] (Section 7). Section 8 presents the considerations in the different design decisions, and Section 9 compares our solution with related works. Finally, Section 10 concludes this paper.

2. Limitations of stateless traits and its monolithic implementation

The stateless traits implementation presents a series of limitations. The analysis of stateless traits limits presented by Bergel et al. in [6] still applies to the traits implementation in Pharo6.0. We do not repeat it in this article but focus on the practical problems that the Pharo development team faced over the years.

2.1. Vocabulary

In this paper, we use the following common vocabulary: A class defines or has methods and instance variables. Each class has a single superclass it inherits from. A class creates instances. Stateless traits define methods. Stateful traits in addition define instance variables. Traits cannot create instances.

Whenever a class is composed out of a trait, we say that the class *uses* the trait. A class using a trait has access to the methods (and instance variables, if applies) defined in the trait.

We use the term *traited class* to refer to classes using traits, and *normal classes* to classes not using traits. For reading convention, we prefix all traits' names with a T (e.g., TNamed). Finally, the *origin* of a method is the class or the trait that defines it.

2.2. Overloaded metaobject protocol

In normal classes, methods originate from two possible locations: Either they are defined in the class itself or they are defined in one of its superclasses *i.e.*, they are inherited. The Metaobject Protocol (MOP) of classes was originally designed to cope with this difference: there exist two families of messages to access the methods [21]:

- One family of messages is used to access the methods defined in the class (*i.e.*, selectors, methods).
- The other family, with messages prefixed with “all”, is used to access both defined and inherited methods (*i.e.*, allSelectors, allMethods).

This separation of messages provides a clear interface showing where the methods are defined.

To make stateless traits and traited classes compatible with existing tools the previously carefully designed MOP has not been revisited and redesigned, but simply overloaded. For example, the method `methods` was modified to return not only the methods defined in the class itself but also the ones contributed by the used traits. This is a problem because certain navigation operations such as editing the original method suddenly became a lot more complex to define. To identify a method origin, the traited class MOP was modified to provide a new family of methods prefixed with “local” (*i.e.*, `localSelectors`, `localMethods`). Similarly to non-prefixed versions of the methods, methods prefixed with “local” return the elements defined in the class itself. However, while the intention was good, the resulting situation is not satisfactory since, having two different MOPs forces tools to use conditional code to detect whether it is manipulating a normal or a traited class to use the correct corresponding API.

2.3. Trait and class coupled implementation

In the current version, the trait implementation is strongly coupled inside the default Pharo metaclass model [5]. There is no way to have a class without trait support. This is a problem because the vast majority of classes do not use traits. In Pharo 6.1, 127 classes use traits over the 6486 classes of the system: it means that less than 1.95% of the classes are using traits.⁵

The original objective of having a coupled implementation was to take advantage of trait reuse capabilities. However, the result, while working, makes the code logic more complex. The code that actually integrates traits with classes is embedded

⁴ This paper always refers to version 6.0 of Pharo when talking about the current implementation. And it refers to version 7.0 of Pharo when describing the new implementation.

⁵ Note that this computation does not take into account the fact that when a class root of inheritance uses a trait, its subclasses should be counted as using traits too.

```
Behavior >> isComposedBy: aTrait
"Answers if this object includes trait aTrait into its composition"
aTrait isTrait ifFalse: [ ^ false].
^ self hasTraitComposition
  ifTrue: [ self includesTrait: aTrait ]
  ifFalse: [ false ]
```

Listing 1: Mixed up implementation in Kernel.

```
AbstractTool >> browseClassVarRefsOf: aClass
| cls |
cls := aClass instanceSide.
cls isTrait
  ifFalse: [ self systemNavigation browseClassVarRefs: cls ]
```

Listing 2: Conditional code in tools.

inside the standard Pharo meta-classes: Behavior, ClassDescription, Class and Metaclass. Internally, these classes manage both the case where a class uses traits and where they do not, mostly using conditionals. Listing 1 shows an example of conditional code that is scattered all over the language kernel.

This scattered implementation prevents Pharo developers from unloading the trait implementation, or having a smaller kernel [37]. This impacts bootstrap time and future kernel evolution possibilities.

Requirements for the new implementation. Since modularity can be interpreted in different ways, we define precisely what are the requirements for a modular trait implementation. The implementation should be modular in the sense that:

- Classes not using traits should not depend on the trait implementation.
- Trait support should be loadable on demand.
- To support traits, the existing hierarchy of a class should not be modified. To benefit from trait support, a class should not be forced to inherit from a special class. This is particularly important since Pharo only supports single inheritance.

From these requirements we see that the granularity of trait support is a class and its subclasses. It means that using a trait in a class only affects this class and its subclasses but not any of its superclasses.

2.4. Wrong unified API between traits and classes

In the second attempt to propose a better API for traits and classes, changes were made by the community to propose a more unified API to the classes Trait and Class. The initial intent was that a class could understand the same messages as a trait, in particular to avoid tool builders to write complex conditional code. However, this attempt did not fully deliver its objectives. In Pharo6.0 implementation, Trait and Class were both users of the trait TBehavior. The trait TBehavior defined common methods that allowed one to use traits and classes polymorphically in some situations, for example, in the compiler.

However, traits and classes did not provide a fully compatible API. Indeed some class behavior such as instance creation did not make sense on traits. Several methods in Trait class were cancelled (*i.e.*, throw an error) or had an incompatible implementation. Since traits cannot be instantiated, their method basicNew threw an exception to avoid the creation of instances.

Other meaningless messages contained a valid but arguable implementation (*e.g.*, superclass returning nil). This inconsistent and possibly accidental behaviour produced lots of conditional code in existing tools to manage such special cases and avoid misbehaviours.

Listing 2 shows an example of conditional code. This implementation is needed because by design traits return nil to the message instanceSide where classes return the instance side class⁶; and traits by design did not support class variables. That's why, first the developer had to guard against that the argument can be a trait and that the message instanceSide returned nil, then that the argument can be a class trait (a trait applied on the class of a class), to finally be sure that the argument is a class. Listing 2 shows that the situation is even more confusing since it uses cls instead of classOrTrait for the temporary variable name.

As shown, having a non carefully thought API brings more complexity in the methods having to deal with classes and traits.

⁶ In Smalltalk, when a programmer defines a class, the system internally creates a class (*e.g.*, Person) instance of another class that is created automatically: its metaclass called Person class. The message instanceSide makes sure that when sent to a class or its metaclass, the class is returned.

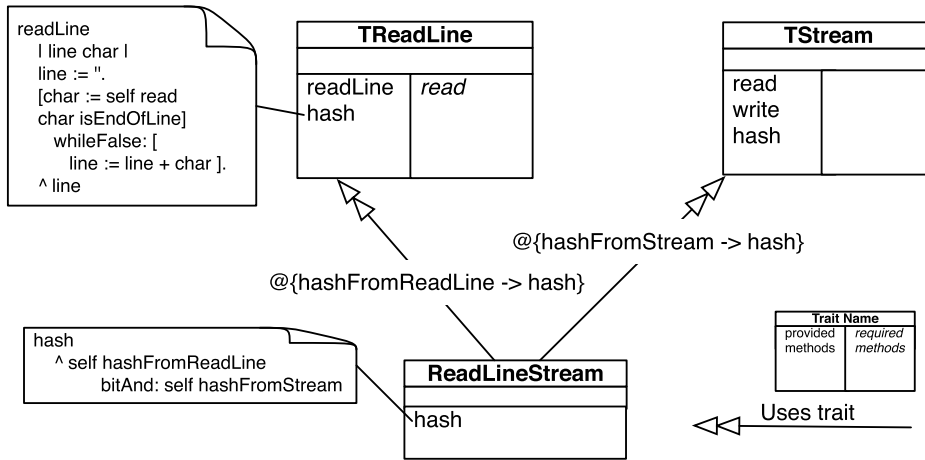


Fig. 1. The class ReadLineStream implements a stream that read lines by using two traits TReadLine and TStream.

3. A modular stateless trait implementation

Implementing a modular and extensible trait implementation presents a number of challenges that we solved in the new implementation presented in this article.

We start with an example of stateless traits as defined by Schärli et al. [44] (Section 3.1). We remind the reader with the stateless trait composition algebra (Section 3.2). Then, since Pharo's kernel is derived from Smalltalk's one, to ease understanding, we explain and sketch a simplified version of the original Smalltalk kernel that does not include traits [21] (Section 3.3). Using the first example, we present our modular stateless trait implementation (Section 3.4).

3.1. Studying a trait usage

To make the implementation description clearer, Fig. 1 describes a trait-based implementation of a Stream library using the mechanisms described by Schärli et al [44]. This library gives the ability to read lines. It reads from a stream until an end of line character is encountered.

The class ReadLineStream uses two traits: TStream and TReadLine. The first trait provides the behaviour required to read from the stream and the second one the ability to handle the reading of text lines.

The trait TReadLine defines only the methods readLine, and hash, and it requires the method read. We define the TReadLine trait as follows:

```
Trait named: #TReadLine
uses: {}
```

The trait TReadLine does not use of any other traits. The uses: clause specifies the trait composition (empty in this case).

The interface for defining a class as composition of traits is an extension of the interface used to define classes without using traits. The extended interface includes the uses: clause to express the trait composition. Here we see how ReadLineStream uses the traits TReadLine and TStream:

```
Object subclass: #ReadLineStream
uses: TReadLine @ {#hashFromReadLine -> #hash }
+ TStream @ {#hashFromStream -> #hash }
instVarNames: ''
....
```

In this example, both traits provide an implementation of hash, producing a conflict as the resulting class cannot have twice the same method. In our example, we decided to have an implementation of hash in ReadLineStream that combine both implementations existing in the traits. This implementation replaces the existing ones in the traits but requires to have access to them. So, we are applying an aliasing operation to the methods contributed by the traits. The hash from TReadLine is renamed as hashFromReadLine and the one from TStream is renamed as hashFromStream. The aliasing of the methods is done through the use of the alias operand (@) in the trait composition of the ReadLineStream class.

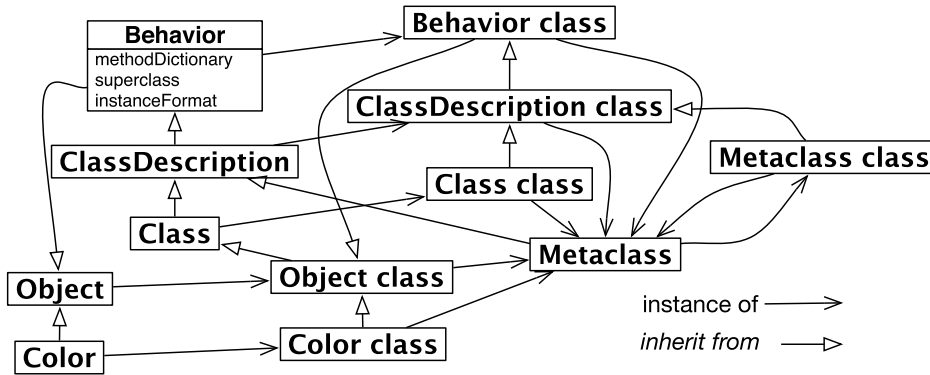


Fig. 2. Initial and traditional Pharo kernel (without traits). The class kernel is decomposed in Behavior (root of all metaclasses and classes), ClassDescription, and Class and Metaclass. Behavior defines the essence of a class (superclass, a method dictionary and a way to describe instance state [5]).

3.2. Stateless trait composition algebra

The implemented trait composition follows the semantics defined in the original traits implementation [44]. The trait composition algebra allows us to compose different traits to form a more complex one. Therefore, classes can use different traits along with a trait composition algebra to express the conflict resolution [14]. This section can be skipped by a reader knowing trait composition. It is a short presentation of the three traits composition operators that are loaded when the traits library is loaded.

Sequence. The sequence operation allows us to combine two or more existing traits. The traits used in the sequence does not have any priority. The resulting combination has all the methods defined in the original traits. Sequences are created using the + operator.

```
Object subclass: #AnExampleClass
  uses: T1 + T2 + T3
  instanceVariableNames: "
  ...
```

Alias Method. When two or more traits define the same method, a conflict occurs. One alternative to solve the conflict is to alias one of the conflicting methods with another name using the @ operator. This operator receives an array of associations (->) of the new name and the original name. The resulting composition has all the methods defined in the trait but with the aliased one renamed.

```
Object subclass: #AnExampleClass
  uses: T1 @ { #new -> #original }
  instanceVariableNames: "
  ...
```

Remove Method. When a conflict exists between two methods, another alternative to solve it is to remove one of the conflicting methods. The - operator allows us to generate a new composition with a method removed. This composition includes all the original methods but the removed one.

```
Object subclass: #AnExampleClass
  uses: T1 - #aMethodToRemove
  instanceVariableNames: "
  ...
```

3.3. Starting with a implicit metaclass kernel

The original Smalltalk kernel that does not include traits [21,9,5]. Fig. 2 presents the original metaclass kernel of Pharo and Smalltalk, its ancestor: as in a traditional implicit metaclass kernel [5]. Each class has its own metaclass e.g., the class Object is instance of the metaclass Object class. The metaclass Object class is instance of the class Metaclass. The class Metaclass is instance of the metaclass Metaclass class and the metaclass Metaclass class is instance of the class Metaclass. Object class inherits from Class.

Fig. 3 shows a simplified view of this original kernel for presentation purposes. We compacted the Behavior, ClassDescription, Class and Metaclass in the class Class as in CLOS [24]. Each class in the system is a first class citizen with a




Fig. 3. Simplified traditional Pharo kernel (without traits). We compacted the Behavior, ClassDescription, Class and MetaClass in the class Class as in CLOS [24].

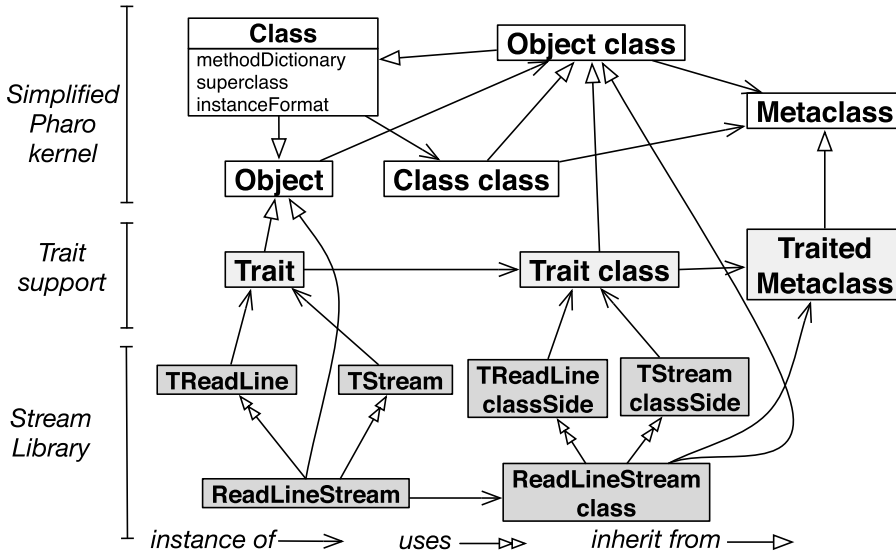


Fig. 4. The traits TStream and TReadLine are instances of Trait, their class-side counter parts, TStream classSide and TReadLine classSide are instances of the class Trait class. The class ReadLineStream is an instance of the class ReadLineStream class which is an instance of the new metaClass TrainedMetaClass that manages classes using traits.

superclass, a method dictionary, and an instance format description. The system uses metaClasses to provide a common behaviour to the classes.

This simplified kernel has no support for traits. Using this simplified kernel we will present the modular implementation of stateful traits that solves the problems mentioned in Section 2.

3.4. Extending the kernel with a modular stateless trait implementation

This section presents the extension of the kernel (Fig. 3) to support traits as well as an instantiation of this meta-model to represent the trait-based stream library (Fig. 1).

To support stateless traits in a modular way, our implementation takes advantage of the existing metaClass support shown in Fig. 2. As shown in Fig. 4, we define a specialized metaClass named TrainedMetaClass to be used by the traited classes and by the Trait class [24,9]. With the new implementation:

- Classes using traits are not instances of MetaClass but of TrainedMetaClass.
- Normal classes are kept unmodified, as instances of MetaClass class.

Following the trait model [25], a traited class uses a trait and its traited metaClass uses the corresponding class-side trait. Indeed when a trait defines a method sending a message to its class side, the corresponding method should be available on the composing class [9]. Therefore, traits are defined and applied as a pair of traits (a trait and its class side trait) to a given class.

All the behaviour related to traits is removed from the kernel classes and it is available in the trait implementation. Neither MetaClass, Class or Behavior have any trait related behaviour. As depicted on Fig. 4, all the trait behaviour is defined in TrainedMetaClass and in the traited classes. Adding the required behaviour to TrainedMetaClass is a simple operation because

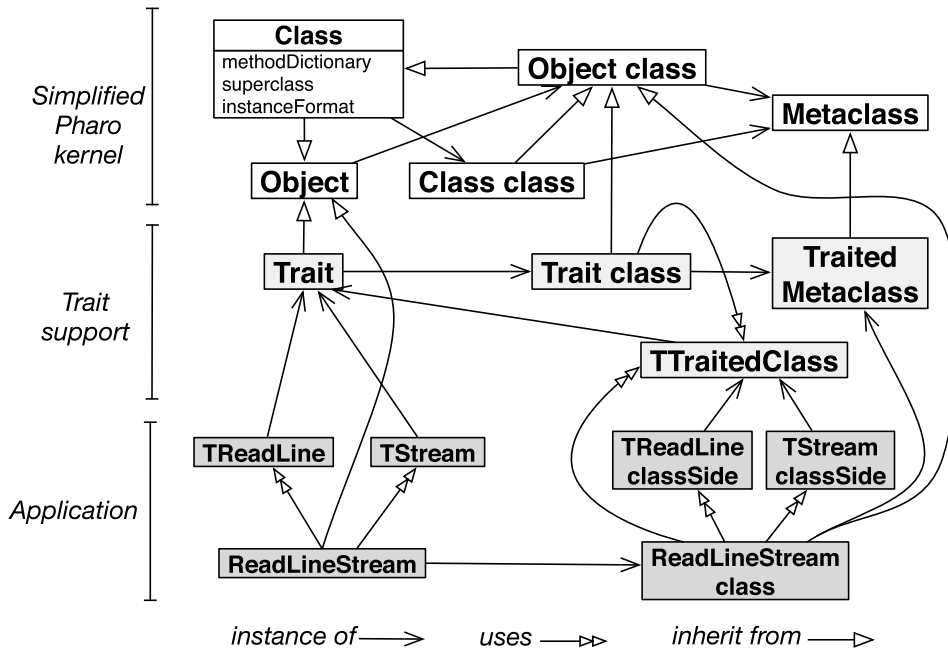


Fig. 5. The final model of our implementation, including the `TTraitedClass` trait. This trait provides the required behaviour to handle traits to the classes using traits.

`TraitedMetaclass` is a regular class in the environment, it extends the behaviour of `Metaclass` through inheritance and overriding. However, integrating the behaviour to support traits into the system is not so simple. We need to integrate this new implementation supporting traits into classes and add new methods and instance variables to handle the trait composition.

In our implementation (presented in this article), the additional behaviour needed in traited classes is factorized in trait named `TTraitedClass`. This trait is automatically included in all metaclasses of traited classes, as shown in Fig. 5. By doing so, the kernel classes of Pharo do not have any reference to the trait implementation and the parallelism of classes and metaclasses hierarchies is preserved.

Using this model, we achieve modularity by packaging `Trait`, `TraitedMetaclass` and `TTraitedClass` in a separate trait library that is imported explicitly by users.

4. Designing a modular metaobject protocol

To have a good integration between the tools and the underlying model, the Metaobject Protocol has to be clear and consistent. To allow the extension of kernel language with a modular trait implementation, the MOP should also support new extensions in a coherent and compatible fashion.

We propose a MOP with two sets of messages, a set for handling *locally defined elements* and other set to handle *contributed elements*. An element is either a method or an instance variable.

- The *locally defined elements* are the elements directly defined in the class or trait.
- The *contributed elements* are the elements defined elsewhere but nevertheless accessible to the class or trait.

Considering the language kernel without traits, the contributed elements of a class are the set of inherited elements. When the trait library is loaded, the contributed elements of a traited class are the list of both inherited elements and elements defined in the used traits. Note that such contributed element list should take into account traits semantics: the fact that local elements (from composer class or trait) take precedence over contributed ones (inherited or used traits).

In the example of `ReadLineStream` (Fig. 1), the `ReadLineStream` class has the locally defined method `hash`, the contributed methods inherited from `Object` (for example `equals`: and `printString`) and the contributed methods from using the `TReadLine` and `TStream` traits.

The Pharo MOP is extended with new messages to support traits. For the sake of clarity, we only present hereafter the messages related to methods. However, the same structure is applicable to other elements of a class such as its instance variables. Note this Metaobject Protocol is shared by classes and traits:

methods. Returns all the methods defined locally in the given class.

selectors. Returns all the method selectors defined locally in the given class.


```
Class >> contributedMethods
"Returns the methods in the superclass chain"
^ self superclass allMethods
```

Listing 3: Implementation in Class.

```
"Defining the class Account without traits"
Object subclass: #Color
instanceVariableNames:
' red blue green alpha '
classVariableNames: ''
package: 'Colors'
```

Listing 5: Class definition without traits.

```
TTraitedClass >> contributedMethods
"Concatenate the methods from superclass and from the
used traits"
^ self superclass allMethods , self traitComposition
allMethods
```

Listing 4: Implementation in TTraitedClass.

```
"Defining the class Employee using a trait composition"
Object subclass: #ReadStream
uses: TReadLine @ {#hashFromReadLine -> #hash }
+ TStream @ {#hashFromStream -> #hash }
instanceVariableNames: ''
classVariableNames: ''
package: 'Streams'
```

Listing 6: Class definition with traits.

contributedMethods. Returns all the methods contributed to this class. In the case of normal classes, it includes the methods in the class hierarchy. In the case of traitled classes, it includes the methods in the class hierarchy and the methods contributed by the traits.

contributedSelectors. Returns the selectors of those methods returned by contributedMethods.

allMethods. Returns all the methods, including the locally defined methods and the contributed methods.

allSelectors. Returns the selectors of those methods returned by allMethods.

originOfMethod: aMethod. This message returns the class or the trait that defines aMethod.

Using this Metaobject protocol, tools access elements in a class and elements contributed to a class without caring about their origin. The origin of an element is used when modifying the original element (i.e., the method defined in a trait in case of a trait). Moreover, this MOP delegates the resolution of the elements to the corresponding class by exploiting polymorphism to have a modular language extension. The contributed elements of a normal class (without using traits) are calculated by the method defined in Class (Listing 3) and the contributed elements of a traitled class are calculated by the method defined in TTraitedClass (Listing 4). So, the implementation in the kernel classes (such as Class) does not know how to include the contributions from traits, this is resolved in the code implemented in TTraitedClass.

4.1. Traited class definition

To have a real modular trait implementation, the syntax to define traitled classes is also loaded with the trait library. Listing 5 and 6 show the definition of a traitled class (ReadStream) and a normal class (Color). As shown, the definition is not the same for traitled classes and normal classes, the class definition should be extended to introduce trait composition.

As the definition of classes in Pharo is done through message-sends to the superclass, a new definition is added as an extension method to Class. This extension method allows all the classes in the system to generate new subclasses using traits. A module in Pharo is able to contribute to existing classes through extension methods [36]. When a module containing extension methods is loaded, the extension methods are installed in the target class.

Each class in the system is responsible to regenerate its own definition. So, each class is asked its definition through the definition message, the new class definition syntax is specialized adequately in the TraitedMetaClass and TTraitedClass.

5. Supporting stateful traits

In addition to offer a modular and extensible stateless trait implementation, our discussions with the designers and engineers of the Moose analysis platform show that there was a real need for traits with state. Moose is a software and data analysis platform [31]. Moose meta models represent a lot of data about the language elements and the stateless traits could not capture this important aspect of the representation.

Following the semantics proposed for *stateful* traits in [6], we extended our new implementation presented in previous sections to support state in traits. This new implementation demonstrates the extensibility of our implementation design by introducing the support of instance and class variables in traits and their initializations.

5.1. Studying a stateful trait example

Before presenting this extension, we first present an example extracted from the original stateful traits article [6]. Fig. 6 shows this example about an implementation of a SyncStream. This class extends the behaviour of a stream adding a lock to synchronize its accesses.

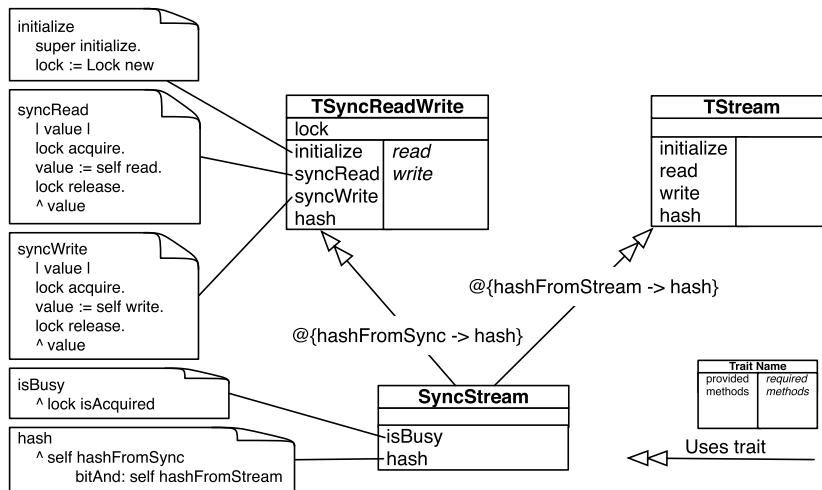


Fig. 6. The class SyncStream uses two traits TSyncReadWrite and TStream.

The class SyncStream uses the traits TStream and TSyncReadWrite. The trait TSyncReadWrite defines the variable lock, three methods syncRead, syncWrite and hash, and requires methods read and write.

Following the specifications of the stateful trait model, our new implementation extends the mechanism for defining traits. We define the TSyncReadWrite trait as follows:

```

Trait named: #TSyncReadWrite
uses: {}
instVarNames: 'lock'
    
```

The trait TSyncReadWrite does not use any other traits and it defines an instance variable named lock. The uses: clause specifies the trait composition (empty in this case), and instVarNames: lists the variables defined in the trait (i.e., the variable, lock). The interface for defining a class as a composition of traits is the same as with stateless traits. Here the definition of SyncStream using the traits TSyncReadWrite and TStream:

```

Object subclass: #SyncStream
uses: TSyncReadWrite @ {#hashFromSync -> #hash }
+ TStream @ {#hashFromStream -> #hash }
instVarNames: "
....
    
```

In our implementation all the instance variables are accessible to methods defined in the class and in all the used traits. In the example, the method isBusy implemented in SyncStream accesses the instance variable defined in TSyncReadWrite. The final set of variables of a traitled class contains the variables defined in its hierarchy and the variables defined in the used traits.

When a traitled class is defined, its class and instance variables are computed using a simple flattening process. Then methods contributed by its used traits are then re-compiled into this traitled class. During this compilation process, the index of the instance variables are recalculated to integrate all the variables defined in the class and contributed by inheritance and used traits.

Note that an implementation based on method copy-down⁷ as used in languages supporting mixins could be evaluated in the future [2,3].

The stateful trait composition algebra extends the one of stateless traits (see Section 3.2). It includes operations required to resolve conflicts arising with instance variables following the semantics described in [6]. The extended operations are: renaming (@@) and removing (--) of instance variables.

⁷ Copy-down is a way to minimize code repetition in mixin implementation. "Methods that do not access instance variables or super are shared in the mixin. Methods that access instance variables may have to be specialized for the invocation, where the instance variable access is customized according to the structure of instance of the invocation." [2] The idea is that since the only method impacted by changes to object internal representation (due to mixin application) are methods accessing fields, it is enough to copy down in the subclasses (classes that represent mixin application to another class) the methods accessing fields. All other methods are then applicable to all the mixin application classes.

```
Trait named: #TNamed
uses: {}
instanceVariableNames: 'name'.
```

```
Trait named: #TWithRole
uses: {}
instanceVariableNames: 'name'.
```

```
Object subclass: #Person
uses: TNamed + (TWithRole @@ {#name -> #roleName})
instanceVariableNames: ""
classVariableNames: ""
package: 'APackage'
```

Listing 7: Example of conflict resolution by renaming the name instance variable of TWithRole into roleName.

```
Trait named: #TNamed
uses: {}
instanceVariableNames: 'name'.
```

```
Trait named: #TFormalName
uses: {}
instanceVariableNames: 'name prefix'.
```

```
Object subclass: #Person
uses: TNamed + (TFormalName -- #name)
instanceVariableNames: ""
classVariableNames: ""
package: 'APackage'
```

Listing 8: Example of conflict resolution by merging the two name instances variables coming from TNamed and TFormalName.

5.2. Conflict resolution

When two or more instance variables with the same name exist in the trait composition, there is a conflict. A conflict occurs in two different scenarios: (1) the instance variables have the same name but their values cannot be shared or their usage is incompatible, and (2) the instance variables should be merged as they contain the same information and it is required to be shared by both traits.

Listing 7 illustrates the first scenario. The Person class uses two traits (TNamed and TWithRole), both traits define and use the name instance variable. However, these instance variables are used in incompatible ways. The TNamed trait uses it to store the person name, but the TWithRole trait uses it to store the role name. In this scenario the instance variables should be renamed. During the flattening process both instance variables (name and roleName) are created and the methods in TWithRole using the original conflicting variable are rewritten to use the new name (roleName). Renaming is achieved using the @@ operator on TWithRole. This operator creates a new composition with the renamed instance variable definition.

Listing 8 shows a conflict solved through the merge of two instance variables. In this scenario, both traits (TNamed and TFormalName) are using a name instance variable. Both are using their variable to store the person name, so their usage is compatible. Having a duplicated variable with the same information is not desired. To solve this problem, we use the – operator to remove the instance variable from TFormalName, leaving the one defined in TNamed. During the flattening process only one instance variable is created and all the methods are using it.

A non-conflicting scenario. This third scenario is about two traits with different instance variables but storing the same information. This is not a conflict, but merging the instance variables is desired to not duplicate the same state. To solve this scenario one of the instance variables is renamed and then the result is merged with the other trait. To merge them, we are removing the name instance variable from the resulting composition.

Listing 9 shows this scenario. The trait TDisplayName stores the person name in the personName instance variable while TNamed stores it in an instance variable named name. As said before, it is not a conflict, but the information should not be duplicated in two different instance variables. To merge, we first rename the instance variable in TDisplayName from personName to name. This will replace all the users of personName with name. Then as both traits have the name instance variable, we remove it from one of the traits. The resulting class will have only one instance variable name.

```
Trait named: #TNamed
uses: {}
instanceVariableNames: 'name'.
```

```
Trait named: #TDisplayName
uses: {}
instanceVariableNames: 'personName'.
```

```
Object subclass: #Person
uses: TNamed + ((TDisplayName @@ {#personName -> #name}) -- #name)
instanceVariableNames: ""
classVariableNames: ""
package: 'APackage'
```

Listing 9: The instance variables in TNamed and TDisplayName are merged to use the same instance variable.

```
TTraitedClass >> new
^self new
initialize;
initializeTrait;
yourself.
```

Listing 10: Implementation of new method.

```
SyncStream >> initializeTrait
self initializeTSyncReadWrite.
self initializeTStream.
```

Listing 11: Generated initializeTrait.

5.3. Trait initialization

The initialization of instances of traitled classes is an important point to address [8,30]. The question is how to compose and invoke the initialization of used traits from the composite one (either a trait or a class). In particular we do not want that the class initialization has to be changed each time a used trait (of the traits the class uses) changes. We propose a simple strategy to support this.

The initialization of the instance variables defined in traits is performed through the redefinition of the initializeTrait method. Each trait requiring the initialization of instance variables overrides this method. If a trait does not include this method, an empty method is generated during the flattening process. When an object of a traitled class is instantiated, the initialization code on the class and the initialization code on the trait is executed. Listing 10 shows how the trait TTraitedClass, which is applied to the class side of all the traitled classes, redefines the new method.

A scenario that should be treated correctly is when a traitled class uses a trait composition including several traits with initialization code. In this scenario, all the initialization code should be installed in the traitled class. To avoid name clashes, the different initializeTrait methods are aliased, and a new initializeMethod calling the aliased methods is generated. In Listing 11 the initializeTrait method invokes the initializeTSyncReadWrite and initializeTStream methods that have been generated by the aliasing.

5.4. Metaobject protocol support

Our proposed MOP supports stateful traits following the same idea expressed for stateless traits. The same operations described for methods are applicable to instance variables. As said in Section 4, the instance variables related operations have the following semantics:

slots. Returns all the instance variables⁸ defined locally in the given class.

contributedSlots. Returns all the instance variables contributed to this class. In the case of normal classes, it includes the instance variables in the class hierarchy. In the case of traitled classes, it includes the instance variables in the class hierarchy and the instance variables contributed by the traits.

allSlots. Returns all the instance variables, including the locally defined methods and the contributed methods.

originOfSlot: anInstanceVariable. Returns the class or trait defining the instance variable.

Using this MOP, the tools are implemented without differentiating if they are using normal classes, traits or traitled classes.

⁸ Slots are the new names for first class instance variables in Pharo7.0.

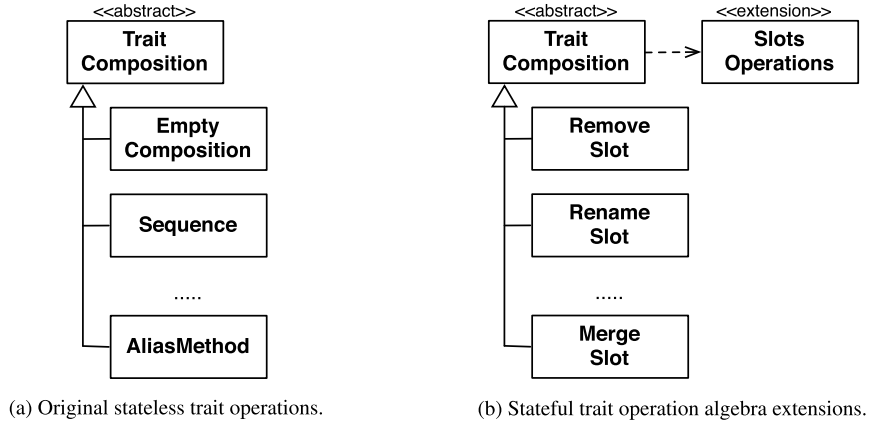


Fig. 7. Extensible trait composition algebra: the original stateless trait operations are implemented in an extensible way (a) and are extended to support stateful trait operations.

6. Implementation details

We validated our solution by replacing the implementation of traits in Pharo 7.0. This implementation is part of the Pharo 7.0 deployed and released version. It includes the implementation described in Section 3, the proposed MOP in Section 4, and the proposed extension in Section 5.

This implementation requires the solution of a number of technical issues. The solutions are presented in this section.

6.1. Polymorphic classes and traits

To simplify the integration of traits with existing tools, we represent traits as regular classes. They are subclasses of the Trait class. The only restriction we impose by design is that traits cannot be extended by means of inheritance. The common Trait superclass allows one to easily specialize trait specific methods and implement a clean interface polymorphic with regular classes. For example, the Trait class overrides compile: to compile and install a method locally in a trait and then propagate such a change to all its users.

6.2. Implementing an extensible algebra

To be able to extend the set of operations supported by the system, we implemented the operations of the trait composition algebra as a set of first-class objects. We represent such first-class objects as classes subclasses of TraitComposition. TraitComposition defines methods to handle the composition. We represented first the stateless operations as shown in Fig. 7a. Then such set of operations is extended to support the stateful traits operations and resolution of conflicts (Section 5).

6.3. Flattening and method dictionaries

Our solution handles the construction of classes in the same way as the Pharo 6.0's implementation: the methods contributed by a trait are flattened in the class method dictionary. So, the virtual machine (VM) executes the methods defined in the class, the ones contributed by the used traits and the inherited ones using the default VM method lookup, without execution penalties. However, traited classes have two method dictionaries.

- A first method dictionary is used to store all the methods visible to the VM. This method dictionary is internal to the implementation and it is not accessible through the MOP.
- To be fully compatible with the existent tools, traited classes and metaclasses have an additional method dictionary with only the methods defined in the class. The second method dictionary is the one that is accessible through the MOP. This way the method dictionary that is visible from the outside has the same semantics as the method dictionary in a normal class. Any change to the available methods of the class is reflected in the hidden method dictionary.

Each method installed in the traited class knows the original trait where it was defined. This information is stored as additional properties in the method. The access to this information is performed through the class metaobject protocol, and the existence of this additional property is an implementation detail that is again not part of the MOP and as such not exposed to the user.

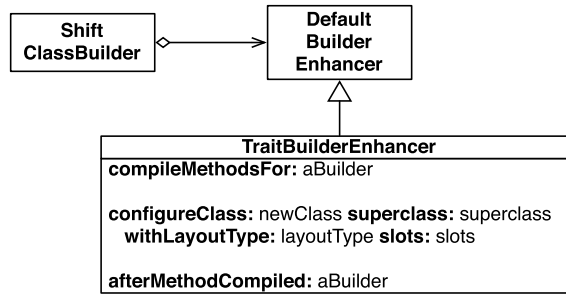


Fig. 8. The ShiftClassBuilder, used by our implementation, is extensible through pluggable build enhancers.

In the Pharo implementation we see that the memory impact of having a second method dictionary is minimal. This second method dictionary shares the symbols and methods with the default one. The implementation only requires extra memory for the collection itself. Measuring the Pharo 7.0 image the average number of local methods in the 545 classes using traits is 2. So, the average impact per class is of 32 bytes in 32 bits images, and 64 bytes in a 64 bits image (4 instance variables for the dictionary itself and 2 instance variables for each of the 2 associations used).

6.4. Supporting live programming

Pharo is a live programming environment. It allows the developer to modify its code while executing it. So, any change made to a trait should be propagated to the users of that trait. To do so, a trait knows all its users and when it is modified (*i.e.*, method or instance variable modification) it then notifies its users. Also the users of a trait are subject to change (*i.e.*, a class is defined to use the trait, or a class using the trait is modified not to use it any more).

This propagation mechanism is implemented using two strategies. First, the notification mechanism is implemented through an *announcement* event system. The announcement event system generates events whenever the system is modified. All the tools requiring to react to these changes are subscribed to these events. Second, a trait intercepts its modifications by overriding the corresponding methods in the Trait class.

The announcement event system is used to update the users of a trait when a traited class is created, modified or removed from the system. The modification overriding is used when a trait is modified (*i.e.*, methods or instance variables) and the users should be updated.

6.5. Class builder and installer replacement

The creation and installation of classes in Pharo are performed by a class Builder and a class Installer. These system components perform all the operations needed to create, modify and install the classes in the environment. To be able to have a modular implementation of traits, some changes to them were needed.

Our solution replaces the class builder and installer with variants that support extensions. It implements a new set of class builder and installer, called the ShiftClassBuilder and ShiftClassInstaller. By default, this new implementation does not have support for traits, as it is intended to be included in Pharo kernel without traits. However, this class builder allows the extension of the building process via plugins, called enhancers.

The class builder and installer are configurable using different builder enhancers [46]. The builder enhancers contribute to the building and installation process (Fig. 8). It also allows one to configure the class of the metaclass to use and how it is created. The implementation of traits is managed by a specific enhancer: TraitBuildEnhancer. It configures the class and metaclass to have the instance variables added by the traits, it also installs all the methods provided by the trait composition. This is performed through the overriding of the defaults actions of the build enhancer.

7. Validation

As said before, the solution has been applied to Pharo 7.0. This allows us to minimize the language kernel that needs to be bootstrapped in the creation of a Pharo 7.0 image.⁹ The Pharo 7.0 image is created after each commit from scratch, generating a Pharo image from source code [35]. In the image generation process, a small image is bootstrapped. This image contains the minimal language kernel to be able to load code. After this minimal image is created all the libraries and tools are loaded, including the new trait implementation. This image is the one that is shipped.

Removing Traits from Kernel. One of the objectives implementing trait support as a library was to remove the trait support from the Pharo language kernel. The language kernel includes all the elements needed to load new pieces of code and to

⁹ A Pharo image is a platform agnostic binary file containing all the objects representing the system: classes of the complete Pharo distribution. A default image consists in around 6000 classes and 120,000 methods.

Table 1
Reduction of code after removing traits from Pharo Language Kernel.

	Old Impl.	New Impl	Diminution
Lines	147,248	124,642	22,606 (15.35%)
Packages	49	44	5 (10.20%)
Classes	694	587	107 (15.42%)
Methods	13,937	11,040	2,897 (20.79%)

execute Pharo code. The removal represents a diminution of 22,606 lines of code (15.36%) and 2,897 methods (20.79%).¹⁰ Table 1 shows a detail of removed code. Having a smaller kernel allows us to create smaller images with only the required functionality.

Removing the trait implementation from the language kernel speeds up the bootstrap process of 30%, going from an initial bootstrap process of 16 minutes to a bootstrap process of 11 minutes. On the contrary, the time to load the libraries has increased of only 1 minute. Producing an overall speed up of 20% for the whole process.¹¹

About traits and classes polymorphism. The previous implementation tried to define a polymorphic API for classes and traits. It leads to ad-hoc situations where while both entities offered the same API, the developers had still to be aware that he manipulated a class or a trait (for example, it does not make sense to display the superclasses of a trait). It produced unnecessary conditional code to differentiate traits and classes. This problem is not only present in the kernel, but also in all the tools and libraries manipulating traits and classes of the new implementation.

Some of the tools affected are the compiler (*Opal*), version managers (*Epicea* and *Monticello*), code navigation (*GT-Spotter*), modeling framework (*Ring*) and class browser and editor (*Nautilus* and *Calypto*). The new implementations represent a reduction of 6,557 lines of code distributed in 89 packages. This impact is visible in the language kernel, but it is more important when analyzing the full system.

Moose. Moose is a software and data analysis platform [31]. Since its version 7.0 it uses stateful traits as a key basic brick for modeling languages (Java, Pharo, PostgreSQL, PowerBuilder, Ada, VisualBasic...) and performing software evolution analyses. The core of Moose defines 123 stateful traits representing elementary program concerns. Such traits are then reused to represent different language and language constructs. This is not a real validation but Moose latest version is a heavy user of stateful traits.

8. Discussion

We discuss now the impact of the new implementation and revisit some design decisions.

8.1. Tool support

Modern IDEs are composed of many dedicated tools. Pharo offers out of the box: message browsers, a cross referencer, navigation tools, package/class browser, change browser, advanced object inspectors, debuggers, automatic completion, VCS support, and static quality rule engines to name a few.

A loadable trait implementation should include all the resources and modifications needed by the tools to fully support them. Due to the fact that our traits are polymorphic to classes, we are able to remove most of the conditional code in the tools. Some tools should be adapted to use the new Metaobject Protocol for normal and traitled classes. However, this topic is still an ongoing effort, as a definitive integration requires modifying existing tools and most of the tools are not prepared to be extended or to allow inclusion of pluggable components. The same consideration is applicable to the new class definition needed by traits. There exists tools that parse the class definition syntax, expecting to receive a given format. These tools should be modified to allow an extensible syntax or to use the messages exposed in the MOP of the classes.

The work left for a definitive integration in Pharo is out of the scope of this work and it will take a couple of years. It is related to the overall quality improvement of existing tools. Pharo is a complex platform composed of about 400 packages and aggregating several subprojects. As such modifying all the required elements is a continuous process requiring time, since such process should always provide useful versions to the users of the platform. This work will be handled with the Pharo community and its industrial consortium.

¹⁰ This information was extracted from the Pharo Github repository calculating the difference before and after applying the changes.

¹¹ Information recovered from the Pharo CI Server Statistics https://ci.inria.fr/pharo-ci-jenkins2/job/Test_pending_pull_request_and_branch_Pipeline/job/development/.

8.2. About design

In our solution we decided to separate the implementation of traits from the kernel classes through the use of a different metaclass for traitled classes. This way all the trait related behaviour is loaded as a library. Other possible alternative is having an extensible metaclass in the kernel of the language. However, making an extensible metaclass in the kernel requires having a more complex kernel than having a simpler implementation.

As explained when presenting our solution (Section 3.4), we have decided to define the behaviour requiring to implement traits in a trait, named `TTraitedClass`, that is included whenever a traitled class is created. The `TTraitedClass` trait includes all the methods that should be overridden in the `Class` class to support traitled classes. By doing so, we allow one to have the ability to inherit from any class in the system and this without modifying the `Class` class.

Our solution extends the existing MOP with messages to calculate the origin of class elements. This decision makes the tools independent of the reuse mechanism used. If a given tool wants to modify the definition of an element, the defining class is easily reachable and modifiable.

We decided to have the definition of traits as classes. Another alternative is the definition of traits as independent objects. As it is desirable to have polymorphism between traits and classes to share tool support, they should implement the same MOP. Taking advantage of traits being classes allows a simple implementation of the shared Metaobject Protocol.

We have decided to have two different method dictionaries. By doing so, the VM uses the method dictionary with the flattened traits, and the tools only see the methods through the MOP. This decision creates a separation between the runtime model expected by the VM and the logical model exposed by the tools. Also by restricting the tools to use the method dictionary with the locally defined methods only, allows us to reuse the implementation of all the methods in `Class` class; as these methods expect the method dictionaries to only have local defined methods.

Our solution requires a modified version of the class builder and installer. This version allows the extension of the class building process.

9. Related work

Other languages, such as Scala [34], Rust [29], Groovy [23] or Self [22] implement traits as a part of their language kernel. They are integrated in the language and cannot be unloaded. In addition traits in Scala, Rust and Self do not support the same composition semantics than traits defined in [44,14]. Their implemented semantics are closer to Mixins [3].

There are languages, such as Javascript [19,47], Python [27] and Racket,¹² that also implement traits in a loadable library. However, these implementations have been developed for a language kernel without trait support. So, there is no need to modify the existing language kernel or the given tools. Finally, the tools for these languages do not support traits.

There are solutions to add traits to existing static-typed languages [10,20] and specially to Java [4,7,32]. However, these solutions focus on the types and type checking challenges introduced by traits. They modify the compiler or generate equivalent code as the one using traits. They do not provide tools or extensible support for development using traits.

Pharo already includes a working traits implementation [14], although this implementation is heavily coupled with the kernel of the language, making impossible to make it a loadable library [25].

Talents are instance specific traits. The early prototype implementation [38] provides ways of expanding Pharo through a loadable library, but this solution has a big execution performance impact as the lookup of methods is solved in the language without taking advantage of the performance optimizations of the VM. In addition, talent methods are defined as plain strings passed as arguments to methods and there is no tool support of any kind.

Razavi et al. [39] already extended an existing language through the definition of custom metaclasses. However, they applied such solution to implement Adaptive Object Models. They do not apply the solution to modify the reuse mechanisms available in the language.

Cazzola et al. [12] also provided support to modify the semantic of the language and the reuse mechanisms through dynamic loadable modules. However, their solution is centred on a modular interpreter architecture. They modify the language modifying the interpreter of it. This solution is not applicable to our case for performance reasons.

Bouraqadi et al. [9,11] show how classes are extended in a reflective language using metaclasses. Metaclasses are then used as an extension mechanism to add new properties to classes in a per-class basis. We leverage also metaclasses to implement traits as they did for example with Mixins. Differently from them, we also analyse the impact of changing the current trait implementation in the programming environment and its provided tools. We also introduce a modular reflective MOP that prevents new language extensions to modify the existing runtime keeping compatibility.

Malayeri et al. present CZ [28]. CZ is an alternative to traits to solve the diamond problem. It uses an alternative way of combining classes extending the multiple inheritance with restrictions and imports. However, this solution is centred in the development of a language and does not provide tools or a runtime library.

Tesone et al. [46] shows a general architecture to implement modular extensions to a programming language preserving the performance while running in a single inheritance VM. It shows how a modular design of the underlying mechanisms (i.e., class builder) supports the introduction of mixins, traits, and multiple inheritance. The implementation presented in the

¹² <http://racket-lang.org/>.

present article uses the same idea (i.e., metaclasses) and a modular class building process. However, it is less general and focused on traits. In addition, it solves the more general problem of class and trait API mismatch and designs an adequate meta-object protocol taking traits into account.

10. Conclusion

In this paper, we addressed the problem of how to modularize a core part of Pharo environment. We took the existing implementation of traits and separated it from the language kernel. Once we had all the mechanisms identified, we proposed a novel implementation that uses a specialization of kernel metaclasses to support traits. This new implementation exposes a common MOP with the existing classes. Since this implementation is based on the class implementation mechanisms, it offers a better support for tools support.

The use of our proposed solution and its implementation allowed us to reduce the size and complexity of Pharo Language Kernel. This reduction sped-up the bootstrapping process in 30%. A faster bootstrapping process improves the ability to modify the language, the infrastructure and the tools. This result is crucial for having a community maintained programming environment that is fully tested and constantly built from sources. Results of this impact are daily noticeably in the results of the Pharo Continuous-Integration process.¹³

Finally, having a loadable independent stateful traits library also achieves other two goals: (1) it produces an increase of the modularity of the system allowing users to select the features used by their applications, and (2) opens the door to the research and implementation of alternative reuse models using similar dynamic and late-bounding techniques. The first objective allows users of Pharo to easily produce tailored configurations for their development and final applications.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work is supported by I-Site ERC-Generator Multi project 2018-2022. We gratefully acknowledge the financial support of the Métropole Européenne de Lille.

References

- [1] Eric Allen, David Chase, Joe Hallett, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L. Steele Jr, Sam Tobin-Hochstadt, Joao Dias, Carl Eastlund, et al., The fortress language specification, *Sun Microsyst.* 139 (140) (2005) 116.
- [2] Lars Bak, Gilad Bracha, Steffen Garup, Robert Griesemer, David Griswold, Urs Hölzle, Mixins in strongtalk, in: *ECOOP'02 Workshop on Inheritance*, June 2002.
- [3] Gilad Bracha, William Cook, Mixin-based inheritance, in: *Proceedings OOPSLA/ECOOP'90, ACM SIGPLAN Notices*, vol. 25, October 1990, pp. 303–311.
- [4] Lorenzo Bettini, Ferruccio Damiani, Xtraitj: traits for the Java platform, *J. Syst. Softw.* 131 (2016) 07.
- [5] Andrew P. Black, Stéphane Ducasse, Oscar Nierstrasz, Damien Pollet, Damien Cassou, Marcus Denker, Pharo by Example, Square Bracket Associates, Kehrsatz, Switzerland, 2009.
- [6] Alexandre Bergel, Stéphane Ducasse, Oscar Nierstrasz, Roel Wuyts, Stateful traits, in: *Advances in Smalltalk – Proceedings of 14th International Smalltalk Conference (ISC'06)*, in: LNCS, vol. 4406, Springer, August 2007, pp. 66–90.
- [7] Lorenzo Bettini, Ferruccio Damiani, Ina Schaefer, Fabio Strocchio Traitrecordj, A programming language with traits and records, *Sci. Comput. Program.* 78 (5) (May 2013) 521–541.
- [8] Viviana Bono, Jarek Kuśmierk, Mauro Mulatero, Magda: a new language for modularity, in: *ECOOP 2012–Object-Oriented Programming*, Springer, 2012, pp. 560–588.
- [9] Noury Bouraqadi, Thomas Ledoux, Fred Rivard, Safe metaclass programming, in: *Proceedings OOPSLA'98*, 1998, pp. 84–96.
- [10] Xuan Bi, Bruno C.d.S. Oliveira, Typed first-class traits, in: *32nd European Conference on Object-Oriented Programming (ECOOP 2018)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [11] Noury Bouraqadi, Safe metaclass composition using mixin-based inheritance, *J. Comput. Lang. Syst. Struct.* 30 (1–2) (April 2004) 49–61.
- [12] Walter Cazzola, Albert Shaqiri, Open programming language interpreters, *The Art, Science, and Engineering of Programming 1* (2017).
- [13] Scott Danforth, Ira R. Forman, Derived metaclass in SOM, in: *Proceedings of TOOLS EUROPE'94*, 1994, pp. 63–73.
- [14] Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, Andrew P. Black, Traits: a mechanism for fine-grained reuse, *ACM Trans. Program. Lang. Syst.* 28 (2) (March 2006) 331–388.
- [15] Stéphane Ducasse, Nathanael Schärli, Roel Wuyts, Uniform and safe metaclass composition, *J. Comput. Lang. Syst. Struct.* 31 (3–4) (December 2005) 143–164.
- [16] Stéphane Ducasse, Roel Wuyts, Alexandre Bergel, Oscar Nierstrasz, User-changeable visibility: resolving unanticipated name clashes in traits, in: *Proceedings of 22nd International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'07)*, ACM Press, New York, NY, USA, October 2007, pp. 171–190.
- [17] Ira R. Forman, Scott Danforth, *Putting Metaclasses to Work: A New Dimension in Object-Oriented Programming*, Addison-Wesley, 1999.
- [18] Matthew Flatt, Robert Bruce Findler, Matthias Felleisen, Scheme with classes, mixins, and traits, in: Naoki Kobayashi (Ed.), *Programming Languages and Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 270–289.
- [19] David Flanagan, *JavaScript: The Definitive Guide*, second edition, O'Reilly & Associates, January 1997.

¹³ <https://ci.inria.fr/pharo-ci-jenkins2/> and <https://github.com/pharo-project/pharo/>.

- [20] Kathleen Fisher, John Reppy, A typed calculus of traits, in: Proceedings of the 11th Workshop on Foundations of Object-Oriented Programming, 2004.
- [21] Adele Goldberg, Dave Robson, Smalltalk-80: The Language, Addison Wesley, 1989.
- [22] Urs Hölzle, Bay-Wei Chang, Craig Chambers, David Ungar, The SELF Manual, Computer Systems Laboratory of Stanford University, 1991.
- [23] König Dierk, Groovy in Action, Manning, 2007.
- [24] Gregor Kiczales, Jim des Rivières, Daniel G. Bobrow, The Art of the Metaobject Protocol, MIT Press, 1991.
- [25] Adrian Lienhard, Bootstrapping Traits, Master's thesis, University of Bern, November 2004.
- [26] Josh Lockhart, Modern PHP: New Features and Good Practices, O'Reilly Media, Inc., 2015.
- [27] Mark Lutz, Programming Python, 2nd edition, O'Reilly & Associates, Inc., 2001.
- [28] Donna Malayeri, Jonathan Aldrich, Cz: Multiple Inheritance Without Diamonds, ACM SIGPLAN Notices, vol. 44, ACM, 2009, pp. 21–40.
- [29] Nicholas D. Matsakis, Felix S. Klock II, The Rust Language, ACM SIGAda Ada Letters, vol. 34, ACM, 2014, pp. 103–104.
- [30] Marco Naddeo, A Modular Approach to Object Initialization for Pharo, Theses, Dipartimento di Informatica, Università degli Studi di Torino; Inria Lille Nord Europe - Laboratoire CRISTAL Université de Lille, November 2017.
- [31] Oscar Nierstrasz, Stéphane Ducasse, Tudor Gîrba, The story of Moose: an agile reengineering environment, in: Michel Wermelinger, Harald Gall (Eds.), Proceedings of the European Software Engineering Conference, ESEC/FSE'05, ACM Press, New York NY, 2005, pp. 1–10. Invited paper.
- [32] Oscar Marius Nierstrasz, Stéphane Ducasse, Stefan Reichhart, Nathanael Schärli, Adding Traits to (Statically Typed) Languages, 2005.
- [33] Oscar Nierstrasz, Stéphane Ducasse, Nathanael Schärli, Flattening traits, J. Object Technol. 5 (4) (May 2006) 129–148.
- [34] Martin Odersky, Scala language specification v 2.4, Technical report, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland, March 2007.
- [35] Guillermo Polito, Stéphane Ducasse, Luc Fabresse, Noury Bouraqadi, Benjamin van Ryseghem, Bootstrapping reflective systems: the case of pharo, Sci. Comput. Program. (2014).
- [36] Guillermo Polito, Stéphane Ducasse, Luc Fabresse, Camille Teruel, Scoped extension methods in dynamically-typed languages, The Art, Science, and Engineering of Programming 2 (1) (August 2017).
- [37] Guillermo Polito, Virtualization Support for Application Runtime Specialization and Extension, PhD thesis, University Lille 1, Sciences et Technologies, France, apr 2015.
- [38] Jorge Ressaia, Alexandre Bergel, Oscar Nierstrasz, Object-centric debugging, in: Proceeding of the 34th International Conference on Software Engineering, ICSE'12, 2012.
- [39] Reza Razavi, Noury Bouraqadi, Joseph Yoder, Jean-François Perrot, Ralph Johnson, Language support for adaptive object-models using metaclasses, Comput. Lang. Syst. Struct. 31 (3) (2005) 199–218.
- [40] Allison Randal, Dan Sugalski, Leopold Totsch, Leopold Tötsch, Perl 6 and Parrot Essentials, O'Reilly Media, Inc., 2004.
- [41] John Reppy, Aaron Turon, Metaprogramming with traits, in: European Conference on Object-Oriented Programming, Springer, 2007, pp. 373–398.
- [42] Markku Sakkinen, Disciplined inheritance, in: S. Cook (Ed.), Proceedings ECOOP'89, Cambridge University Press, Nottingham, July 1989, pp. 39–56.
- [43] Michel Schinz, Compiling scala for the Java virtual machine, EPFL Supercomput. Rev. (2005).
- [44] Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, Andrew P. Black, Traits, Composable units of behavior, in: Proceedings of European Conference on Object-Oriented Programming, in: LNCS, vol. 2743, Springer Verlag, July 2003, pp. 248–274.
- [45] Ghan Bir Singh, Single versus multiple inheritance in object oriented programming, SIGPLAN OOPS Mess. 5 (1) (January 1994) 34–43.
- [46] Pablo Tesone, Guillermo Polito, Luc Fabresse, Noury Bouraqadi, Stéphane Ducasse, Implementing modular class-based reuse mechanisms on top of a single inheritance VM, in: SAC 2018: SAC 2018: Symposium on Applied Computing, Pau, France, April 2018.
- [47] Tom van Cutsem, Mark Miller, Traits.js – robust object composition and high-integrity objects for ecmaScript 5, in: Proceedings of PLASTIC'11, 2011.