



HAL
open science

Uniform and safe metaclass composition

Stéphane Ducasse, Nathanael Schärli, Roel Wuyts

► **To cite this version:**

Stéphane Ducasse, Nathanael Schärli, Roel Wuyts. Uniform and safe metaclass composition. *Computer Languages, Systems and Structures*, 2004, 31, pp.143 - 164. <10.1016/j.cl.2004.11.003>. <hal-05022588>

HAL Id: hal-05022588

<https://hal.science/hal-05022588v1>

Submitted on 8 Apr 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Computer Languages, Systems & Structures 31 (2005) 143–164

COMPUTER
LANGUAGES,
SYSTEMS &
STRUCTURES

www.elsevier.com/locate/cl

Uniform and safe metaclass composition[☆]

Stéphane Ducasse^{a,*}, Nathanael Schärli^a, Roel Wuyts^b

^aSoftware Composition Group, IAM-Universität Bern, Neubruckstr. 10, 2012 Bern, CH, Switzerland

^bDecomp Laboratory, Université Libre de Bruxelles, Belgium

Accepted 2 November 2004

Abstract

In pure object-oriented languages, classes are objects, instances of other classes called *metaclasses*. In the same way as classes define the properties of their instances, metaclasses define the properties of classes. It is therefore very natural to wish to reuse class properties, utilizing them amongst several classes. However this introduced *metaclass composition problems*, i.e., code fragments applied to one class may break when used on another class due to the inheritance relationship between their respective metaclasses.

Numerous approaches have tried to solve metaclass composition problems, but they always resort to an *ad-hoc* manner of handling conflicting properties, alienating the meta-programmer. We propose a *uniform* approach that represents class properties as *traits*, groups of methods that act as a unit of reuse from which classes are composed. Like all the other classes in the system, metaclasses are composed out of traits. This solution supports the reuse of class properties, and their *safe* and *automatic* composition based on *explicit* conflict resolution. The paper discusses traits and our solution, shows concrete examples implemented in the Smalltalk environment *Squeak*, and compares our approach with existing models for composing class properties.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Metaclass composition; Traits; Reflective kernel; Reuse; Mixins

[☆] We gratefully acknowledge the financial support of the Swiss National Science Foundation for the projects “Tools and Techniques for Decomposing and Composing Software” (SNF Project No. 2000-067855.02) and “Recast: Evolution of Object-Oriented Applications” (SNF 2000-061655.00/1).

* Corresponding author. Tel.: +41 31 6314903; fax: +41 31 6313355.

E-mail address: ducasse@iam.unibe.ch (S. Ducasse).

1. Reusing class properties

In class-based object-oriented programming, classes are used as instance generators and to implement the behavior of objects. In object-oriented languages such as CLOS, Smalltalk or Ruby, classes themselves are first-class objects, and instances of the so-called *metaclasses* [1–5]. In the same way that classes define the properties for their instances (objects), metaclasses implement the properties for their instances (classes). Examples of class properties are *Singleton*, *Final*, *Abstractness* ... [6].

Treating classes as first-class objects and having metaclasses is important for two main reasons:

- *Uniformity and control*. In a pure object-oriented language it is natural for classes to be instances of metaclasses. The uniformity defines metaclasses as the natural place to *specify* and *control* object creation and other class behavior.
- *Reuse of class behavior*. Since a metaclass is just like any other class, class behavior is reused and conventional reuse and decomposition techniques are applied to the metaclasses [6]. Hence the same techniques that are available for base classes (inheritance and overriding of methods, for example) are applicable at the meta level.

When a language has metaclasses, those metaclasses can be *implicit* or *explicit*. With *implicit* metaclasses the programmer cannot specify the metaclass for a class [7]. As such, implicit metaclasses successfully address the goal of “uniformity and control”, but they fall short for achieving “reuse of class behavior”. *Explicit metaclasses* avoid this limitation because the programmer can explicitly state from which metaclass his or her classes are instances [1–4].

Languages without explicit metaclasses suffer from the fact that class properties cannot be reused across classes, and that they cannot be combined. For example, every time one needs a class with the Singleton behavior, the same code needs to be implemented over and over again. With explicit metaclasses the singleton class property can be factored out to a Singleton metaclass, which can then be used to instantiate classes that exhibit the Singleton behavior. However languages with explicit metaclasses suffer from the fact that composition can be *unsafe* [2,8] or are based on *non-uniform* mechanisms, i.e., the meta-programmer cannot use the same composition mechanism used for programming at the base level than for programming at the meta level. This is clearly a problem, since metaclasses originate from the wish of uniformity in OOP (see Section 2).

To address these problems we propose to use the *general-purpose* object-oriented language feature *traits* [9]. Traits are composable units of behavior that close the large conceptual gap between a single method and a complete class. Our approach models class properties with traits, and uses trait composition to safely combine and reuse properties in metaclasses. Consequently, metaclass composition (like class composition) enjoys all the conceptual benefits of the traits composition model. In particular, composition conflicts that occur when composing two properties that do not quite fit together are detected automatically and the conflict resolution is *explicit* and under *control* of the composing entity.

As we will show in the rest of the paper, our solution supports the reuse of class properties, their safe and *automatic* composition with *explicit* conflict resolution, and the usage of the same mechanism (traits) for both the base and metalevel. As safety is a broad term we follow the definition of safe metaclass composition as defined in [10] and that we present in the following section. Now we start by identifying precisely.

2. Explicit metaclass problems

Having explicit metaclasses promotes reuse but introduces several problems summarized in this section and detailed in the rest of the paper.

Unsafe composition. Some approaches sacrifice the compatibility between the class and the metaclass level [2,8]. Unsafe metaclass composition means that code fragments applied to a class may break when used on another one due to the inheritance relationship between the metaclasses of the classes involved (See Section 3).

Ad-hoc and non-uniformity. There are some approaches that are specifically designed to avoid the compatibility problems raised in the first point. Their solutions, however, rely on *ad-hoc* composition mechanisms that are based on automatic code generation and dynamically changing the meta-metaclass [10]. Not only does this make it hard to understand the resulting code, it also leads to problems in case of conflicting properties and results in hierarchies that are fragile with respect to changes. Note that MetaclassTalk by using mixin composition at the metalevel is the only solution that solves this problem [11].

The solutions are not satisfactory from a conceptual point of view either, because the meta level (or meta meta level) does not employ object-oriented techniques (such as inheritance or instantiation) but *ad-hoc* mechanisms only applicable for metaclass composition. This breaks the fundamental idea of reflective programming that uses the *available* features of a language to define and control the behavior of the language itself [3].

Limited composition. Other approaches used in the specific context of metaobjects use *chain of responsibility* [12] or composite metaobjects [13] to compose metaobjects. The first approach does not provide full control over the composition. The second approach forces the programmer to develop specific metaobjects to compose others, even when the reuse of these composite metaobjects is unclear.

The ideal metaclass composition solution would make composition be *automatic*. However, as we will discuss in Section 11 a simple solution does not exist in a context where new properties can be defined and composed, and where their semantics can severely conflict. So the solution is a mechanism that is both safe and uniform, i.e., one that does not require the developer to make a paradigm shift and where the development of base-level applications and meta-level applications is the same.

3. Qualifying composition

Offering explicit metaclasses is a way to reuse class properties but it also opens the door for *metaclass compatibility problems* [8]. This section defines criteria by which approaches that solve metaclass composition problems can be characterized and distinguished. We start by listing two criteria that were already identified in [10] (*upward*, *downward* compatibility and *per class property*), and then introduce three new ones that were not previously considered (*property composition*, *property application*, and *control of the composition*), but that qualify the problem in a more detailed way.

Upward compatibility. The fact that classes are instances of other classes which define their behavior introduces hidden dependencies in the inheritance relationships between the classes and their metaclasses. Careless inheritance at one level (be it the class or metaclass level), can break inter-level communication. Bouraqadi et al. [10] refined the metaclass compatibility problems in two precise cases named *upward* and *downward* compatibility.




Fig. 1. Left: Upward compatibility—dependencies on the base level need to be addressed at the meta level. Right: Downward compatibility—dependencies on the meta level need to be addressed at the base level.

Let B be a subclass of A , $MetaB$ the metaclass of B , and $MetaA$ the metaclass of A . Upward compatibility is ensured for $MetaB$ and $MetaA$ iff: every possible message that does not lead to an error for any instance of A , will not lead to an error for any instance of B .

Fig. 1 left illustrates upward compatibility. When an instance of B receives the message `i-foo`, the message `c-bar` is sent to B . The composition of A and B is upward compatible, if B understands the message `c-bar`, i.e., $MetaB$ should implement it or somehow inherit it from $MetaA$.

Downward compatibility. Let $MetaB$ be a subclass of the metaclass $MetaA$. Downward compatibility is ensured for two classes B , instance of $MetaB$ and A , instance of $MetaA$ iff: every possible message that does not lead to an error for A , will not lead to an error for B .

Downward compatibility is illustrated in Fig. 1 right. When B receives the message `c-foo`, the message `i-bar` is sent to a newly created instance of B . The composition of $MetaA$ and $MetaB$ is downward compatible, if that new instance of B understands the message `i-bar`, i.e., B should implement it or somehow inherit it from A .

Definition. Metaclass composition is *safe* when it supports downward and upward compatibility.

Per class property. Different metaclass properties can be assigned to different classes in an inheritance hierarchy. Some systems such as NeoClasstalk and MetaClasstalk allow one to assign a property to a *single* class without it being inherited by its subclasses [10,11]. The authors of NeoClasstalk and MetaClasstalk, Bouraqadi et al. defined *class property propagation* as follows: “A property assigned to a class is automatically propagated to its subclass”. We name this criteria *per class property*. For example, it is possible to define that a class is abstract and its subclasses are not abstract and this without having to redefine the property at the subclasses level.

Property composition. One of the main motivations for having explicit metaclasses is to *combine class properties*, as shown in Fig. 2, so that one class can, for example, be both a Singleton and Final. Hence a mechanism is needed that supports such property composition. This can be a general-purpose language mechanism such as multiple inheritance [3,4], mixin composition [11], chain of responsibility [12], or an *ad-hoc* mechanism such as generation of new classes and methods [10].

Property application. Property application is the mechanism by which the composed properties are applied to classes. As shown in Fig. 2 we distinguish the *composition* of properties from the *application* of a property to a specific class because some approaches employ different techniques for these two purposes. As an example, SOM uses ordinary multiple inheritance to compose class properties

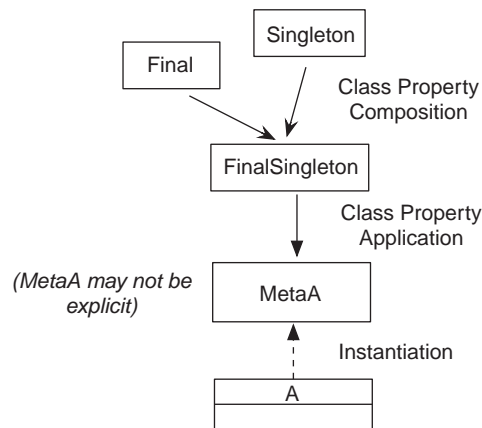


Fig. 2. Property Composition and Property Application: two different stages in the process of reusing class properties.

but it employs a combination of multiple inheritance and code generation to apply a class property to a class.

Control. The mechanism used to apply and combine class properties can be *implicit* or *explicit*. We call the mechanism *implicit* if the system automatically combines or applies the class properties and implicitly resolves conflicts in a way that may or may not be what the programmer intends. We call the mechanism *explicit* if the system gives the programmer explicit control over how the properties are combined and applied. In particular, the programmer should have *explicit control* over how conflicts are resolved. For many approaches, this is not the case because the composition of properties is based on a chain of responsibility which does not provide full control of the composition.

4. Traits in a Nutshell

Traits [9] are an extension of single inheritance with a similar purpose as mixins but avoiding their problems. Traits are essentially groups of methods that serve as building blocks for classes and are primitive units of code reuse. As such, they allow one to factor out common behavior and form an intermediate level of abstraction between single methods and complete classes. A trait consists of *provided methods* that implement its behavior, and of *required methods* that parameterize the provided behavior. Traits cannot specify any instance variables, and the methods provided by traits never directly access instance variables. Instead, required methods can be mapped to state when the trait is used by a class.

With traits, the behavior of a class is specified as the composition of traits and some *glue methods* that are implemented at the level of the class. These glue methods connect the traits together and can serve as accessors for the necessary state. The semantics of such a class is defined by the following three rules:

- *Class methods take precedence over trait methods.* This allows the glue methods defined in the class to override equally named methods provided by the traits.
- *Flattening property.* A non-overridden method in a trait has the same semantics as the same method implemented in the class.

- *Composition order is irrelevant.* All the traits have the same precedence, and hence conflicting trait methods must be explicitly disambiguated.

Because the composition order is irrelevant, a *conflict* arises if we combine two or more traits that provide identically named methods that do not originate from the same trait. Traits enforce explicit resolution of conflicts by implementing a glue method at the level of the class that overrides the conflicting methods, or by *method exclusion*, which allows one to exclude the conflicting method from all but one trait. In addition traits allow *method aliasing*. The programmer can introduce an additional name for a method provided by a trait to obtain access to a method that would otherwise be unreachable, for example, because it has been overridden. Traits can be composed from subtraits. The composition semantics is the same as explained above with the only difference being that the composite trait plays the role of the class.

5. Using traits to reuse and compose class properties

Our approach is based on using traits to compose and reuse class properties within the traditional parallel inheritance schema proposed by Smalltalk (see Fig. 8 left). Therefore our approach is safe, i.e., it supports downward and upward compatibility. But on top of that it promotes the reuse of class properties. Composition and application of class properties are based on trait composition, which gives the programmer explicit control in a uniform manner.

We represent class properties as traits, which are then used to compose metaclasses as shown in Fig. 3. Since traits have been fully implemented in the open-source Squeak Smalltalk environment [14], we implemented all the examples shown here in Squeak. During our refactoring of Squeak code we identified the following class properties: TAbstract, TSingleton, TRememberInstances, TCreator, and TFinal which we explain below. We start with a simple example illustrating how a class is composed by reusing a class property, then we look how the traditional Boolean hierarchy [6,10] is re-expressed with traits and finally Section 6 shows that traits provide a good basis to engineer the meta level.

5.1. Singleton

To represent the fact that a class is a Singleton, we define the trait TSingleton. This trait defines the following methods: `default` which returns the default instance, `new` which raises an error, and `reset` which invalidates the current singleton instance. It requires `basicNew` which returns a newly created instance,¹ and the methods `uniqueInstance` and `uniqueInstance:`. Note that these accessor methods are needed because traits cannot contain instance variables. Fig. 4, left, shows the trait TSingleton.

```
Trait named: #TSingleton uses: {} category: 'Traits-Example'
TSingleton >>default
  self uniqueInstance isNil
    ifTrue: [self uniqueInstance: self basicNew].
↑ self uniqueInstance
```

¹ Using `basicNew` is the traditional way to implement Singleton in Smalltalk when we want to forbid the use of the new method [15]. `basicNew` allocates objects without initializing them. It is a Smalltalk idiom to never override methods starting with 'basic' names.

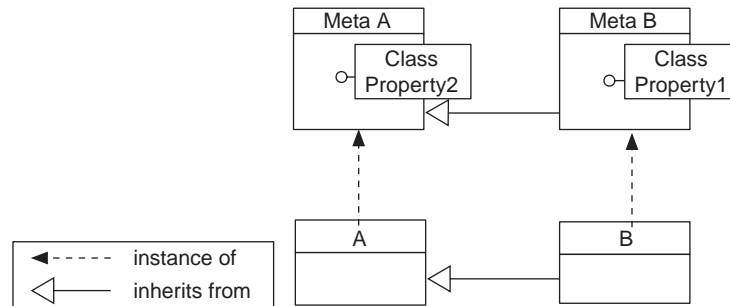


Fig. 3. Metaclasses are composed from traits representing class properties. Traits supports upward and downward compatibility.

```

TSingleton>>new
  self error: 'You should use default'
TSingleton>>reset
  self uniqueInstance: nil
  
```

As an example, suppose that we want to specify that a certain class `WebServer` is a Singleton. First of all we define the class `WebServer` in the traditional Smalltalk way as shown in Fig. 4. Then we specify at the metaclass level, i.e., in the class `WebServer` class, that the class is a Singleton by specifying that the class is composed from the trait `TSingleton`. The metaclass defines an instance variable `uniqueInstance` to maintain the singleton instance. It also defines two glue methods `uniqueInstance` and `uniqueInstance:` as accessor methods for the instance variable `uniqueInstance`. These two glue methods fulfill the required methods with the same name of the trait `TSingleton`. Note that the required method `basicNew` is provided by the class `Behavior`, of which `WebServer` class, is an indirect subclass (see Fig. 4, right).

5.2. The Boolean hierarchy

The Smalltalk Boolean hierarchy consists of the abstract class `Boolean`, that has two subclasses `True` and `False` that are singleton classes. Traits allow the boolean hierarchy to be refactored as shown in Fig. 5. Note that the refactored solution is backwards compatible with the idioms existing in the current Smalltalk implementation and literature [15]. So we assume that a method `basicNew` is defined on the class `Behavior` that can always be invoked to allocate instances and that should not be overridden.

Boolean. The class `Boolean` is an abstract class, so we compose its class `Boolean` class from the trait `TAbstract`.

```

Trait named: #TAbstract uses: {} category: 'Traits-Example'
TAbstract>>new
  self error: 'Abstract class. You cannot create instances'
TAbstract>>new: size
  self error: 'Abstract class. You cannot create instances'
  
```

False and True. The classes `False` and `True` are Singletons so their classes `False` class and `True` class are composed from the trait `TSingleton` which is then reused in the two classes.

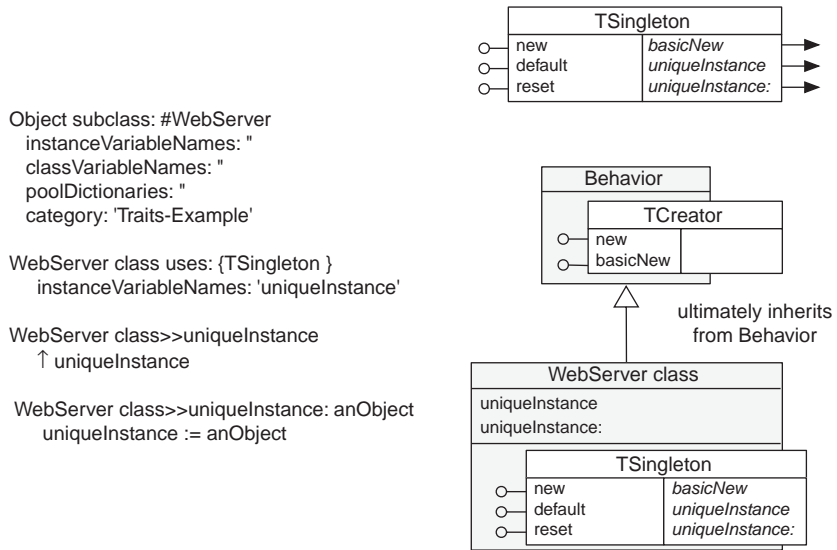


Fig. 4. Left. The trait TSingleton. Right. The class Behavior, the root of metaclasses in Smalltalk, is composed from the trait TCreator and as such provides the method basicNew.

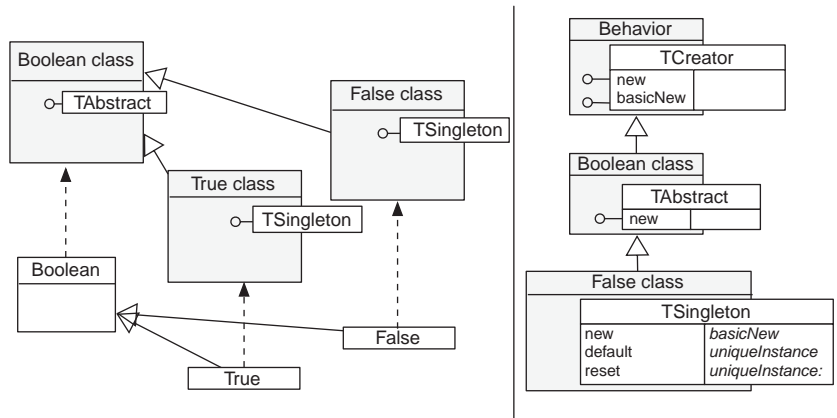


Fig. 5. Left: Boolean hierarchy refactored with traits. Right: The complete picture for the Boolean hierarchy solution.

As mentioned above, the trait TSingleton requires the methods basicNew, uniqueInstance, and uniqueInstance:. Therefore the class False class (resp. True class) has to define an instance variable uniqueInstance and the two associate accessor methods uniqueInstance and uniqueInstance:. Note that the method basicNew does not have to be redefined locally in the class False or True class as it is inherited ultimately from the class Behavior, the inheritance root of the metaclasses [7] (see Fig. 5 right). This example shows that class properties are reused over different classes and that metaclasses are composed from different properties.

```
False class
  uses:{TSingleton}
  instanceVariableNames: 'uniqueInstance'
```

```
False class>>uniqueInstance
  ↑ uniqueInstance
False class>>uniqueInstance: anObject
  uniqueInstance := anObject
```

6. Engineering the meta level

So far we presented simple examples that show how traits are well-suited to model class properties, which can then be combined or applied to arbitrary classes. In this section, we show that traits also allow more fine-grained architectures of class properties. We also want to stress that the techniques used here at the meta level are exactly the same as those used at the base level. As such, traits provide a uniform model.

Since many of these properties are related to instance creation, and we perform our experiments in Smalltalk, we first clarify the basic instance creation concept of Smalltalk. In Smalltalk, creation of a new instance involves two different methods, namely `basicNew` and `new`.² The method `basicNew` is a low-level primitive which simply allocates a new instance of the receiver class. The method `new` stands at a conceptually higher level and its purpose is to return a usable instance of the receiver class. For most classes, `new` therefore calls `basicNew` to obtain a new instance and then initializes it with reasonable default values.

6.1. Class properties

Fig. 6 gives an overview of the class properties we identified (see Section 10 for a deeper discussion). Note that all of these properties are traits, and that they are therefore composed using trait composition.

Allocation. As indicated by its name, the trait `TAllocator` provides the behavior to allocate new instances. In our case, this is the standard Smalltalk `basicNew` method, but of course we could also create another trait with an alternative allocation strategy.

Instantiation. The traits `TInstantiator` and `TInitInstantiator` are two class properties for instance creation. The trait `TInstantiator` uses the trait `TAllocator` and implements the method `new` in the traditional Smalltalk manner, which means that it does not initialize the newly created instance. The trait `TInitInstantiator` uses the trait `TAllocator`. However, as suggested by its name, it actually initializes the newly created instance by calling the method `initialize` before the instance is returned.

```
TInstantiator>>new
  ↑self basicNew
TInitInstantiator>>new
  ↑self basicNew initialize
```

Note that the method `initialize` is called on the new instance, which means that the requirement for `initialize` in the trait `TInitInstantiator` is actually a requirement for the instance side.

Remembering instances. The trait `TRememberInstances` represents an instance creation property that remembers all the instances created by a class. It uses the trait `TInitInstantiator` and aliases the

² Note that there are also the methods `basicNew:` and `new:`, which are used to create objects with indexed fields (i.e., arrays). For sake of simplicity, we do not take these methods into account here.

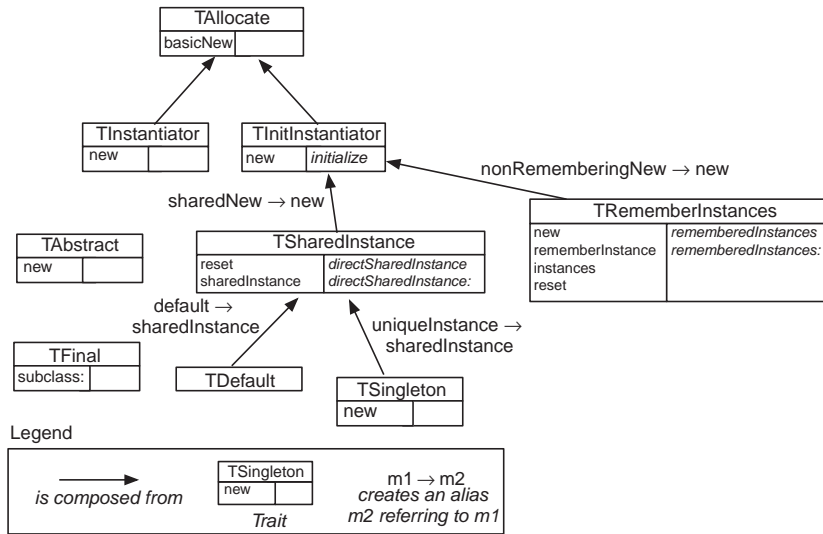


Fig. 6. A fine-grained architecture of class properties based on traits.

method `new` of the traits **TInitInstantiator** which is then available as `nonRememberingNew`. This aliasing allows one to access the original `new` method of the trait **TInitInstantiator** while leaving the option to override the method `new` in the trait **TInitInstantiator**. It requires the methods `rememberedInstances` and `rememberedInstances:` to access a collection storing the created instances. Then, it implements the methods `new`, `rememberInstance:`, `instances`, and `reset` as follows:

```

TRememberInstances>>new
  ↑ self rememberInstance: self nonRememberingNew
TRememberInstances>>rememberInstance: anObject
  ↑ self instances add: anObject
TRememberInstances>>instances
  self rememberedInstances ifNil: [self reset].
  ↑ self rememberedInstances
TRememberInstances>>reset
  self rememberedInstances: IdentitySet new
    
```

Note that another implementation could be to define the methods `reset` and `rememberedInstances:` as trait requirements. This would leave the class with the option to use other implementations for keeping track of the created instances.

Default and Singleton. The traits **TDefault** and **TSingleton** implement the class properties corresponding to the Default Instance and Singleton design patterns. Whereas a Singleton can only have one single instance, a class adhering to the Default Instance pattern has one default instance but can also have an arbitrary number of other instances.

Since these two properties are very similar, we factored out the common code into the trait **TSharedInstance**. To get the basic instantiation behavior, this trait uses the property **TInitInstantiator** and again applies an alias to ensure that the method `new` is available under the name `sharedNew`. Then, it implements

the methods `reset` and `sharedInstance` as follows:

```
TSharedInstance>>reset
  self directSharedInstance: self sharedNew.
TSharedInstance>>sharedInstance
  self directSharedInstance ifNil: [self reset].
  ↑ self directSharedInstance.
```

The property `TDefault` is then defined as an extension of the trait `TSharedInstance` that simply introduces the alias default for the method `sharedInstance`. Similarly, the property `TSingleton` introduces the alias `uniqueInstance` for the same method. In addition, `TSingleton` overrides the method `new` so that it cannot be used to create a new instance:

```
TSingleton>>new
  self error: 'Cannot create new instances of a Singleton.
  Use uniqueInstance instead'.
```

Another useful class property popularized by Java is the class property `TFinal` which ensures that a class cannot have subclasses. In Smalltalk, this is achieved by overriding the message `subclass:`.³ Note that unlike all the other properties presented in this section, `TFinal` is not concerned with instance creation and therefore is entirely independent of the other properties. In Section 10 we discuss the relevance of the class properties we presented.

6.2. Advantages for the programmer

Having an architecture of class properties has many advantages for a programmer. Whenever a new class needs to be created, a choice can be made regarding the creation of instances, and whether or not the class should be final. Besides having the obvious advantage of avoiding code duplication, it also makes the design much more explicit and therefore facilitates understandability of the class. The level of abstraction of the trait design is at the right level: the traits correspond to the class properties, and the class properties can be combined into metaclasses.

In addition, factoring out the properties in such a fine-grained way still gives the user a lot of control about some crucial parts of the system. Suppose for example that at first we would have decided to use the trait `TInitInstantiator` as the basis for all the other instance creation properties. If later on, we would decide to comply to the Smalltalk standard to create uninitialized instances by default, then we could make this change without modifying any of the involved methods. We would just need to make sure that the traits `TRememberInstances` and `TSharedInstance` use the trait `TInstantiator` instead of `TInitInstantiator`.

Explicit composition control power. By providing several different properties that are all related to instance creation behavior, this example also shows why it is so important to have explicit control over composition and application of class properties. In our example, there are many different properties which essentially introduce variants of the method `new`, and therefore, combining these properties typically leads to conflicts that can only be resolved in a *semantically* correct manner if the user has explicit control over the composition. In case of traits, this is ensured by allowing partially ordered compositions, exclusions, and aliases.

³ In reality, the method to create a subclass takes more arguments but this is not relevant here.

As an example, imagine that we want to combine the properties `TDefault` and `TRememberInstances` to get a property that allows both a default instance and also remembers all its instances. With our trait-based approach, we do this by creating a new trait `TDefaultAndRememberInstances` which uses `TRememberInstances` and `TDefault` as follows:

```
Trait named: TDefaultAndRememberInstances
  uses: {TDefault @ {# defaultReset → # reset}.
        TRememberInstances – {#new}
        @ {# storeNew→#new.
          #storeReset→# reset}}
```

```
TDefaultAndRememberInstances>>sharedNew
  ↑self storeNew
```

```
TDefaultAndRememberInstances>>reset
  self storeReset.
  self defaultReset
```

Since both traits provide a method `new`, we exclude this method from the trait `TRememberInstances` when it is composed. As a consequence the trait contains the `new` method provided by `TDefault`, which uses `sharedNew` to create a new instance. Since we want to make sure that each new instance is also stored, we override `sharedNew` so that it calls `storeNew`, which is an alias for the `new` method provided by `TRememberInstances`.

Because the method `reset` is also provided by both traits, we use aliasing to make sure that we can access the conflicting methods. Then, we resolve the conflict by overriding the method `reset` so that it first removes the stored instances (by calling `storeReset`) and then creates a new default instance (by calling `defaultReset`). Note that the newly created instance will be remembered as the default instance and will also be stored in the collection with all the instances of the class.

7. A new kernel

In this section we present briefly the key implementation aspects of the new Smalltalk kernel that is bootstrapped with traits. Fig. 7 shows the new class hierarchy that we obtain. As the new kernel with traits is an extension of the traditional Smalltalk kernel, we get the traditional classes: `Behavior`, `ClassDescription`, `Metaclass`, and `Class` which now deal with the fact that a class may be composed of traits.

To model traits we then follow the previous design of the kernel and mimic the classes `Behavior`, `ClassDescription` and `Class`. Three classes `TraitBehavior`, `TraitDescription`, and `Trait` are introduced. `Trait` represents a trait and is applied to both the class and instance side. In addition the class `Behavior`, root of the instantiation graph, uses two important traits: `TInstantiator` and `TInitInstantiator` as presented in previous section. The class `Behavior` in Smalltalk also defines information about the state and behavior related to superclass and instance variables (`format`). As this is not needed for traits, we introduce a new abstract superclass `PureBehavior`, which factors out the common code between `Behavior` and `TraitBehavior`.

Some traits. The traits we identified and used in this new kernel are not really remarkable in the sense of new MOP entries. In fact we mainly use traits to reuse code between the classes `TraitDescription` and

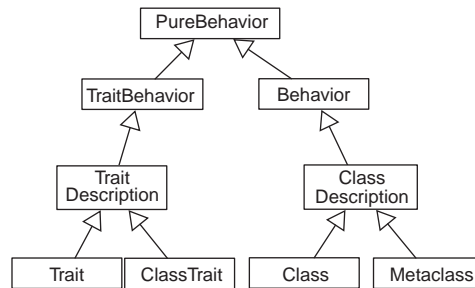


Fig. 7. Hierarchy of the new Smalltalk kernel supporting traits.

ClassDescription for the following reasons.

- The new kernel is based on the traditional Smalltalk kernel, which uses inheritance as the primary reuse mechanism. A lot of polymorphic methods are used among the classes Behavior and ClassDescription, and the classes Class and Metaclass. As a consequence, there is not much need for introducing traits to share this functionality.
- In the traditional Smalltalk kernel, the class Behavior only defines the minimal state and behavior to support classes as run-time entities. For example, Behavior does not define the notion of named instance variables but just knows the format of the instances the class will create in terms of the number and kind of instance variables [7]. Support for named instance variables is (amongst other things) implemented by ClassDescription, a subclass of the class Behavior. Because of this inheritance structure, some code that is implemented in ClassDescription (and uses state defined in ClassDescription) cannot and should not be reused by pushing it up to PureBehavior. Instead, we use traits to share behavior between ClassDescription and TraitDescription.

The class PureBehavior uses the following traits:

- TBasicCompile supports the compilation and decompilation of methods in a class.
- TTestingSelector supports the testing of selectors of methods (e.g., canUnderstand:, ...).
- TCompiledMethodAccess supports the access into the method dictionary and source code access.
- TMethodIterating supports the iteration over compiled methods.

The class Behavior uses the following traits:

- TInstantiator and TInitInstantiator implement the creation of objects as described in the previous section.
- TFamilyAccess supports the access and enumeration of superclasses and subclasses.
- TInstanceEnumerator supports the enumeration of instances of the class.
- TMethodTesting supports the querying of methods.

The traits resulting from the decomposition of the classes PureBehavior and Behavior are not currently used by any other classes. In contrast, the following traits are used to share behavior between the two

Table 1

Comparison of the models from Section 8 on how they handle the composition problems described in Section 2

	Up	Down	Per class	Application	Composition	Control
Smalltalk	Yes	Yes	No	No	No	No
CLOS	Yes	Yes	No	Multiple inheritance	Multiple inheritance	Explicit + linearization
SOM	Yes	No	No	Multiple inheritance + code generation	Multiple inheritance	Implicit
NeoClasstalk	Yes	Yes	Yes	Inheritance + generation	Inheritance + code generation	Implicit
MetaClassTalk	Yes	Yes	Yes	Inheritance	Mixin composition	Mixin linearization
Traits	Yes	Yes	No	Trait composition	Trait composition	Explicit

classes TraitDescription and ClassDescription.

- TClassComment supports the management of comments.
- TMethodDictionaryManagement supports the management of methods categories.
- TOrganization supports how methods are sorted into method categories.
- TCodeReformatting supports reformatting of source code.
- TCodeFileOut supports filing-out (saving) of classes.
- TBehaviorCopy supports the copying of methods and their organization.
- TOrganizedCompilation supports the compilation of methods within the context of categories.

It should be noted that in the Smalltalk metaclass kernel, identifying traits that can be reused independently of each others is difficult because the behavior of the kernel is based on inheritance and the code was tightly coupled.

8. Related work

This section shows how the main approaches that support explicit metaclasses address the problems described in Section 2. We also discuss the solution offered by Smalltalk (although it has implicit metaclasses) since it forms the basis for the NeoClasstalk solution and our own solution. Table 1 summarizes the comparison of these approaches. Note that the table shows the influence of the CLOS approach based on multiple inheritance to support metaclass composition in SOM.

8.1. Metaclass composition

8.1.1. Smalltalk

In Smalltalk (and more recently in Ruby), metaclasses are *implicit* and created *automatically* when a class is created [7]. Each class is the sole instance of its associated metaclass. This way the two hierarchies are parallel (see Fig. 8 left). Hence the architecture is safe as it addresses compatibility issues but completely prevents class property reuse between several hierarchies.

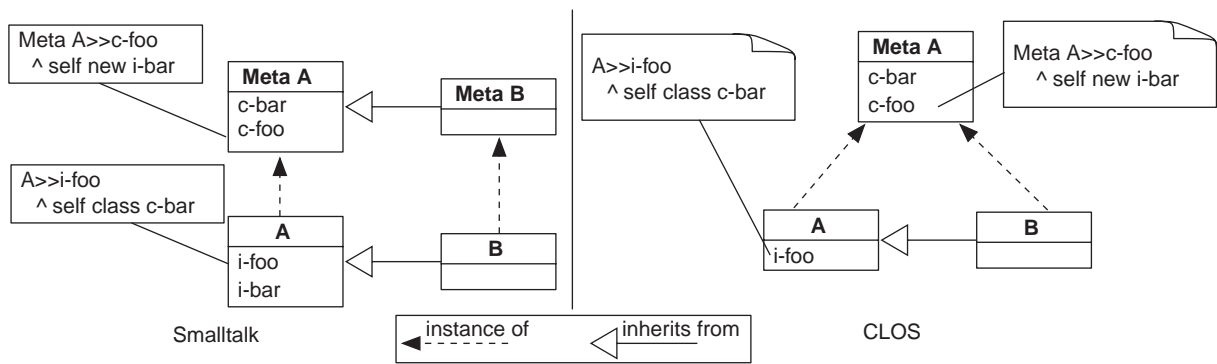


Fig. 8. Left: Smalltalk addresses compatibility issues by preventing reuse using implicit metaclasses and parallel hierarchies. Right: By default CLOS addresses compatibility issues by preventing subclasses to have different metaclasses than their superclasses.

8.1.2. CLOS

CLOS's approach could be summarized as “do it yourself”. Indeed by default in CLOS, a class and its subclasses must be instances of the same metaclass, prohibiting classes in the same hierarchy from having different class properties. For example, in Fig. 8 right, class B has by default the same metaclass as its superclass A, and this cannot be changed. So class B always has the same class properties as class A. Note that since CLOS has explicit metaclasses, multiple inheritance can be used for composing class properties. For example, in the context described by Fig. 2 it is possible to use multiple inheritance to explicitly combine the two properties Final and Singleton expressed as metaclasses into a new class SingletonFinal. Note that such an implementation suffers from the same problems as multiple inheritance based on linearization occurring at the base level [16].

The general CLOS rule that a class and its subclasses must be instances of the same metaclass can be circumvented using CLOS's metaobject protocol (MOP). Indeed, the generic function `validate-superclass` [3] offers a meta-programmer the possibility to specify that a class and its subclasses can be instances of different classes. However, this comes at a very high price because the CLOS MOP does not provide predefined strategies for avoiding compatibility problems or for dealing with possible conflicts. Hence the semantics of the composition has to be implemented manually, a far from trivial undertaking.

This means that by default CLOS is upward and downward compatible but it prevents usage of different metaclasses within an inheritance hierarchy and reuse of class properties. Both the composition of class properties and the application of properties are done with multiple inheritance. The control of the composition is explicit, because the user has to use multiple inheritance to create a new metaclass. However, since multiple inheritance in CLOS uses implicit linearization, the well-known problems associated with this form of conflict resolution also apply to the meta level [16].

8.1.3. SOM

The solution proposed by SOM (System Object Model) [5] is based on the automatic generation of *derived metaclasses*, that inherit multiply from the metaclasses to compose class properties. When at compile time a class is specified to be an instance of a certain metaclass, SOM automatically determines whether upward compatibility is ensured and if necessary creates a derived metaclass. In Fig. 9 left, the class B (originally an instance of MetaB), inheriting from class A (instance of MetaA) finally becomes

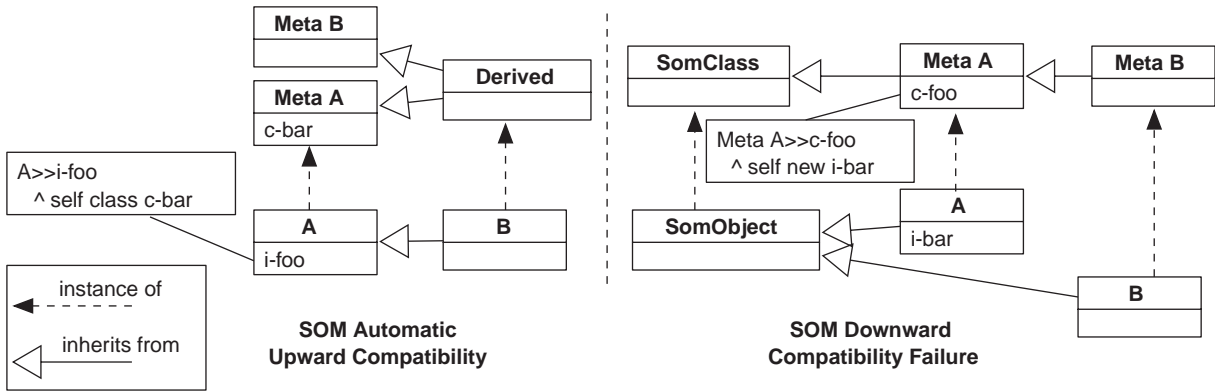


Fig. 9. Left: SOM supports upward compatibility by automatically deriving new metaclasses and changing the class of the inheriting class B. Right: SOM downward compatibility failure example.

an instance of a derived metaclass inheriting from MetaA and MetaB. Note that SOM ensures that the existing metaclass MetaB takes precedence over MetaA in case of multiple inheritance ambiguities (since B is a subclass of A).

While SOM supports upward compatibility as shown in Fig. 9 left, it does not support downward compatibility [10] as shown in Fig. 9 right. When the class B receives the c-foo message, a run-time error will occur because its instances do not understand the i-bar message. However, in SOM, contrary to CLOS, two distinct classes need not have the same metaclass. But as in CLOS, the composition of class properties is based on multiple inheritance. The application of a class property is done by a combination of multiple inheritance and automatic class generation. This happens at compile time, and the programmer has no explicit control over how possible conflicts are resolved.

8.1.4. NeoClasstalk

NeoClasstalk’s approach is interesting since it supports both downward and upward compatibility and enables class property reuse between different hierarchies [10,17,18]. NeoClasstalk uses two techniques to accomplish this: *dynamic change* of classes and the composition of metaclasses by *code generation*. It generalizes the parallel inheritance solution of Smalltalk by enabling class properties reuse, but it also introduces some problems of its own that we discuss in detail after explaining the basic principles.

NeoClasstalk allows properties to be assigned to classes. Fig. 10 shows what happens when assigning a property to Meta B. B inherits from class A and is an instance of the class Meta B before the new property is assigned to Meta B. When assigning the property, the system automatically creates a new metaclass Property m + Meta B (called a *property metaclass*), which inherits from the metaclass Meta B and defines the property code. It then *changes the class* of B to be that newly created metaclass. NeoClasstalk supports also per class property, i.e., a property added to a class does not get automatically propagated to its subclasses.

To be able to reuse the property classes, NeoClasstalk stores the class properties in strings on methods of the so-called *meta-metaclasses*. The actual metaclasses are then generated from these strings, as shown for our example in Fig. 10. For example, the Property m represented by a meta-metaclass is used to generate a new metaclass named Property m + Meta B from the metaclass Meta B and the Property m.

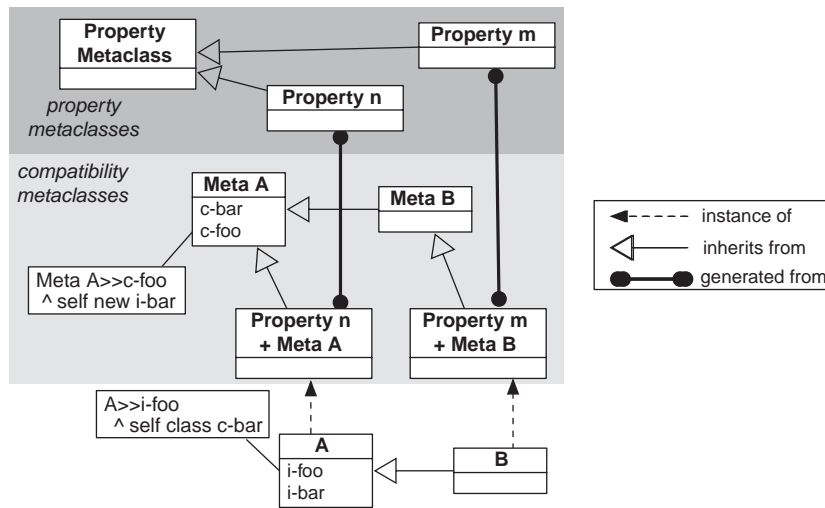


Fig. 10. Assigning the property *m* to class Meta B and property *n* to the class Meta A in NeoClasstalk. The light grey area denotes the metaclass area. The dark grey area is the realm of the class properties.

Besides the intrinsic complexity of NeoClasstalk's approach, it has the following drawbacks:

8.1.5. Dynamic class creation and dynamic change of class

The approach relies on the dynamic creation of classes and the dynamic changing of classes. It induces a complex management of meta-metaclass changes that should be propagated to the generated instances. Moreover as programming at the meta meta level is based on manipulating the strings that represent bodies of methods of metaclasses, it is not the same as programming at the metaclass or the base level. Basically, despite the name, the property metaclasses are not really at the meta-metaclass level, but merely storage holders for strings. The relation between the meta-metaclass level and the metaclass level is therefore not instantiation, as one would expect, but code generation. This breaks the uniformity of the model.

8.1.6. Ad-hoc and implicit composition

Property metaclasses are composed by code generation and applied implicitly by defining them in an inheritance chain. The composition is based on the assumption that a metaclass is designed to be plugged in this inheritance chain and that other composed behavior can be reached via super invocations. The composite metaclass has only limited control over the composed behavior as it can only invoke overridden behavior but does not have the full composition control.

As a summary, NeoClasstalk provides both downward and upward compatibility, and it allows one to assign class properties on a per-class basis. The composition of class properties is implicit and based on code generation and chain of responsibility. The application of class properties is based on dynamic class changes and code generation.

8.1.7. MetaclassTalk

MetaclassTalk follows the architecture of NeoClasstalk by offering compatibility and property metaclasses. MetaclassTalk uses mixin composition to compose metaclass properties [11]. This experiment

makes MetaclassTalk the closest model to our own approach as it supports both downward and upward compatibility while allowing the reuse of class properties. However in MetaclassTalk, property composition is based on mixin linearization. As such it has the same problems as the ones we present in [9]: the composite entities do not have the full control of the composition, and the glue code is spread over multiple classes. These problems are solved by traits.

8.2. Metaobjects

Other approaches such as CodA [19], Mostrap [12], Iguana/J [20] support the composition and reuse of metaobjects. Such a composition is often based on chain of responsibility [12], i.e., a metaobject is designed to be composed in a chain of metaobjects by invoking the overridden functionality. The problem with chain of responsibility is that it forces all the metaobjects to follow a certain architecture. It more importantly gives the composing metaobject only a very limited control over the composition: it can invoke the rest or do nothing. In contrast, traits composition is automatic when there is no conflict, and when conflicts arise, the composing metaclass has complete control over all the composed class properties.

The authors of Guarana [13] and Reflex [21], introduce *composite metaobjects*, i.e., a metaobject that define the composition semantics of several metobjects. This approach works well for coarse-grained composition, such as for making changes to the message passing semantics (broadcast, concurrent dispatch, or remote invocations). However, it is too heavyweight to compose class properties, since it would force the developer to define an explicit *composite metaclass* for all simple conflicts whose reuse is even questionable.

CodA [19] structures the meta-level architecture around several metaobjects responsible for the different actions. However it raises the issue of compatibility between all the metaobjects associated to a given object. The solution is to manually define a semantically coherent configuration of metaobjects implementing the desired semantics [19]. This solution shows again that there is no magic and that composing operations with conflicting semantics cannot be achieved in an automatic fashion.

9. Advantages and disadvantages

Advantages. Traits support the decomposition of class properties as reusable units of behavior. Since metaclasses are composed of traits and the model is based on the parallel hierarchy of Smalltalk, it is upward and downward compatible and supports the reuse of class properties across different hierarchies.

In addition the proposed model is uniform with respect to the concepts used at the base level and the meta level (like CLOS). Both levels use the same concepts (traits and inheritance). Furthermore, the model is simple, and there is no need for on-the-fly code generation (as in SOM or NeoClasstalk) or for dynamic changes to classes (as in NeoClasstalk).

Class properties can be composed of traits that represent those properties. The application of the properties to an actual metaclass is accomplished by using the appropriate composite trait in the metaclass definition. The composing metaclass has *complete* control of the composition, and possible conflicts are resolved *explicitly* when the property is applied to a metaclass.

Having explicit control over the composition is especially important because it allows a programmer to freely adapt the behavior of the composite metaclass and to compose class properties that may not quite fit together. This means that our approach lets system designers ship their class hierarchies together with

a set of *prefabricated* class properties in the form of traits, which can then be used and combined by the programmers. In case some class properties built by different vendors do not quite fit together, the traits model not only indicates the resulting conflicts, but also provides the programmer with the necessary means to resolve the conflicts to achieve the expected semantics.

Disadvantages. Glue methods and state have to be redefined in the metaclass where a property is applied. For example, the instance variable `uniqueInstance` and the two accessors methods have to be defined in all classes that implement a Singleton. We consider this to be a limit of the traits model and the price to pay to have the minimal mechanism supporting traits composition. Introducing state into traits would solve this but would introduce other problems, such as the well-known *diamond problem* of multiple inheritance [22], where state gets inherited through different paths.

It may happen that instance variables defined in a superclass are not necessary in the subclasses. For example, if the superclass implements a Singleton and the subclasses do not, then the instance variable that holds the Singleton instance as well as the methods to access it will be inherited by the subclasses. However, this problem is not due to traits by itself but is a result of using the inheritance mechanism in general. Table 1 compares the approaches.

10. Class properties

Traits let us decide if a given functionality is defined as a trait or as a class. When defining functionality as a trait we automatically offer the possibility for future classes to use the identified behavior. One might wonder why we have so few class properties. First of all we chose to reengineer the current implementation of Squeak and not to design a new MOP. In this article we present the main *class properties* that we identified during our implementation, and we did not invent new ones. Secondly, we deliberately took heavily conflicting class properties, so that we could clearly show the conflict resolution advantage of traits. Composing non-conflicting properties is trivial. Thirdly, other important efforts to build metaclass libraries, such as SOM [5], present nearly the same set of class properties.

Another point to consider is the role of the classes in the context of a MOP [3]; we believe that a lot of class properties identified in [6] are due to the fact that the classes were the single entry point in their MOP, while certain responsibilities are definitely the responsibilities of other meta-entities such as methods. It is also out the scope of this paper to present a new MOP based on traits, even if this is definitely future work.

Class property propagation. Our approach does not support per class property because we did not want to change the class creation protocol of Smalltalk-80. However, there is nothing in our approach that prevents us to support per class property in a similar way than the compatibility model does [10,11]. Fig. 11 shows that the intermediate metaclasses Boolean class + Abstract, False class + Singleton, and True class + Singleton are composed of traits and that properties such abstractness of the class Boolean, are not propagated to its subclasses. Note that metaclasses such as Boolean class could also be composed of traits if the property have to be propagated to the subclasses.

11. Automatic composition

An important difference between traits and most other approaches lies in the fact that traits have automatic conflict *detection* but expect *explicit* conflict resolution controlled by the programmer. Most approaches aim for a completely automatic resolution of conflicts, where possible conflicts are resolved

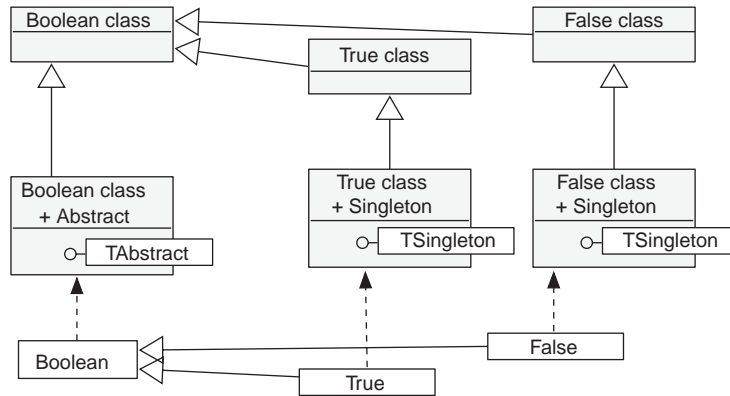


Fig. 11. Controlling class property propagation and trait composition.

according to some automatic scheme. If such schemes would work in all possible situations and in such a way that the programmer can easily foresee the result of a composition, developing software would be much easier. However, fully automatic resolution of conflicts is no panacea. It is only trivial when the composed semantics are orthogonal, so that conflicts can simply not occur. But it becomes extremely complex or even impossible when the semantics overlaps, which is the case for class properties. Looking in other areas such as multiple inheritance conflicts resolution we see that techniques based on automatic linearization techniques are not always satisfying [16] and often lead to unpredictable method invocations. The same applies here.

Nowadays the problems of composition of services or overlapping aspects is difficult and nearly impossible without the use of metadata. For example, Kienzle and Guerraoui demonstrates that trying to automatically compose transactions with other simpler aspects such as notification is doomed to failure [23]. In a similar vein authentication and encryption composition can only be a success when the encryption is invoked first, but authentication should take precedence over persistency and transactions.

Note that in the context of metaclass composition, the set of metaclass behavior is not predefined and fixed, as such it is possible to load a package in which another meta-programmer has developed new class properties with sensible composition exigence. Therefore any clever composition engine based on meta-information would have to deal with the openness of the set of class properties. Our solution, based on traits, differs from the other approaches since trait composition is automatic as long as there are no conflicts. Conflicts are detected automatically. When there is a conflict, then the traits model offers mechanisms to solve the conflict. This contrasts with the approaches that use an automatic scheme to handle conflicts.

12. Conclusion and future work

The need to reuse class properties led to meta-level architectures based on explicit metaclasses [1,2]. While offering reuse of class properties, such models introduced metaclass composition problems [8]. Different approaches exist that try to solve metaclass compositions problems, based on multiple inheritance, code generation or automatically changing metaclasses [4,10,11]. However, the definition, the

composition and the application of the class property were not controllable by the developer or meta programmer.

Our solution models class properties with *traits* (first class groups of methods), and uses trait composition to safely combine and reuse them. Using traits to compose class properties first of all solves the metaclass composition problems (upward and downward compatibility is ensured) while supporting the reuse of class properties. In addition, composition and conflict resolution are *explicit* and under *control* of the composing entity. Thirdly, traits is a general-purpose composition mechanism for object-oriented languages that we have already applied successfully at the base level (for example, to refactor the Smalltalk collection hierarchy [24]).

We implemented all the examples shown in this article using the Squeak implementation of traits and we started to refactor the kernel of Squeak using traits. Our next step is to use traits to define a new metaobject protocol for Smalltalk.

Acknowledgements

We also like to thank Eric Tanter, Noury Bouraqadi and the anonymous reviewers for their valuable comments and discussions.

References

- [1] Ingalls D. The Smalltalk-76 programming system design and implementation. In: POPL'76; 1976. p. 9–16.
- [2] Cointe P. Metaclasses are first class: the objvlisp model. In: OOPSLA '87; 1987. p.156–67.
- [3] Kiczales G, des Rivières J, Bobrow DG. The art of the metaobject protocol. Cambridge, MA: MIT Press; 1991.
- [4] Danforth S, Forman IR. Derived metaclass in SOM. In: TOOLS EUROPE '94; 1994. p. 63–73.
- [5] Forman IR, Danforth S. Putting metaclasses to work: a new dimension in object-oriented programming. Reading, MA: Addison-Wesley; 1999.
- [6] Ledoux T, Cointe P. Explicit metaclasses as a tool for improving the design of class libraries. In: ISOTAS '96, LNCS 1049. Berlin: Springer; 1996. p. 38–55.
- [7] Goldberg A, Robson D. Smalltalk-80: the language. Reading, MA: Addison-Wesley; 1989.
- [8] Graube N. Metaclass compatibility. In: OOPSLA '89; 1989. p. 305–16.
- [9] Schärli N, Ducasse S, Nierstrasz O, Black A. Traits: composable units of behavior. In: ECOOP 2003, LNCS 2743. Berlin: Springer; 2003. p. 248–74.
- [10] Bouraqadi-Saadani NMN, Ledoux T, Rivard F. Safe metaclass programming. In: OOPSLA '98; 1998. p. 84–96.
- [11] Bouraqadi N. Safe metaclass composition using mixin-based inheritance. Journal of Computer Languages, Systems and Structures 2004;30(1–2):49–61.
- [12] Mulet P, Malenfant J, Cointe P. Towards a methodology for explicit composition of metaobjects. In: OOPSLA '95; 1995. p. 316–30.
- [13] Oliva A, Buzato LE. The design and implementation of guarana. In: USENIX conference on object-oriented technologies and systems (COOTS'99), 1999.
- [14] Ingalls D, Kaehler T, Maloney J, Wallace S, Kay A. Back to the future: the story of Squeak. A practical Smalltalk written in itself. In: OOPSLA '97; 1997. p. 318–26.
- [15] Alpert SR, Brown K, Woolf B. The design patterns smalltalk companion. Reading, MA: Addison-Wesley; 1998.
- [16] Ducournau R, Habib M, Huchard M, Mugnier M. Monotonic conflict resolution mechanisms for inheritance. In: OOPSLA '92; 1992. p. 16–24.
- [17] Rivard F. Évolution du comportement des objets dans les langages à classes réflexifs. PhD thesis, Ecole des Mines de Nantes, Université de Nantes, France; 1997.

- [18] Ducasse S. Evaluating message passing control techniques in Smalltalk. *Journal of Object-Oriented Programming (JOOP)* 1999;12(6):39–44.
- [19] McAffer J. Meta-level programming with coda. In: *ECOOP '95, LNCS 952*. Berlin: Springer; 1995. p. 190–214.
- [20] Redmond B, Cahill V. Supporting unanticipated dynamic adaptation of application behaviour. In: *ECOOP 2002, LNCS 2374*. Berlin: Springer; 2002. p. 205–30.
- [21] Tanter E, Bouraqadi N, Noye J. Reflex—towards an open reflective extension of java. In: *International conference on metalevel architectures and separation of crosscutting concerns, LNCS 2192*. Berlin: Springer; 2001. p. 25–43.
- [22] Snyder A. Inheritance and the development of encapsulated software systems. In: *Research directions in object-oriented programming*. Cambridge, MA: MIT Press; 1987. p. 165–88.
- [23] Kienzle J, Guerraoui R. Aop: Does it make sense? The case of concurrency and failures. In: *ECOOP '2002, LNCS 2374*. Berlin: Springer; 2002.
- [24] Black AP, Schärli N, Ducasse S. Applying traits to the Smalltalk collection hierarchy. In: *OOPSLA '03; 2003*. p. 47–64.