



**HAL**  
open science

# Scoped Extension Methods in Dynamically-Typed Languages

Guillermo Polito, Stéphane Ducasse, Luc Fabresse, Camille Teruel

► **To cite this version:**

Guillermo Polito, Stéphane Ducasse, Luc Fabresse, Camille Teruel. Scoped Extension Methods in Dynamically-Typed Languages. *The Art, Science, and Engineering of Programming*, 2017, 2 (1), 10.22152/programming-journal.org/2018/2/1 . hal-01609310

**HAL Id: hal-01609310**

**<https://hal.science/hal-01609310>**

Submitted on 6 Oct 2017

**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.

# Scoped Extension Methods in Dynamically-Typed Languages

Guillermo Polito<sup>a</sup>, Stéphane Ducasse<sup>b</sup>, Luc Fabresse<sup>c</sup>, and Camille Teruel<sup>1b</sup>

a Univ. Lille, CNRS, Centrale Lille, Inria, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille

b Inria Lille - Nord Europe, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille

c Mines Douai, Mines-Telecom Institute

**Abstract Context.** An extension method is a method declared in a package other than the package of its host class. Thanks to extension methods, developers can adapt to their needs classes they do not own: adding methods to core classes is a typical use case. This is particularly useful for adapting software and therefore to increase reusability.

**Inquiry.** In most dynamically-typed languages, extension methods are globally visible. Because any developer can define extension methods for any class, naming conflicts occur: if two developers define an extension method with the same signature in the same class, only one extension method is visible and overwrites the other. Similarly, if two developers each define an extension method with the same name in a class hierarchy, one overrides the other. To avoid such “accidental overrides”, some dynamically-typed languages limit the visibility of an extension method to a particular scope. However, their semantics have not been fully described and compared. In addition, these solutions typically rely on a dedicated and slow method lookup algorithm to resolve conflicts at runtime.

**Approach.** In this article, we present a formalization of the underlying models of Ruby refinements, Groovy categories, Classboxes, and Method Shelters that are scoping extension method solutions in dynamically-typed languages.

**Knowledge.** Our formal framework allows us to objectively compare and analyze the shortcomings of the studied solutions and other different approaches such as MultiJava. In addition, language designers can use our formal framework to determine which mechanism has less risk of “accidental overrides”.

**Grounding.** Our comparison and analysis of existing solutions is grounded because it is based on denotational semantics formalizations.

**Importance.** Extension methods are widely used in programming languages that support them, especially dynamically-typed languages such as Pharo, Ruby or Python. However, without a carefully designed mechanism, this feature can cause insidious hidden bugs or can be voluntarily used to gain access to protected operations, violate encapsulation or break fundamental invariants.

ACM CCS 2012

- Software and its engineering → Modules / packages;
- Theory of computation → Formalisms;

**Keywords** class extensions, scope, packages, method lookup, dynamic languages

## The Art, Science, and Engineering of Programming

Submitted March 31, 2017

Published August 7, 2017

DOI 10.22152/programming-journal.org/2018/2/1



© Guillermo Polito, Stéphane Ducasse, Luc Fabresse, and Camille Teruel  
This work is licensed under a “CC BY 4.0” license.

In *The Art, Science, and Engineering of Programming*, vol. 2, no. 1, 2018, article 1; 26 pages.

### 1 Introduction

Extension methods are a popular feature in dynamically-typed object-oriented languages. An extension method is a method that a developer adds to a class which he does not own. Variants of extension methods are available in many dynamically-typed languages: it is known as *open classes* in Ruby [10], *categories* in Objective-C [12] and Groovy [9], and *extension methods* in Smalltalk [8] and Pharo [5]. Other languages such as Golo [13], which is a dynamically-typed language but offering only static method resolution, offers class extensions named class 'augmentations' but only supporting function addition.

Extension methods are globally visible in most existing implementations, causing two problems: *accidental overwrites* and *accidental overrides*. An *accidental overwrite* happens when two developers define an extension method with the same signature in the same class: in this case a conflict occurs and one method overwrites the other. An *accidental override* happens when two developers define an extension method with the same signature in the same class hierarchy: one method overrides the other. We call such overrides *accidental* because they happen silently and unintentionally. Another common problem is the absence of dependency declarations between extension methods and their callers. Together with the global visibility of extension methods, this promotes the emergence of hidden dependencies that are difficult to track, especially in a dynamically-typed language.

One way to solve these problems is to assign each extension method a particular scope. Variants of scoped extension methods have already been discussed in the literature with the Classbox model [2, 3, 4], the Method Shelters model [1] and Matriona class extensions [15]. In addition, scoped extension methods have been implemented in Ruby since version 2.1 and in Groovy. These variants, however, have different semantics that must be well understood by the developers. To the day of this writing, there is no clear description and comparison of their semantics as well as pros and cons of their impact on the way applications are built. In addition, these variants rely on dedicated method lookup algorithms to resolve conflicts at runtime and tend to have a negative impact on speed.

In this article, we study the semantic differences between variants of scoped extension methods. We scope our analysis to solutions in dynamically-typed languages. We acknowledge solutions for this problem exist also in the context of statically-typed languages [6, 7, 16], but they are not directly applicable to dynamically-typed languages because they rely on static type information. For the sake of completion, we finally compare solutions for both dynamically and statically typed languages.

The main contributions of this paper are:

- a precise study of the problems induced by extension methods (section 2);
- an in-depth study of existing solutions analyzing their main characteristics (section 3);
- a formalization and comparison of the underlying models of these solutions (section 4 and section 5);

- an example of how such formalisation can be used in the form of a metric to estimate the risk of accidental overrides (section 6).

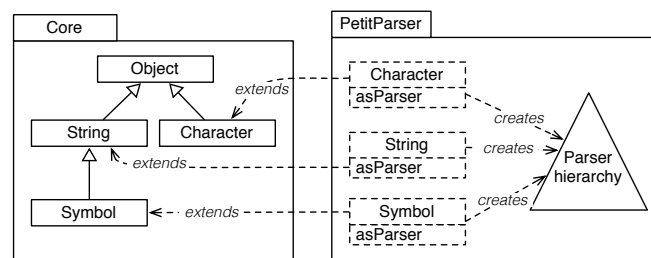
## 2 Extension Methods

This section presents the extension method mechanism in detail. First, we give some common use cases of extension methods. Then, we show some problems induced by globally visible extension methods.

### 2.1 Usage of Extension Methods

We show the advantages of extension methods with examples taken from *PetitParser*, a parser combinator library for Pharo [14]. In *PetitParser*, parsers are modeled as objects and parser combinators accept one or several parsers to produce a new composed parser. Examples of combinators include “,” to sequence two parsers and “star” to repeat a parser zero or more times. For example, the following piece of code shows how we can create a parser that accepts the regular expressions of the form **ab**\*<sup>1</sup>:

```
1 (PPLiteralObjectParser on: $a), (PPLiteralObjectParser on: $b) star.
```



■ **Figure 1** The *PetitParser* parser combinator library defines `asParser` extension methods on core classes to create various kinds of parsers.

**As syntactic sugar.** In addition to these combinators, *PetitParser* defines convenient `asParser` extension methods to some core classes. These extension methods create parsers depending on the receiver (see figure 1). For example, the `asParser` extension method defined in the class `Character` returns<sup>2</sup> a parser that accepts the receiver character.

```
1 Character >> asParser
2   ^ PPLiteralObjectParser on: self
```

<sup>1</sup> The syntax to denote a character in Smalltalk is the character itself preceded by a dollar sign (\$).

<sup>2</sup> The character `^` stands for return in Smalltalk syntax.

## Scoped Extension Methods in Dynamically-Typed Languages

Together, combinators and asParser extension methods give a readable DSL-like syntax. For example, the following expression creates the same parser as before for `ab*`:

```
1 $a asParser, $b asParser star
```

In this example, extension methods asParser act as syntactic sugar *i.e.*, `$a asParser` has the same meaning as `PPLiteralObjectParser on: $a`.

**To improve extensibility.** Extension methods can also improve code quality by making classes polymorphic together. Consider the following code:

```
1 MyParser>>one: a thenMany: b
2   ^ a asParser , b asParser star.
3
4 MyParser>>id           MyParser>>int
5   ^ self              ^ self
6     one: #uppercase    one: ($1 to: $9)
7     thenMany: #letter  thenMany: #digit
```

In the MyParser class, the one:thenMany: method takes as parameter two objects that can be converted into parsers and returns a new parser. The id and int methods use that first method to build custom parsers. The method id sends the message one:thenMany: with two symbols (uppercase and letter) while the method int sends the same message with an interval and a symbol. Extension methods are useful here as they allow any developer to add the method asParser to any class and pass instances of this class to one:thenMany:.

**To adapt classes interface.** The Adapter pattern adapts the interface of an existing class to work with other classes without modifying its source code. The classic realization consists in creating an adaptor class whose instances are used to wrap the instances of the adapted class whenever needed at the expenses of obtaining a different identity. Extension methods permit a class interface to be adapted without relying on an adapter class. Instances of the adapted class can be used directly as they do not need to be wrapped with an adapter object. Therefore, object identity is preserved and less objects are created (no adapters).

**Monkey-patching.** If a third-party library or framework has a bug, a developer can create overwriting extension methods to correct the bug. This technique is known as *monkey patching*. While monkey patching is often recognized as a bad practice in developer communities, it is occasionally useful.

### 2.2 Problems of Globally Visible Extension Methods

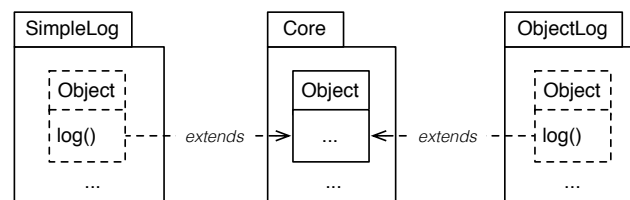
Despite all the benefits that extension methods bring to developers, they can also cause several conflicts and headaches, specially when their usage is not controlled or scoped. Most implementations of extension methods, such as the ones present in various Smalltalk dialects, Ruby (before the introduction of *refinements*) and Objective-C,

make extension methods globally visible. This can lead to undeclared dependencies, *accidental overwrites* and *accidental overrides*.

**Undeclared dependencies.** Once an extension method is loaded it is globally visible. The method can be called from any class of any loaded package without any form of declaration. This means that an application can work correctly in the developer's environment and fail once deployed because the application depends on an extension method from a non-loaded package. The absence of declaration favors the emergence of such hidden dependencies.

**Accidental overwrites.** Extension methods defined by different packages may conflict in two different ways. The first kind of conflict arises when two packages each define an extension method with the same signature in the same class. In this case, one extension method replaces the other. We call this situation an *accidental overwrite*. We show an example in figure 2. The class Object is part of the package Core. A package SimpleLog declares an extension method log for the class Object. This package is a logging framework that records the string representation of an object in a log file. The package ObjectLog declares another extension method log for the class Object. This latter package is another logging framework that serializes objects in a log file for later analysis. Both extension methods conflict and the two logging frameworks cannot be loaded at the same time.

Even though these name clashes happen sparingly, they are difficult to anticipate, especially when considering package co-evolution in large projects involving several packages.



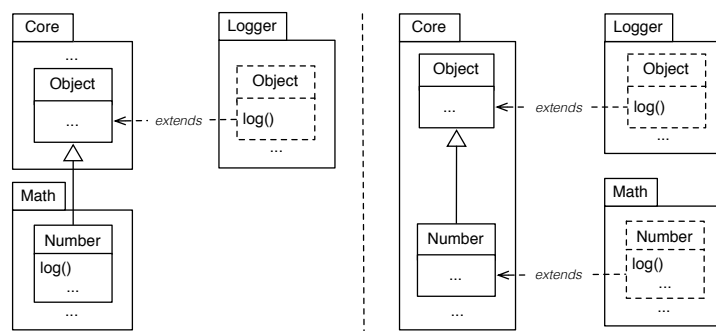
■ **Figure 2** An example of accidental overwrite. Two packages each declare an extension method log for the class Object.

**Accidental overrides.** The second kind of conflict arises when an extension method overrides another extension method defined higher in the class hierarchy. We depict two examples of such overrides in figure 3. On the left part, a regular method log in package Math accidentally overrides an extension method in its superclass declared in package Logger. While Logger's extension method log prints the receiver object in some log file, Math's extension method log computes the logarithm of a number. When they send the message log to an object, users of the Logger package expects that the extension method of Logger is taken into account. However, Number class, as a subclass of Object overrides that log method in package Math. On the right part of figure 3, an extension method in package Math overrides another extension methods in package Logger. The situation is very similar to the previous one. The package Math

## Scoped Extension Methods in Dynamically-Typed Languages

and `Logger` are unaware of each other so none of them know that `Math`'s extension method overrides `Logger`'s.

Large programs usually involve multiple concerns and domains, each coming with its own terminology. Accidental overwrites and overrides happen when these terminologies overlap. In the context of extension methods, the probability of accidental overwrites and overrides is large because any package can declare an extension method for any class. Accidental overrides are a form of interference between packages which is more insidious than accidental overwrites. Indeed, an accidental overwrite is easily noticeable because the client packages are likely to break upon the first invocation of the overwriting method. Accidental overrides are much less noticeable because they affect only instances of the class defining the overriding method. Note that this problem only appears because defining a method implies overriding methods with the same signature upper in the hierarchy.



■ **Figure 3** Two examples of accidental overrides. To the left, a regular method accidentally overrides an extension method. To the right, an extension method accidentally override another extension method.

Since multiple parties can enhance the interface of any class, one party should not be able to override the methods defined by an unrelated party it is not aware of. In other words, extension methods need to be scoped.

### 3 Existing Mechanisms for Scoped Extension Methods

Because extension methods with global visibility exhibit the above-mentioned problems, several implementations propose a narrower visibility. This section describes five of these solutions we selected because they are representative of five different scoping strategies. Depending on the solutions, the scope of activation of extension methods is either lexical or dynamic. In solutions where the scope of activation is lexical, the set of extension methods that are active at a given point is determined statically. In solutions with a dynamic scope of activation, the set of extension methods active at a given point depends on a dynamic context: the call stack. Dynamic scoping is necessary to support a property called *local rebinding*.

First we present the local rebinding property and some of its weaknesses. Then, we show three solutions that expose the local rebinding property: *Classboxes* [3, 4],

*Method Shelters* [1] and Groovy's *categories* [9]. Finally, we present Ruby's *refinements* and Selector Namespaces [17] where extension method activation is determined lexically.

### 3.1 Definitions

In the following, we use the following terms:

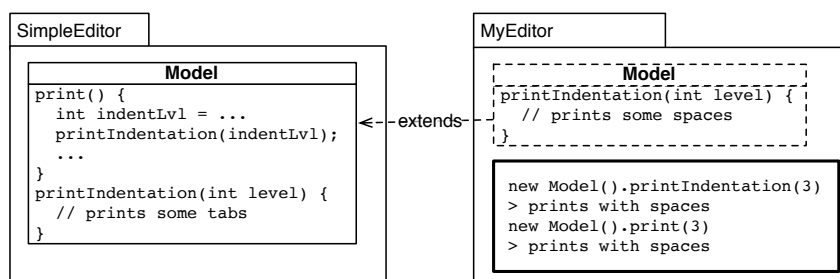
*Package.* We call package, the language-specific unit of deployment that gathers definitions of classes and other constructs from the language. Different packages are potentially maintained by different parties. A package also declares dependencies to other packages by importing some definitions.

*Class extension.* A class extension is a named set of extension methods that apply to the same class. We do not consider addition of instance variables.

*Extension group.* An extension group is a named set of extension methods that may apply to different classes.

### 3.2 Introduction to Local Rebinding

*Local rebinding* is a method-lookup algorithm first defined in the Classbox model [3, 4] and refined in the Method Shelters model [1] and hierarchical layer-based class extensions [15]. This property permits extension methods to override regular methods in a contextual manner. An active extension method takes precedence over regular methods, even for indirect calls. In figure 4, the MyEditor package defines an extension method `printIndentation(int)` that redefines the one in the original package. This extension method prints spaces instead of tabs. When invoked from within this package, this redefinition is taken into account, even in indirect calls: when invoking the `print()` method defined in the SimpleEditor package, the redefined version of `printIndentation(int)` will be executed and not the one defined in the SimpleEditor package.



■ **Figure 4** With local rebinding, changes made by an extension method are applied in case of indirect calls.

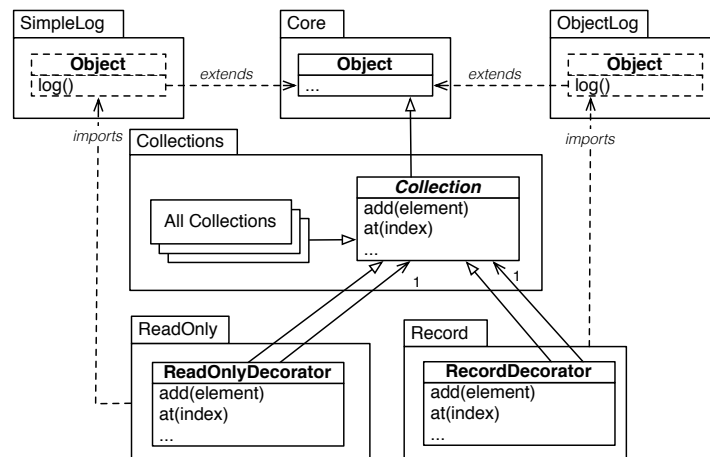
With local rebinding, the lookup algorithm may have to dispatch to different methods in different contexts. In technical terms, when the signature of an extension method  $e$  matches the one of a method  $m$  of the extended class, local rebinding ensures that  $e$  overrides  $m$  during the dynamic extent of message sent by importers of  $e$ . The method lookup algorithm has to access this contextual information to determine



## Scoped Extension Methods in Dynamically-Typed Languages

the active extension methods. Such a method lookup algorithm can be implemented either by inspecting the call stack or by storing the set of active extension methods in a dynamic variable.

### 3.3 Illustrating local rebinding stack behavior



■ **Figure 5** Decorating collections: two potentially conflicting extensions selected by on stack state.

Consider the example depicted in figure 5. A Collections package defines common collections and an abstract class Collection. Two packages ReadOnly and Record each define a collection decorator.

The read-only decorator disables all operations that mutate the decorated collection. When one of these operations is invoked, the read-only decorator logs the attempt using the logging facility of the SimpleLog package and throws an error. The record decorator just logs the operations done on the decorated collection using the logging facility of the ObjectLog package for latter analysis as shown below in pseudo-code.

1	ReadOnlyDecorator>>at(index)	ReadOnlyDecorator>>add(element)
2	return decoree.at(index)	'adding failed'.log();
3		throw IllegalWrite()
4		
5	RecordDecorator>>at(index)	RecordDecorator>>add(anObject)
6	{ 'accessing'. decoree. index }.log();	{ 'adding'. decoree. anObject }.log();
7	return decoree.at(index)	return decoree.add(index)

If a client application uses both decorators together, one log() method is likely to contextually override the other. This is the case when one decorator decorates the other. In this case, the composition order matters because it impacts the call stack and thus the extension methods that are active when looking up log().

```
list = new List([1,2,3,4]);
Case 1 (new ReadOnlyDecorator(new RecordDecorator(list))).at(3);
Case 2 (new RecordDecorator(new ReadOnlyDecorator(list))).add(5);
```

<i>Stack for Case 1</i>	<i>Stack for Case 2</i>
2. RecordDecorator.at()	2. ReadOnlyDecorator.add()
1. ReadOnlyDecorator.at()	1. RecordDecorator.add()

In Case 1, a read-only decorator decorates a record decorator that decorates a list. When sending the `at(3)` message to the read-only decorator, first its `at()` method transfers the request to the record decorator. The `at()` method of the record decorator then tries to log this operation. At this point, two method activations are at the top of the call-stack: first an activation of the `at()` method of record decorator, then an activation of the `at()` method of the read-only decorator. Since each package defining the `at()` method imports a different `log()` extension method, the lookup algorithm must decide which one to select. A similar situation occurs with Case 2 with another call order.

We now study two strategies to select a method in case of ambiguities: *bottom-up local rebinding* and *top-down local rebinding*.

### 3.4 Bottom-up Local Rebinding

The first strategy gives precedence to extension methods imported by callers (*i.e.*, appearing first in the call stack). We refer to this strategy as *bottom-up local rebinding*. This is the strategy of the Classbox [3, 4] and Method Shelters [1] models. In the context of figure 5, this means that the `log()` method of the SimpleLog package is selected in Case 1 and the `log()` extension method of the ObjectLog package is selected in Case 2. This strategy implies that *client code may override other extension methods defined in any package*. As the developer of a package, your methods can be overridden by a package that is indirectly using yours. Consequently, this forces a developer to know the implementation of all the packages it uses (even indirectly) to prevent himself from creating accidental contextual overrides. This raises a tension with information hiding at the package-level and precludes local reasoning.

#### 3.4.1 Classboxes

A classbox is a modular construct defining classes and class extensions, taking the role of a package. A classbox can define at most one *class extension* per imported class. This prevents useful ways to group related extension methods (See section 5.1). A classbox can import class extensions from other classboxes. Classboxes have been devised to facilitate handling of unanticipated changes *i.e.*, a client classbox pushes modifications to other classboxes. Used sparingly, classboxes allows developers to customize the implementation of external packages. However, if used extensively, accidental contextual overrides are likely to occur.

#### 3.4.2 Method Shelters

The Method Shelters model [1] builds upon the Classbox model by adding the ability to protect some extension methods from accidental contextual overriding. A method shelter is a package that contains an *exposed chamber* and a *hidden chamber*. Each chamber declares classes and methods, and they can import other method shelters.

## Scoped Extension Methods in Dynamically-Typed Languages

```
1 shelter :MathShelter
2   class Fixnum #fixed size integer in Ruby
3     def /(x)
4       Rational(self,x)
5     end
6   end
7 end
8
9 shelter :AverageShelter do
10  class Array
11    def avg
12      sum = self.inject(0){|r,i| r + i}
13      sum / self.size
14    end
15  end
16  hide
17  import :MathShelter
18 end

shelter :ClientShelter do
  import :AverageShelter
  def calc
    p([1,2,3,4,5,6,7,8,9,10].avg) #prints "(11/2)"
    p(55/10) #prints 5
  end
end

shelter_eval :ClientShelter do
  calc
end
```

■ **Figure 6** Example taken from [1]. Method shelters provide the ability to control which method can be overridden: extension methods declared after `hide` cannot be overridden by client shelters

Importing a method shelter brings the extension methods of its exposed chamber into the importing chamber. Thus, only methods imported or declared in the exposed chamber can be contextually overridden by other method shelters.

We illustrate the behaviour of hidden chambers in figure 6. In the figure, two definitions of division (`/`) over integers coexist without accidental contextual override. The default `/` method of the `FixNum` class defines euclidian division. The `Math` shelter redefines `/` as exact division: the method returns a rational number. The `average` shelter imports the `math` shelter in its hidden chamber. The `avg` method of `Array` uses the exact division of the `Math` shelter to compute the average of an array of integers. Finally a client shelter imports the `average` shelter and computes the average of an integer array: the computation results in a rational number. The client shelter is oblivious of the fact that the `average` shelter uses the `Math` shelter. From its point of view, `/` still refers to the standard euclidian division.

Method shelters work as classboxes if a program uses only the exposed chambers, and thus, this means that the same problems arise. On the other hand, putting all methods inside the hidden chamber prevents the redefinition of methods, avoiding the local rebinding property.

### 3.5 Top-Down Local Rebinding

The second strategy gives priority to extension methods imported by callees. With this priority strategy, an extension method can be overridden in a called method. In the context of figure 5, this means that the `log()` method of the `ObjectLog` package is selected in Case 1 and the `log()` extension method of the `SimpleLog` package is selected

in Case 2. We refer to this strategy as *top-down local rebinding*. This is the strategy of *Categories* in Groovy [9].

### 3.5.1 Groovy Categories

Groovy developers can define scoped extension methods in *categories*. A category defines a named extension that can be put into the scope of a block of code using the `use` keyword. When a use block is entered, the category is activated by pushing it onto a thread-local stack variable. This extension is popped from the thread-local stack of active extensions when the block is exited. Upon method lookup, a method redefined in a category takes priority over the original method in the extended class. In case of conflict between two extension methods in two categories, the method defined in the lastly-activated category (the one that is nearest to the top of the stack) is selected. Moreover, a use block can activate several categories. If there is conflicting methods in these categories, the first definition hides the others.

## 3.6 Lexical extension activation

This section presents scoped extension mechanisms using a lexical scope of activation, in contrast to the already presented models providing local-rebinding. In these solutions, only extensions defined and imported explicitly in the current lexical scope are active during the execution of a program. This kind of scoped extension methods is provided by *refinements* in Ruby, and *selector namespaces* in *SmallScript*. In the rest of this section we describe ruby refinements and selector namespaces as significant examples of these solutions.

### 3.6.1 Subsystems and Selector Namespaces

We report on the Subsystem proposal since it is probably the source of inspiration for *SmallScript* [17]. In the subsystem proposal [17], method signatures (selectors) are decomposed in two parts: a *value* and a *name*. A selector *value* is the key that is used to actually identify methods. A message send uses a selector *value* to select a method with the same selector *value*. This value is not known to the programmer. A selector *name* is the actual symbolic name used in Smalltalk code to refer to a selector *value*.

A message send using a selector name dispatches to whichever selector value that is bound to it in the lexical scope. In other words, message name resolution is static. Selectors are organized into hierarchies that support redefinition and shadowing. When the Smalltalk compiler processes a message, it looks up the selector names in the current scope and any of its enclosing scopes and uses the selector value that is found.

Unfortunately, the lack of a more clear documentation for *selector namespaces* prevents us from analyzing its properties more in detail.

### 3.6.2 Ruby Refinements

Since its first versions, Ruby supports globally visible extension methods under the name of *open classes*. Ruby 2.1 introduces scoped class extensions under the name

## Scoped Extension Methods in Dynamically-Typed Languages

of *refinements*. In refinements, only modules and classes importing a refinement can call its extension methods. In addition, if a class imports a refinement in its body, this refinement is also active in the scope of the subclasses, even when the subclasses are defined in another package. This propagation of visibility provides some common facilities to subclasses, a feature that may be useful in frameworks where an abstract class of the framework is subclassed by users. Also, developers who subclass an external class should be aware of the refinements that are active in that class. Surprisingly, while the sequence of active refinements can be determined statically, the implementation of refinements does the resolution dynamically with a dedicated and slower method lookup. This choice may be due to other implementation constraints.

### 4 Method Lookup Formalization of Scoped Extension Solutions

We presented in section 3 several models of scoped extension methods. To study the different design choices of each model, we present in this section a formal specification of a method lookup algorithm for scoped extension methods. Using this specification, we formalized the different strategies of the studied solutions to select active extensions and then select the method to execute.

#### 4.1 Notations and Base Model

We use five semantic domains and three functions to model language entities and their relations. We use  $\rightarrow$  to denote partial functions.

$$\begin{array}{ll} c \in \mathcal{C} & \text{superclass} : \mathcal{C} \rightarrow \mathcal{C} \\ s \in \mathcal{S} & \text{method} : \mathcal{C} \times \mathcal{S} \times \mathcal{E} \rightarrow \mathcal{M} \\ m \in \mathcal{M} & \text{imports} : \mathcal{M} \rightarrow \mathcal{E}^* \\ e \in \mathcal{E} & \\ \bar{m} \in \Gamma \text{ where } \Gamma = \mathcal{M}^* & \end{array}$$

Classes are elements of  $\mathcal{C}$ . Signatures are elements of  $\mathcal{S}$ . In a dynamically-typed language a signature usually consists of a name and a number of parameters. Methods are modeled as elements of  $\mathcal{M}$  and extension groups are modeled as elements of  $\mathcal{E}$ .

The partial function *superclass* denotes the class-superclass relationship. Because a class cannot inherit from itself, this function is acyclic. For a class  $c$  that has no superclass,  $\text{superclass}(c)$  is undefined: usually, only one such class exists in programming languages.

The partial function *method* denotes the existence of a method in a given context. This function returns a method if a given extension defines such method with a given signature for a given class. It is undefined if the extension defines no such method.

The function *imports* returns a sequence that models which extension imports are effective in the context of a method, in order of decreasing priority (we note  $\mathcal{X}^*$  the set of finite sequences of elements of  $\mathcal{X}$ ). These imports could be declared for a single

method, for a whole class, for a whole class hierarchy, for a whole package, *etc.* What matters is which one affects a method and this is what *imports* indicates.

Finally, call stacks are elements of  $\Gamma$ . For the purpose of modelling lexical and local-rebinding mechanisms, knowing the method of a stack frame is enough so call stacks are modeled as a finite sequence of methods ( $\Gamma = \mathcal{M}^*$ ). The first element of such a method sequence corresponds to the bottom of the call stack and the last element corresponds to the method that sent the message being looked up. We use the notation  $\bar{x}$  for sequences (*i.e.*,  $\bar{x} = \langle \bar{x}_1, \dots, \bar{x}_{|\bar{x}|} \rangle$ ).

**Standard Lookup.** The standard lookup algorithm for class-based dynamically-typed languages with single dispatch depends on the class of the receiver object and a method signature. To take the dynamic scoping of classboxes and method shelters into account, a lookup algorithm must also consider the call stack. Consequently, we model the lookup as:

$$lookup : \mathcal{C} \times \mathcal{S} \times \Gamma \rightarrow \mathcal{M}$$

A method lookup fails whenever the *lookup* function is undefined. To better distinguish between the different kinds of lookup algorithms, we divide the lookup in two steps. The first step determines the sequence of active method extensions from a call stack. The second step selects a suitable method to be executed among a sequence of extensions. The *lookup* function is then defined using two functions: *activeExts* and *select* that represent these two steps.

$$lookup(c, s, \bar{m}) = select(c, s, activeExts(\bar{m})) \quad \text{where: } \begin{cases} activeExts : \Gamma \rightarrow \mathcal{E}^* \\ select : \mathcal{C} \times \mathcal{S} \times \mathcal{E}^* \rightarrow \mathcal{M} \end{cases}$$

We can now describe different versions of the *activeExts* and *selection* functions separately. We call the different versions of *activeExts*: *active extensions strategies*; and the different versions of *select*: *method selection strategies*.

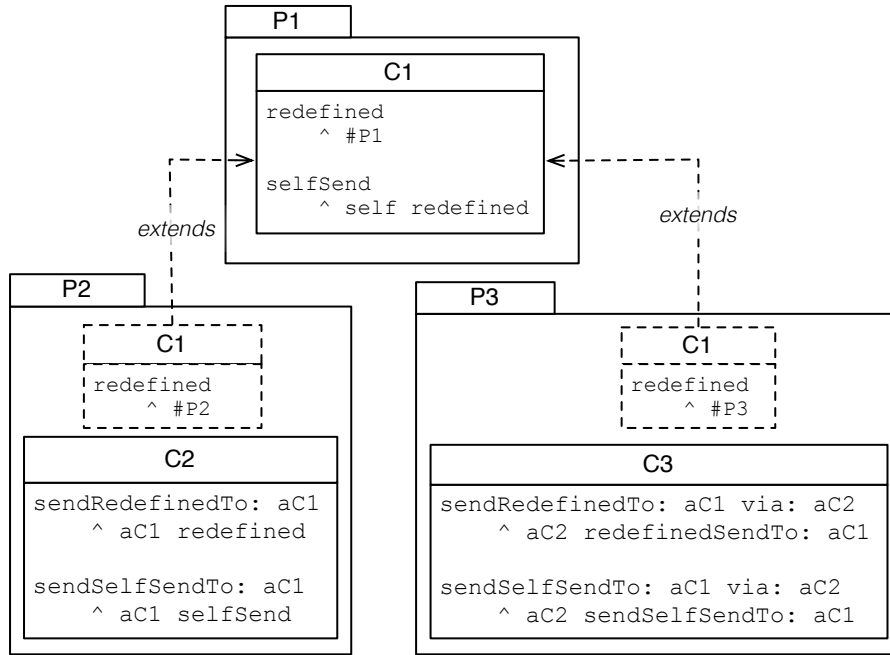
## 4.2 Active Extensions Strategies

We now review the different active extension strategies. In the context of local rebinding, the lookup has to consider the chain of callers to find if one imports an extension with an overriding extension method. The extension activation is dynamically-scoped. This means that the lookup algorithm traverses the call stack or uses a thread-local variable to determine active extensions. The call-stack can be traversed bottom-up giving priority to callers imports, or top-down, giving priority to callees imports. Without local rebinding, the extension activation is said to be lexical. For each strategy, we consider that a global extension named *global* contains all regular methods and it is implicitly imported by default.

Besides the formalization, figure 7 summarizes and illustrates active extension strategies through an example. In the example, two packages P2 and P3 extend class C1 from package P1 with an override. The table illustrates how four different method invocations behave in this scenario: (a) a class in P2 calls a redefined method of C1;

## Scoped Extension Methods in Dynamically-Typed Languages

(b) a class in P2 calls a method of P1 calling a redefined method of C1; (c) a class in P3 calls (a); (d) a class in P3 calls (b). At runtime, this generates overrides between the redefined method in P2 and P3. We see that lexical activations use the definition available in the current lexical scope. Local-rebinding, on the other side, will depend on the order of message sends in the stack. Cases (c) and (d) give precedence to the method defined in P2 or P3, depending on the local-rebinding strategy.



	LR↑	LR↓	Lex
aC2 sendRedefinedTo: aC1	#P2	#P2	#P2
aC2 sendSelfSendTo: aC1	#P2	#P2	#P1
aC3 sendRedefinedTo: aC1 via: aC2	#P3	#P2	#P2
aC3 sendSelfSendTo: aC1 via: aC2	#P3	#P2	#P1

■ **Figure 7** Active Extension Strategies by Example. This figure shows what are the results of four different expressions in the presence of the different active extension strategies.

**Bottom-up local rebinding.** We first consider the extension activation strategy of bottom-up local rebinding as exemplified by Classboxes and also by Method Shelters exposed chambers. The selection of active extensions for method shelters is more refined as it stops searching if one of the shelters is imported in a hidden chamber. Here is the definition of the  $activeExts_{lr\uparrow}$  that computes the active extensions following this strategy:

$$activeExts_{lr\uparrow}(\bar{m}) = imports(\bar{m}_1) \frown \dots \frown imports(\bar{m}_{|\bar{m}|}) \frown \langle global \rangle$$

If the call stack is empty, that is if this is the first lookup of the associated thread, the function just returns the implicitly imported global extensions. Otherwise, the function recursively concatenates the imports of each method in  $\bar{m}$  from the oldest method activation ( $\bar{m}_1$ ) to the newest ( $\bar{m}_{|\bar{m}|}$ ). Concatenation of sequence is noted “ $\frown$ ”. As a result of this bottom-up approach, extensions imported in the methods of the oldest stack frames come first.

**Top-down local rebinding.** Now, we consider top-down call-stack traversal as exemplified by Groovy categories. Here is the definition of the  $activeExts_{lr\downarrow}$  that computes the active extensions following this strategy:

$$activeExts_{lr\downarrow}(\bar{m}) = imports(\bar{m}_{|\bar{m}|}) \frown \dots \frown imports(\bar{m}_1) \frown \langle \text{global} \rangle$$

Like with the function  $activeExts_{lr\downarrow}$ , if call stack is empty  $activeExts_{lr\downarrow}$  returns the implicitly imported global extension. Otherwise, the function recursively concatenates the imports of each method in  $\bar{m}$  from the newest method activation to the oldest. As a result of this top-down approach, extensions imported in the methods of the newest stack frames come first.

**Lexical extension activation.** We finally consider the lexical extension activation strategy as exemplified by Ruby refinements. A lexical extension activation means that the call-site is enough to know the active extensions. It also means that the sequence of active extension is known statically. The active extensions are the ones imported by the calling method, that is the first element of the sequence  $\bar{m}$ .

$$activeExts_{lex}(\bar{m}) = imports(\bar{m}_1) \frown \langle \text{global} \rangle$$

Choosing one of these three active extensions determination strategies (bottom-up local rebinding, top-down local rebinding, lexical) determines which method extensions are active during a message send. The next step of the lookup is to choose a method among these extensions.

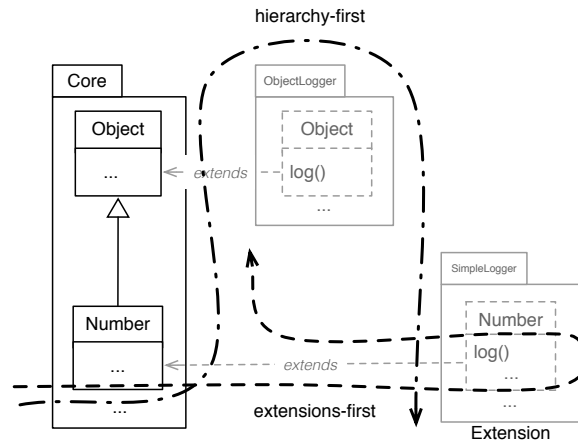
### 4.3 Method Selection Strategy

Once the sequence of active extensions are determined according to one of the previous strategies, the second step is to select one method from all these extensions. One strategy is to lookup a method in the first active extension throughout the hierarchy and then continue with following extensions (See figure 8). We refer to this strategy as *hierarchy-first method selection strategy*. Another solution is to lookup the method in the receiver class for each active scope in order and then continue to the superclass. We refer to this strategy as *extensions-first selection strategy*.

The choice of the method selection strategy has a big impact for the accidental override depicted in figure 3. Indeed, given a sequence of active extensions, these strategies determine whether in a hierarchy two extension methods with the same name from different extensions have an override relationship or not.



## Scoped Extension Methods in Dynamically-Typed Languages



■ **Figure 8** Two method selection strategies: extensions-first and hierarchy first.

**Extensions-first.** It searches for a suitable method in each active extension before searching in the superclass of the receiver class. This is the strategy used by all solutions presented in section 3.

$$select_{ext}(c, s, \bar{e}) = \begin{cases} lookupInClass(c, s, \bar{e}) & \text{if it is defined} \\ select_{ext}(superclass(c), s, \bar{e}) & \text{else if } superclass(c) \text{ is defined} \\ \text{is undefined} & \text{otherwise} \end{cases}$$

The function  $select_{ext}$  first looks if the first of extension in  $\bar{e}$  defines a method for the provided class and signature using the function  $lookupInClass$ . If no method is found (i.e.,  $lookupInClass(c, s, \bar{e})$  is undefined),  $select_{ext}$  continues recursively with the superclass of  $c$  if it exists. Otherwise, it is undefined if  $superclass(c)$  is undefined. The function  $lookupInClass$  searches for the first suitable method defined for a given class in a given sequence of extensions. It is defined as follow:

$lookupInClass(c, s, \langle \rangle)$  is undefined

$$lookupInClass(c, s, \bar{e}) = \begin{cases} method(c, s, \bar{e}_1) & \text{if it is defined} \\ lookupInClass(c, s, tail(\bar{e})) & \text{otherwise} \end{cases}$$

With extension-first method selection, a method can be overridden in extensions with higher priority in the class of the method or in any active extensions in subclasses of that method class.

**Hierarchy-first.** It first looks up for the whole hierarchy of the receiver class in the context of the first extension and then consider the other extensions. As we will see later, this solution permits to limit occurrence of accidental overrides.

$select_{hrc}(c, s, \langle \rangle)$  is undefined

$$select_{hrc}(c, s, \bar{e}) = \begin{cases} lookupInExtension(c, s, \bar{e}_1) & \text{if it is defined} \\ select_{hrc}(c, s, tail(\bar{e})) & \text{otherwise} \end{cases}$$

The function  $select_{hrc}$  first looks if the first scope of  $\bar{e}$  defines a method in the provided  $c$  or in class  $c$  inherits from thanks to the function  $lookupInExtension$ . If no method is found, it continues recursively with the remaining scopes if there is some. The function  $lookupInExtension$  searches for the first method defined in the hierarchy of a given class in a given extension. It is defined as follow:

$$lookupInExtension(c, s, e) = \begin{cases} method(c, s, e) & \text{if it is defined} \\ lookupInExtension & \text{else if } superclass(c) \text{ is defined} \\ (superclass(c), s, \bar{e}) & \\ \text{is undefined} & \text{otherwise} \end{cases}$$

With hierarchy-first method selection, an extension method imported in a subclass can be overridden by extension methods imported by superclasses, if the extensions in the superclasses have higher priority.

## 5 Comparison and Discussion

This section extends the comparison criterion with import granularities. Then, we provide a comparison of existing solutions to expose the concepts. We include in this comparison solutions for statically-typed languages to show how our conceptual decomposition captures also their semantics. We present an analysis of the studied approaches.

### 5.1 Declaration of Dependencies

Once extension methods are local to their users it is mandatory for the users to declare which extension methods they bring into scope. Hence, all the existing solutions here solve the problem of hidden dependencies. These dependencies are usually declared with some form of import statements. Such an import statement between a user (the *importer*) and a set of extension methods (the *importee*) requires answering two questions: “What is imported?” and “Where is it imported?”.

**Importee granularity.** How can a developer declare which extension methods should be imported? Many different granularities can be considered. Importing extension methods one by one is tedious: the solutions presented here offer means to group related extension methods together. One possible grouping is at the class-extension level (used by Classboxes for example) *i.e.*, extension methods are grouped by the class they extend. This kind of grouping is simple but cannot specify a set of related methods in different classes (such as the *asParser* methods presented in section 2.1). Also, with classboxes and method shelters this class-centric grouping cannot specify different sets of methods for the same class. Being able to make different groups for the same class can be useful: one group for a public API while another group is not meant to be exposed because it is for implementation purpose only. Another

## Scoped Extension Methods in Dynamically-Typed Languages

possible grouping is the extension group (used by Method Shelters, Refinements and Categories) where extension methods are grouped under a named extension and can affect different classes.

**Importer granularity.** To which extent/scope is visible an extension? With Classboxes for example, a class extension is imported and visible for all methods in the importing Classbox. With Groovy Categories, extension methods are active during the execution of an importing block. With Ruby Refinements, imported extension methods are visible in the importing class and all its subclasses.

### 5.2 Comparison of Solutions

We summarize in table 1 a comparison of existing scoped extension methods solutions according to the following criteria previously discussed: the importee granularity, the importer granularity, the extension activation strategy and the method selection strategy.

We observe in the table that solutions for dynamically-typed languages (Matriona and Classboxes on Squeak, Method Shelters and Categories on Groovy) use mainly local rebinding solutions. On the other hand, solutions for statically-typed languages (MultiJava and Expanders on Java, PRM Refinements) use lexical activations. The one exception is Ruby Refinements that uses lexical activations on a dynamically-typed language.

■ **Table 1** Comparison of the different approaches to scoped extension methods

	Importee granularity	Importer granularity	Extension activation strategy	Method selection strategy
Classboxes [2, 3, 4]	one class extension per class per package	package	bottom-up local rebinding	extensions-first
Method Shelters [1]	one extension per package	package	controlled bottom-up local rebinding	extensions-first
Groovy Categories [9]	many extensions per package	block of code	top-down local rebinding	extensions-first
Matriona [15]	many extensions per package	package	controlled top-down local rebinding	extensions-first
Ruby Refinements	many extensions per package	class and its subclasses	lexical	extensions-first
MultiJava [6]	many extensions per package	package	lexical	hierarchy-first
Expanders [16]	many extensions per package	package	lexical	hierarchy-first
PRM [7]	many extensions per package	package	lexical	hierarchy-first

### 5.3 Discussion

As one of the authors of local-rebinding [3, 4], and after several years gathering more experience on the topic, we believe that local-rebinding brings more problems than solutions. Accidental contextual overriding asides, local-rebinding violates object encapsulation since the same object can behave differently depending on the caller's code. Lexical activation of extension does not have this problem. Indeed, we consider that lexically-activated extension methods do not cause accidental overrides but just normal intentional overrides because developers know beforehand which extensions are active in a scope and how they may override each other. Therefore they have a simpler and more predictable behavior that allows for local reasoning of a program.

The design space of scoped method extensions is wider than one can expect. For instance, the active extension strategy is not the only design choice. A language designer should also think about method-selection, import granularity, security and so on. For example, we determined that while local-rebinding improves code adaptability it causes too much encapsulation problems. Despite lexical extension methods cannot modify the behavior of an object in a contextual manner like local-rebinding, they are easier to reason about.

The import relationship granularity has consequences on expressivity and segregation of extension methods in meaningful groups. From the importee perspective, we saw that being able to define extensions *i.e.*, named groups of extension methods is the best solution. Extension groups are more powerful than class extensions because:

- an extension can specify a set of related methods in different classes (such as the `asParser` methods presented in section 2.1),
- different extensions can specify different sets of methods for the same class,
- and the previous class-based grouping can be realized with extensions whose methods all belong to the same class.

Methods, classes and packages are all valuable importers granularities and a solution should support all of them. However, this requires the language to support grouping of methods. The imports taken into account for a regular method would be the imports declared at the method-level plus the imports declared at the class-level, plus the imports declared at the package-level. This includes the associated trade-off of increasing the language complexity.

Finally, whereas we called some overrides as “accidental”, they can also be malicious, *e.g.*, voluntarily corrupting a class behavior to gain access to protected operations or break fundamental invariants. This is why the design of scoped extension methods and method selection strategy are crucial because they have a strong impact on accidental overrides as we will see in the next section.

### 6 Our Formal Framework In Action: Example of an Analysis to Minimize Accidental Overrides

As discussed in previous sections, extension methods are useful but accidental overrides are insidious, hard to detect and may be frequent when several packages are involved in a program. For example, in Pharo 3,<sup>3</sup> 4.7% of all methods are extension methods, 16.7% of all classes and traits are extended, 48.1% of all packages define an extension method and 31.7% of all packages define a class or a trait that is extended by another package. These numbers illustrate that in such a practical context, extension methods pose a high risk of accidental overrides limited thanks to coding conventions.

In this section we show how we can use our defined formal framework to propose a metric to estimate the risk of accidental overrides for the two method selection strategies: extension-first and hierarchy-first. We define our metric and use it to determine which strategy provides the least risk. The objective of this section is not to provide bullet-proof metric but to show how our formal framework can be used for this purpose. Language designers are free to not follow this metric.

**Accidental Overriding Space (AOS).** We call accidental override space (AOS) the set of all possible locations where a method could accidentally override another method for a given message. For example, let us consider an arbitrary message  $mess = (c_{rcv}, s, \bar{e})$  with signature  $s$  sent to an instance of  $c_{rcv}$  with the sequence of active extensions  $\bar{e}$ . Let  $meth = method(c_{def}, s, e_i)$  be the method this message dispatches to. This method is declared in the extension  $e_i$ , the  $i$ -th extension of  $\bar{e}$  (possibly global if it is a regular method) for the class  $c_{def}$  (i.e.,  $c_{rcv}$  or one of its superclasses). Now let us consider the addition of an arbitrary method  $new = method(c_{new}, s, e_j)$  with the same signature  $s$  in extension  $e_j$ . We want to model the set of method locations where this new method would cause an accidental override, i.e., the set of method locations that would cause  $mess$  to dispatch to  $new$  instead of  $meth$ . Since the method  $new$  has the same signature as  $meth$ , a method location only consists of a class and an extension. If  $new$  overrides  $meth$  and are defined in the same extension  $e_i$ , this override is intentional, so we only consider locations where  $j \neq i$ .

**AOS of Extension-First Strategy ( $AOS_{ext}$ ).** For extension-first method selection,  $new$  accidentally overrides  $meth$  if: (1)  $new$  is defined for a subclass of  $c_{def}$  in an extension  $e_j$  in  $\bar{e}$  where  $j \neq i$ , or (2)  $new$  is defined for  $c_{def}$  in an extension  $e_j$  where  $j < i$ . If we note  $superclass^{-1+}(c)$  all the subclasses of a class  $c$  (transitive closure of the inverse of  $superclass$ ), we have:

$$AOS_{ext}(mess, meth) = \{(c_{new}, s, e_j) \mid (c_{new} \in superclass^{-1+}(c_{def}) \wedge i \neq j) \vee (c_{new} = c_{def} \wedge i < j)\}$$

The size of  $AOS_{ext}(mess, meth)$  is given by:

$$|AOS_{ext}(mess, meth)| = |superclass^{-1+}(c_{def})| \times (|\bar{e}| - 1) + (i - 1)$$

<sup>3</sup> build #860, which contains 4115 classes/traits and 74648 methods

**AOS of Hierarchy-First Strategy ( $AOS_{hrc}$ ).** For hierarchy-first method selection, *new* accidentally overrides *meth* if *new* is defined for any class in  $c_{def}$  hierarchy in an extension  $e_j$  in  $\bar{e}$  where  $j < i$ . If we note  $superclass^+(c)$  all the superclasses of a class  $c$ , we have:

$$AOS_{hrc}(mess, meth) = \{(c_{new}, s, e_j) \mid c_{new} \in (superclass^{-1+}(c_{def}) \cup c_{def} \cup superclass^+(c_{def})) \wedge i < j\}$$

The size of  $AOS_{hrc, meth}(mess)$  is given by:

$$|AOS_{hrc}(mess, meth)| = (|superclass^{-1+}(c_{def})| + |superclass^+(c_{def})| + 1) \times (i - 1)$$

**Comparison of  $|AOS_{ext}|$  and  $|AOS_{hrc}|$ .** We can now compare the AOS of each method selection strategy. We ask ourself when the hierarchy-first strategy is better than the extension-first strategy *i.e.*, when  $|AOS_{hrc}(mess, meth)| \leq |AOS_{ext}(mess, meth)|$ .

**AOS Comparison in the Pharo Language.** To compare the two metrics defined above, we must have an idea of the average number of subclasses and superclasses a class has. This section shows how this metric is concretized in the case of Pharo. A more general language-independent approach follows after this analysis.

In Pharo the average number of subclasses of a class is 8.82 and the average number of superclasses of a class is 3.83. With these numbers our inequality reduces to  $1.43i - 0.43 \leq |\bar{e}|$ . In the table below, the first row shows  $|\bar{e}|$  ranging from 1 to 10. Remember that  $i$  can range from 1 to  $|\bar{e}|$ . For each value of  $|\bar{e}|$ , the second row shows the maximum value of  $i$  that still satisfies the inequation above. This table shows that hierarchy-first strategy (second row) has less risk to cause accidental overrides than extension-first strategy (first row).

$ \bar{e} $	1	2	3	4	5	6	7	8	9	10
$i \leq$	1	1	2	3	3	4	5	5	6	7

For example, when 5 extensions are active ( $|\bar{e}| = 5$ ), the hierarchy-first strategy has less risk to cause an accidental override than extension-first strategy whenever *meth* belongs to one of the first 3 active extensions ( $i \leq 3$ ). Following the table, we observe that for this example in 67% of total cases hierarchy-first strategy has less accidental method override risks. . Extension-first strategy performs better when *meth* belongs to the last extensions *i.e.*, the ones that have smaller priority. But it also means that using extension-first strategy, accidental overrides can happen for extension with a lower priority. So, hierarchy-first strategy has also the advantage that an accidental method override can only happen for extension with a higher priority. All of these reasons show that the hierarchy-first strategy is better to limit accidental overrides.

**Generalization.** The above analysis relies on average number of subclasses and superclasses of a given class. The following question then arises: Does hierarchy-first strategy always provide a smaller risk of accidental overrides than extension-first strategy? If we make these two numbers varying from 1 to 10 and compute all results,

## Scoped Extension Methods in Dynamically-Typed Languages

it appears that: hierarchy-first performs better when the average number of subclasses is greater than the average number of superclasses (and extension-first on the opposite way). And this assertion is usually true because object-oriented hierarchies are built by extension *i.e.*, subclassing.

**Conclusion.** This analysis shows that generally, the hierarchy-first strategy minimizes the risk of “accidental” overrides in comparison the extension-first strategy.

### 7 Related Work

We have already shown and analyzed in previous sections existing solutions for scoping extension methods. In this section, we compare this work with respect to the problem of conflicts and other module related formalizations.

**Module Taxonomy.** Bergel, Ducasse and Nierstrasz present a module taxonomy in their work “Analyzing Module Diversity” [2]. They present a simple module calculus consisting of a small set of operators over environments and modules. Using these operators, they are then able to specify a set of module combinators that capture the semantics of Java packages, C# namespaces, Ruby modules, selector namespaces, gbeta classes, classboxes, MZScheme units, and MixJuice modules. The article develops a simple taxonomy of module systems. Even if the paper covers Classboxes and selector namespaces, it does not specifically focus on method extensions and does not cover some of the more recent languages supporting class extensions such as Groovy and Method Shelters. In addition, their semantics does not capture the fine grained aspects of the local rebinding lookup stack traversal.

**Accidental overrides.** Simple changes can have unexpected effects due to implicit contracts between a class and its subclasses. This well-known problem, coined as the *fragile base class problem* [11], is due to the fact that current languages do not support well extension contracts: Just changing the calling structure of a method without changing its external behavior may have unexpected effects in presence of subclasses. C# is the one of the rare languages that offers a way to control unintended name capture (called accidental overrides in this paper). C# allows the programmer to qualify a method with the keyword `new` (rather than `override`) to declare that while the newly defined method has the same name as the one in its superclasses, it is used for a different concept than in the superclasses. As such, all calls in the superclass hierarchy that would invoke a method with the same name will not consider the new method.

### 8 Conclusion

Globally-visible extension methods can lead to conflicts: accidental overrides and overwrites. These conflicts pose class encapsulation problems that can lead to subtle

bugs or be exploited by malicious parties. In this article, we analyzed multiple solutions that propose to scope extension methods in dynamically-typed languages: Classboxes, Ruby Refinements, Method Shelters, and Groovy Categories.

We defined scoped extension mechanisms as a combination of a *active extension strategy* and a *method selection strategy*. An active extension strategy defines what extension methods are available in a given context. We identified lexical activation as well as two flavours of local-rebinding activations that were partially described in the literature. A method selection strategy defines how a method is selected when there are multiple active extensions defining methods with the same signature. Method selection strategies can give precedence to the class hierarchy (hierarchy-first) or to the extensions (extensions-first).

We then used these formal semantics to characterize other solutions such as MultiJava, Expanders and Matriona. We show that the semantics of scoped extension methods has a big impact on accidental overrides, and concluded that the combination of lexical extension methods with the hierarchy-first method selection strategy gives the best results to minimize them.

**Acknowledgements** We thank the french DGA (Direction Générale de l'Armement) for the PhD grant of Camille Teruel.

## References

- [1] Shumpei Akai and Shigeru Chiba. “Method Shelters: Avoiding Conflicts Among Class Extensions Caused by Local Rebinding”. In: *AOSD '12*. Potsdam, Germany: ACM, 2012, pages 131–142. ISBN: 978-1-4503-1092-5. DOI: 10.1145/2162049.2162065.
- [2] Alexandre Bergel, Stéphane Ducasse, and Oscar Nierstrasz. “Analyzing Module Diversity”. In: *Journal of Universal Computer Science* 11.10 (Nov. 2005), pages 1613–1644. DOI: 10.3217/jucs-011-10-1613. URL: <http://rmod.inria.fr/archives/papers/Bergo5cModuleDiversity.pdf>.
- [3] Alexandre Bergel, Stéphane Ducasse, and Oscar Nierstrasz. “Classbox/J: Controlling the Scope of Change in Java”. In: *OOPSLA'05*. 2005. ISBN: 1-59593-031-0. DOI: 10.1145/1094811.1094826.
- [4] Alexandre Bergel, Stéphane Ducasse, and Roel Wuyts. “Classboxes: A Minimal Module Model Supporting Local Rebinding”. In: *JMLC'03*. Volume 2789. LNCS. Springer-Verlag, 2003, pages 122–131. ISBN: 978-3-540-45213-3. DOI: 10.1007/b12023.
- [5] Andrew P. Black, Stéphane Ducasse, Oscar Nierstrasz, Damien Pollet, Damien Cassou, and Marcus Denker. *Pharo by Example*. Kehrsatz, Switzerland: Square Bracket Associates, 2009, page 333. ISBN: 978-3-9523341-4-0. URL: <http://files.pharo.org/books/pharo-by-example/>.




## Scoped Extension Methods in Dynamically-Typed Languages

- [6] Curtis Clifton, Gary T. Leavens, Craig Chambers, and Todd Millstein. “MultiJava: Modular Open Classes and Symmetric Multiple Dispatch for Java”. In: *OOP-SLA’00*. 2000, pages 130–145. ISBN: 1-58113-200-X. DOI: 10.1145/353171.353181.
- [7] Roland Ducournau, Floréal Morandat, and Jean Privat. *Modules and Class Refinement: a Meta-Modeling Approach to Object-Oriented Programming*. Technical report 07-021. LIRMM, Université Montpellier II, 2007.
- [8] Adele Goldberg. *Smalltalk 80: the Interactive Programming Environment*. Reading, Mass.: Addison Wesley, 1984. ISBN: 0-201-11372-4.
- [9] Dierk König. *Groovy in action*. Shelter Island, NY: Manning, 2007. ISBN: 978-1-932394-84-9.
- [10] Yukihiro Matsumoto. *Ruby in a Nutshell*. Sebastopol, CA, USA: O’Reilly & Associates, Inc., 2001. ISBN: 0-596-00214-9.
- [11] Leonid Mikhajlov and Emil Sekerinski. “A Study of the Fragile Base Class Problem”. In: *ECOOP’98*. LNCS 1445. Springer-Verlag, 1998, pages 355–383. ISBN: 978-3-540-64737-9. DOI: 10.1007/BFb0054083.
- [12] Lewis J. Pinson and Richard S. Wiener. *Objective-C*. Boston, MA, USA: Addison Wesley, 1988. ISBN: 0-201-50828-1.
- [13] Julien Ponge, Frédéric Le Mouël, and Nicolas Stouls. “Golo, a Dynamic, Light and Efficient Language for Post-Invokedynamic JVM”. In: *PPPJ’13*. Stuttgart, Germany: ACM, Sept. 2013. ISBN: 978-1-4503-2111-2. DOI: 10.1145/2500828.2500844. URL: <https://hal.inria.fr/hal-00848514>.
- [14] Lukas Renggli, Stéphane Ducasse, Tudor Gîrba, and Oscar Nierstrasz. “Practical Dynamic Grammars for Dynamic Languages”. In: *DYLA’10*. June 2010.
- [15] Matthias Springer, Hidehiko Masuhara, and Robert Hirschfeld. “Hierarchical Layer-based Class Extensions in Squeak/Smalltalk”. In: *Modularity’16*. Málaga, Spain: ACM, 2016, pages 107–112. ISBN: 978-1-4503-4033-5. DOI: 10.1145/2892664.2892682.
- [16] Alessandro Warth, Milan Stanojević, and Todd Millstein. “Statically scoped object adaptation with expanders”. In: *OOPSLA ’06*. 2006, pages 37–56. ISBN: 1-59593-348-4. DOI: 10.1145/1167473.1167477.
- [17] Allen Wirfs-Brock. *Subsystems — Proposal*. OOPSLA 1996 Extending Smalltalk Workshop. Oct. 1996.

## About the authors

**Guillermo Polito** is research engineer at CNRS working currently in the RMoD (<http://rmod.lille.inria.fr>) and Emeraude (<http://www.cristal.univ-lille.fr/emeraude/>) teams. His research targets programming language abstractions and tool support for modular long-lived systems. For this, he studies how reflective systems can evolve while maintaining these properties. He is interested in how these concepts combine with distribution and concurrency. Contact him at [guillermo.polito@univ-lille1.fr](mailto:guillermo.polito@univ-lille1.fr).



**Stéphane Ducasse** is directeur de recherche at Inria. He leads the RMoD (<http://rmod.lille.inria.fr>) team. He is expert in two domains: object-oriented language design and reengineering. He worked on traits, composable groups of methods. Traits have been introduced in Pharo, Perl, PHP and under a variant into Scala, Fortress of SUN Microsystems. He is also expert on software quality, program understanding, program visualisations, reengineering and meta-modeling. He is one of the developer of Moose, an open-source software analysis platform <http://www.moosetechnology.org/>. He created <http://www.synectique.eu/> a company building dedicated tools for advanced software analysis. He is one of the leader of Pharo (<http://www.pharo.project.org/>) a dynamic reflective object-oriented language supporting live programming. The objective of Pharo is to create an ecosystem where innovation and business bloom. He wrote several books such as Functional Programming in Scheme, Pharo by Example, Deep into Pharo, Object-oriented Reengineering Patterns, Dynamic web development with Seaside. Contact him at [stephane.ducasse@inria.fr](mailto:stephane.ducasse@inria.fr)



**Luc Fabresse** is associate professor in the CAR research theme (<http://car.mines-douai.fr>) at the Mines-Telecom Institute, Mines Douai, France. His researches aims at easing the development of mobile and constrained software using dynamic and reflective languages such as Pharo. One of his goal is to support live programming of mobile and autonomous robots in an efficient way. He is the co-author of multiple research papers (<http://car.mines-douai.fr/luc>) and he concretizes all these ideas (models and tools) in the PhaROS platform (a Pharo client for the Robotics Operating System) to develop, debug, test, deploy, execute and benchmark robotics applications. Each year, Luc also gives computer science lectures, co-organizes events (technical days, conferences, ...) and promotes Smalltalk as an ESUG (European Smalltalk User Group) board member. Contact him at [luc.fabresse@mines-douai.fr](mailto:luc.fabresse@mines-douai.fr)



## Scoped Extension Methods in Dynamically-Typed Languages

**Camille Teruel** is a young researcher and software engineer at Foretagsplatsen. His works reflection control for security. His main domain of work are on Meta-Object protocols for security and isolation. Contact him at [camille.teruel@gmail.com](mailto:camille.teruel@gmail.com)

