



HAL
open science

A VM-Agnostic and Backwards Compatible Protected Modifier for Dynamically-Typed Languages

Iona Thomas, Vincent Aranega, Stéphane Ducasse, Guillermo Polito, Pablo Tesone

► **To cite this version:**

Iona Thomas, Vincent Aranega, Stéphane Ducasse, Guillermo Polito, Pablo Tesone. A VM-Agnostic and Backwards Compatible Protected Modifier for Dynamically-Typed Languages. *The Art, Science, and Engineering of Programming*, 2023, 10.22152/programming-journal.org/2024/8/2. hal-04119017v2

HAL Id: hal-04119017

<https://hal.science/hal-04119017v2>

Submitted on 23 Dec 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

A VM-agnostic and backward compatible protected modifier for dynamically-typed languages

Iona Thomas^a, Vincent Aranega^a, Stéphane Ducasse^a, Guillermo Polito^a, and Pablo Tesone^a

^a Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 – CRISTAL, France

Abstract

Context. In object-oriented languages, method visibility modifiers hold a key role in separating internal methods from the public API. Protected visibility modifiers offer a way to hide methods from external objects while authorizing internal use and overriding in subclasses.

Inquiry. However, while present in main statically-typed languages, visibility modifiers are not as common or mature in dynamically-typed languages.

Approach. In this article, we present `PROTDYN`, a self-send-based visibility model calculated at compile time for dynamically-typed languages with its formal semantics. `PROTDYN` relies on code rewritings and syntactic differentiation of self and object sends, thus determining method visibility at compile time.

Knowledge. We present `#PHARO`, a `PROTDYN` implementation of this model that is backward compatible with existing programs, and its port to Python.

Grounding. Using these implementations we study the performance impact of `PROTDYN` on the method lookup, in presence of global lookup caches, and polymorphic inline caches. We show that it has a very low impact on performance and keeps the benefits from global lookup cache and polymorphic inline cache.

Importance. Protected modifier semantics enforces encapsulation such as private but allow developers to still extend the class in subclasses. `PROTDYN` offers a VM-agnostic and backward compatible design to introduce protected semantics in dynamically-typed languages.

ACM CCS 2012

- **General and reference** → *Computing standards, RFCs and guidelines*;
- **Applied computing** → **Publishing**;

Keywords Dynamically-typed languages, protected, encapsulation

The Art, Science, and Engineering of Programming

Perspective The Empirical Science of Programming

Area of Submission Social Coding, General-purpose programming



© Iona Thomas, Vincent Aranega, Stéphane Ducasse, Guillermo Polito, and Pablo Tesone
This work is licensed under a “CC BY 4.0” license.
Submitted to *The Art, Science, and Engineering of Programming*.

Protected Modifier

1 Introduction

Method visibility modifiers such as `public`, `protected`, and `private` are present in mainstream statically-typed languages like Java, C#, and C++. These modifiers answer the need to specify if and how a method should be used in other parts of an application. It is well-known that some methods may be defined for internal usage only and therefore should be differentiated from the ones that are available publicly [33, 25, 34, 28]. Visibility modifiers also define whether a method can be overridden.

In dynamically-typed object-oriented languages, visibility modifiers are not as common and not as mature. Ruby defines dynamically enforced `private` and `protected` modifiers [29]. Python enforces `private` methods through name-mangling and does not enforce `protected` methods [27]. An ECMAScript proposal in stage 4 as of the writing of this paper proposes the inclusion of `private` fields [13]. In Smalltalk and its descendants [18, 35, 5] all methods are `public`.

Protected method modifiers in their general form offer two interesting facets: (1) they hide methods from external objects, while allowing the class and its subclasses to invoke `protected` methods, and (2) they allow redefinition in subclasses to support reuse [34]. This dual aspect of `protected` modifiers makes them an interesting concept that fits well with full late-bound object-oriented languages. This raises the question of their introduction in object-oriented dynamically-typed languages (*e.g.*, without static type checking).

In this paper, we present `PROTDYN`, a self-send-based visibility model calculated at compile time that introduces `protected` methods in dynamically-typed languages. Following Schärli *et al.* analysis [31], the model is based on the *syntactic* distinction between object and self-sends [36]. Object-sends can only execute `public` methods, while self-sends can execute `protected` and `public` methods. In addition, subclasses are able to make method access more permissive, by changing a modifier from `protected` to `public`, but not the opposite. This makes our visibility solution determinable at compile-time based on syntactic information.

We present an implementation of `PROTDYN` in Pharo named `#PHARO`. See appendix E for a port to Python. These implementations are (1) optionally loadable, (2) do not require any changes to the default method lookup supported by a virtual machine, and (3) are backward compatible with existing code. Using this `#PHARO` implementation, we validate our model.

We show that our solution based on code-rewriting at compile time introduces a marginally noticeable overhead of around 1% compared to using Pharo without our modifier. `#PHARO` is behaving well when introducing optimizations such as global lookup caches and polymorphic inline caches. These results show that our implementation is a viable one (Section 5).

To summarize, the contributions of this article are:

- the definition of a `protected` method model named `PROTDYN`;
- `#PHARO`, an implementation for Pharo [5] (a Smalltalk descendant) that supports optional `protected` modifiers, on default method lookup with no run-time overhead. A similar implementation in Python is presented in Appendix.

In Section 2, we start by presenting briefly the current protected modifiers in a selection of statically-typed and dynamically-typed languages and defining our model. In Section 3, we describe informally the core elements of our model `PROTDYN`. (In appendix D, we extend `SmalltalkLITE` [4] to formally define the semantics of `PROTDYN`). In Section 4, we describe `#PHARO`, the implementation of this model in Pharo. In Section 5, we show with a performance analysis that introducing protected methods do not have noticeable run-time penalties. Then, we discuss alternative implementations in Section 6 as well as the related work in Section 7 before concluding Section 8.

2 About protected methods

Protected is an important method modifier since it supports encapsulation [31, 30] and at the same time, it enables white-box reuse via overriding in a Hook/Template [9, 34]. The protected method modifier is largely spread over mainstream languages, but even if its semantics are always tied in some way to the visibility of a method across the inheritance tree, they are slightly different in every language. The design of visibility modifiers has two key aspects. First, the visibility semantics specify what methods are accessible from what call-sites. Second, a visibility mechanism determines if a given method is accessible from a given call-site.

2.1 Existing protection modifiers

While the protected modifier generally conveys that a method is hidden from clients and accessible from subclasses, several variations exist in mainstream languages. Programming languages propose different visibility modifiers that have similar names yet they have different visibility semantics, redefinition semantics, and enforcing mechanisms. Notably:

- In Java, a protected method is visible from its class and subclasses and from classes within the same package. They are overridable by subclasses. Their access is statically enforced by the type system [37].
- C++'s protected modifier is similar to Java's, except for visibility by friends function and classes that replaces package visibility [23].
- In Ruby, PHP and C# protected methods have neither package visibility nor friend mechanism. In C# their access is statically enforced by the type system like previous protected modifiers, while in Ruby and PHP they are dynamically enforced.

In Section 7.1 we provide a larger and deeper comparison of modifier semantics.

2.2 Protected semantics for dynamically-typed languages

In this paper, we propose a self-send-based visibility semantics with a compile-time implementation mechanism for dynamically-typed languages based on the syntactic differentiation between self message sends and object message sends. Note that Appendix A contains definitions of terms we use in the following of this paper.

Protected Modifier

Our design choice to distinguish between self message sends and object message sends is guided by the analysis of Schärli *et al.* [30] that we summarize below.

2.2.1 Visibility semantics in Schärli *et al.*

Schärli *et al.* [30] studied different semantics for a modifier's visibility: based on static types, class-based, identity-based, and self-send based. While the modifier example is for locally-bound private methods, the comparison between the different syntax options is still relevant.

Static type. The first option relies on a static type system. In terms of scope, a visibility modifier can be implemented to restrict access to instances of a single class. However, the static type requirement is opposed to the nature of dynamically-typed languages that we are targeting.

Dynamic Class-based. The second option is a class-based modifier: the visibility of a method is assessed at run time according to the class of the sender object and the class where the method has been defined. This can be used to restrict visibility to a single class, or a class hierarchy (the class where the method is defined and its subclasses). With this option, having a version of a method local to a class (or its subclasses) overriding a public method available on superclasses' instances makes reasoning on the program execution more difficult.

Dynamic Identity-based. The third option is an identity-based modifier: a sender object can send a protected message *only to itself*. The receiver's identity of a protected method is compared at run time with the sender's identity. Such a semantic brings subtle differences in method executions that may make programs difficult to predict. It can lead to a program working with some instance of a class and failing with other instances of the same class. In addition, using the same selector on the same object can produce a different result if a protected method is accessed in one case, and a public method for the same selector in another case.

Dynamic Self-send-based. The final option is self-send-based visibility: protected messages can only be sent to self. As this does not rely on dynamic information, it allows one to decide at compile time which message-send sites can access methods with restricted visibility. Self-send-based visibility makes it easier for the programmer to understand quickly where protected methods can be accessed and where they can not. It makes reasoning on code easier.

Schärli *et al.* still points out some issues with symmetry properties. For example `self = arg` might not be equivalent with `arg = self` even if `arg` and `self` point to the same object. This happens for the same reason as the problems with the two previous visibility semantics, but at least, in this case, the developer can identify syntactically whether a protected method will be invoked.

2.2.2 Chosen semantics for a protected modifier

We want a protected modifier with the following properties:

- A **protected** method is visible from a message-send site if the receiver is the same as the current method receiver (*i.e.*, the object referenced by `self` or `this`). This means that a protected method is not visible from a message-send site if the message receiver is not the current method receiver, even though they are of the same class.

- A **protected** method is overridable from the subclasses of the defining class.

We chose a self-send-based visibility semantics as it creates fewer ambiguities than the identity-based one. We say that the semantics above are self-send-based because they do not take into account lexical scoping such as the class, package, or namespace as a design choice and they only rely on the fact that the receiver is or is not syntactically self or super. We also add the following rule:

- A **public** method can not be overridden by a **protected** method in any subclasses. This rule supports subtyping between a class and its subclasses. In addition, it ensures that the symmetry issues will not arise with self = arg and arg = self as long as arg and self point to the same object.

ProtDyn in a nutshell. As we will detail in Section 3, ProtDyn is a self-send-based time visibility model determined at compile where:

- A **protected** method is visible from a self-send site.
- A **public** method (*i.e.*, non-protected) is visible from all sites (object-send and self-send sites).
- A **protected** method is overridable from the subclasses of the defining class.

Notice that this model trades-off programmers' flexibility for a compile-time mechanism based on syntactic information: protected methods will not be visible from the receiver object unless it's syntactically self or this.

3 PROT-DYN: a protected modifier model

Our solution is based on the syntactical differentiation of object-sends and self-sends [30]. Appendix D contains a formal definition of the model.

3.1 Object-send and self-send lookup semantics by example

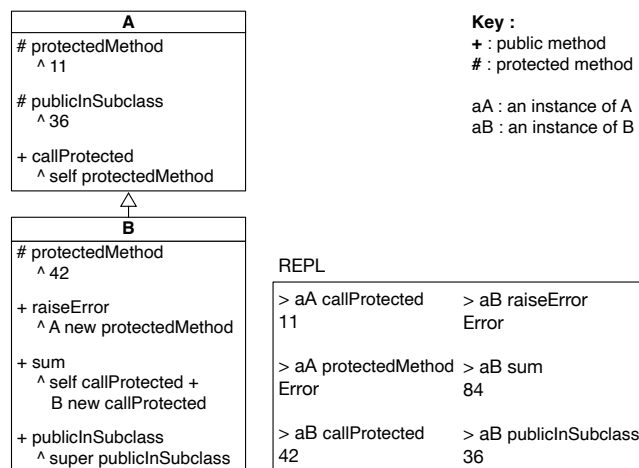
PROT-DYN differentiates send sites by their syntactic receiver, thus object-sends and self-send behave differently:

- Object-sends see only *public* methods.
- Self-sends see *public and protected* methods.

Figure 13 illustrates, with six simple scenarios, the distinction between object-sends and self-sends. In those scenarios, *aA* represents an instance of the A class, and *aB*, an instance of the B class. We also denote with *A»foo* the method with selector *foo* in class A. Additional explanations on each example are available in Appendix B.

These examples illustrate that our compile-time mechanism for protected methods is tied to the syntactic differentiation between object and self-send sites. As we stated in the previous section, this model trade-off programmer's flexibility for a compile-time and syntactically based visibility mechanism. For example, a statement sequence such as `temp := self. temp foo` will have more restricted visibility than directly doing `self foo`. Syntactic differentiation is a simple rule that can be easily learned by developers and makes program understanding easier.

Protected Modifier



■ **Figure 1** Message sending is modified to distinguish between object-sends and self-sends: only object-sends can invoke protected methods which can also be overridden and taken into account by default method lookup.

3.2 About changing visibility in subclasses

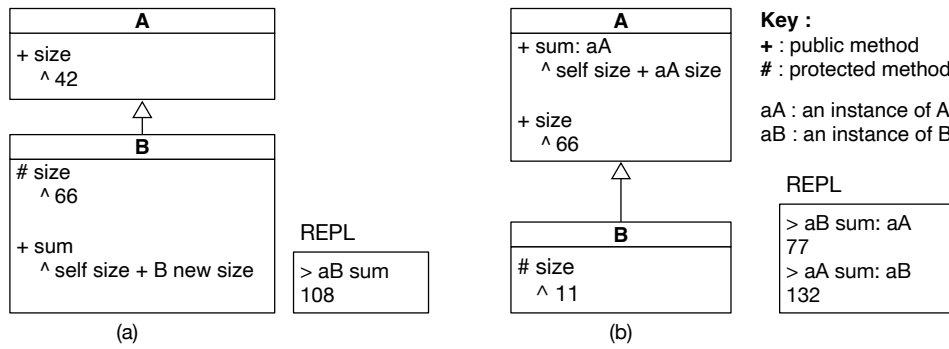
Reducing the visibility of a method in a subclass makes programs difficult to reason about. Figure 2 shows two different examples where reducing visibility creates ambiguities when the method lookup in object-sends skips protected methods. In Figure 2(a), `aB sum` returns 108. The self `foo` self-send lookup finds the `B#foo` method as expected, but the object-send `B new foo` cannot find the `B#foo` method as it is protected. The lookup finds the `A#foo` method which is public. In Figure 2(b), reducing the visibility of the `sum:` method makes the result of a `sum:` send asymmetrical:

- `aB sum:` `aA` would return 77 as expected. The self-send finds the protected `B#foo` method and the object-send lookup finds the `A#foo` method.
- `aA sum:` `aB` would return 132. The self-send lookup finds the `A#foo` method as expected, but the object-send cannot find the `B#foo` method as it is protected. Therefore, the lookup finds the `A#foo` method a second time.

There are two different alternatives to avoid this unintuitive behavior (See Figure 2), that we leave as an implementation choice for language designers:

- *Raising an error at runtime.* For object-sends, the semantics of this solution is that the lookup throws an exception if it finds a protected method. This is the strategy chosen by the Ruby language.
- *Statically forbidding the visibility in subclasses.* This solution is to avoid the problem by construction. This is the solution of statically-typed languages such as Java.

On the other way around, opening up the API by making a protected method public in a subclass provides additional flexibility to the system without leading to the issues discussed earlier. A developer opens the API by redefining a protected method as public in a subclass. Therefore, it does not trade-off encapsulation for flexibility but provides a better way to expose a previously protected method if needed. The ability



■ **Figure 2** Overriding a public method by a protected one leads to buggy situations where self-sends and object-sends can yield two different results. (a) superclass method is used, (b) asymmetrical results.

to increase the visibility of a method is for example used in Java language, where less restrictive modifiers can be applied to an overridden method.

4 Implementation

In this section, we present #PHARO, an implementation for Pharo that performs compile-time restriction of method visibility thanks to the syntactic differentiation between self-send and object-send sites. We first present the design principles behind our implementation and an overview of our solution. We then delve into the two key aspects of our implementation: double public method registration and self-send selector mangling. appendix E presents a Python port of our solution.

4.1 Design Principles

We now establish the principles and constraints that our protected modifier implementation should offer:

- **Backward compatible:** Existing programs (*i.e.*, not using protected modifiers) should continue to work and expose the same behavior under the presence and absence of protected method support.
- **Not requiring a new runtime:** The solution should not be based on changing the virtual machine execution logic, to ease portability and deployment.
- **No run-time penalty when...**
 - **not using protected methods:** A program not using protected modifiers should not be impacted by the presence of the protected modifier implementation.
 - **using protected methods:** A program using protected methods should not have a run-time penalty compared to public method only.

Protected Modifier

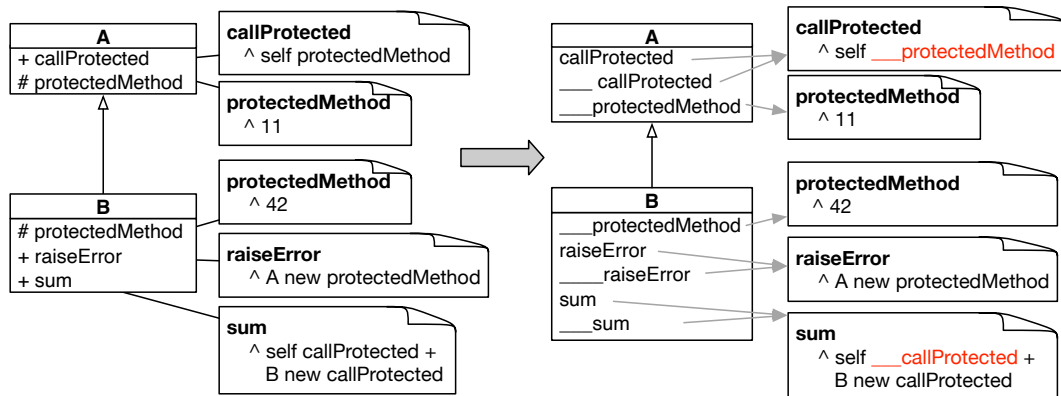
4.2 Implementation overview

```
1 B » bar [  
2   <protected>  
3   ^ 42 ]  
4 B » raiseError [  
5   ^ A new bar ]  
6 B » sum [  
7   ^ self foo + B new foo ]  
8 B » foobar [  
9   ^ super foobar ]
```

■ **Listing 1** Excerpt of the example in Figure 2.

Our implementation relies on the following techniques: *double registration of public methods* and *selector mangling of self-sends*. First, protected methods are identified by developers using method annotations, for example relying on the Pharo annotation system [11]. Figure 1 illustrates the usage of the `<protected>` annotation.

Protected methods are registered in the method dictionary with a mangled selector. Public methods are registered two times, once with their original selector and once with their mangled selector. All self-sends sites are rewritten by mangling the message selector. These two transformations are shown in Figure 3. Then Section 4.5 presents how to limit the recompilation to classes using protected methods and their subclasses.



■ **Figure 3** From the model to the actual implementation: Selector mangling and public method double addition to their class method dictionary.

4.3 Double public registration

Protected methods are added to method dictionaries with a mangled selector (e.g., prefix `__` in the following examples and Figure 3). Mangled selectors are forbidden by the compiler and are not meant to be used directly, as self-sends are modified at compile time to look for methods with this prefix (see Section 4.4). Protected methods are therefore visible from self-sends, but not from object-sends.

In addition, public methods are visible from both self-sends and object-sends (i.e., all sends). Self-sends see the methods installed with the mangled selector while object-sends see methods installed with their original selector. Figure 3 illustrates method registration.

- Protected methods are added once in the defining class with their mangled selector. There are only two entries corresponding to the `__bar` selector, respectively in the dictionaries of classes A and B.
- Public methods entries are added twice in the method dictionary of the defining class: once with their normal selector and once mangled. Their definitions are shared.

4.4 Selector mangling for self-sends sites

During compilation, all self-sends are mangled. In Appendix Section 4.4, we formalize this transformation. Here are key points related to selector mangling for self-sends:

- Self-send sites are rewritten to use mangled selectors to see protected methods. For example, `A»foo` definition contains a self-send to `bar`, hence `A»foo` is rewritten as a send to `__bar`. The same goes for super-sends.
- Mangled self-sends see public methods because they are installed with the mangled selector in addition to the original selector. For example, `B»sum` definition contains a self-send to `foo`, hence it is rewritten as a send to `__foo`. The same goes for super-sends.

The rewrite unit of our implementation is a hierarchy. We chose to recompile all methods in the class and its subclasses and add all the duplicated mangled entries when we add a protected method for the first time. This avoids recalculating the exact set of methods to be recompiled for each new protected method which would be required for lazy recompilation.

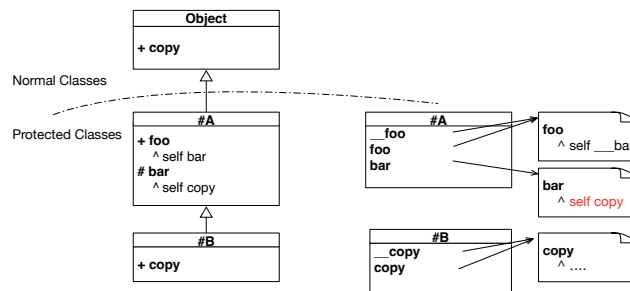
4.5 Preventing selector mangling propagation to the whole system

Our implementation rewrites all self-sends in a whole *descending* hierarchy to avoid unnecessary rewrites of the existing system. Superclasses with no protected methods do not have double registration for public methods, thus self-sends will not find mangled entries in their method dictionaries. Adding selector mangling in superclasses would propagate our transformation upwards in the hierarchy and consequently to the full system.

To avoid this problem our rewriting strategy performs an additional check at compile-time when mangling self-sends. We handle public methods defined above the topmost classes with a protected method differently from public methods in classes using the protected modifier. This is the case of the method `copy` in Figure 4: it is publicly defined in `Object` and the topmost class using a protected method is class `A`. For such a method, we do mangle the self-send in the classes defining protected methods. In Figure 4, the method `bar` invokes `copy` without the mangling as we do otherwise (and explained above). This way we do not have to recompile `Object` and avoid the propagation of the double registration to the entire system.


As the modifier of method can only change from protected to public in subclasses, any redefinition of `copy` is public (as in class `B`). This ensures that invoking `bar` on an instance of `B` correctly finds the redefined version of `copy` in `B`.

Protected Modifier



■ **Figure 4** Limiting the propagation of recompilation with selector mangling for protected to the top of the hierarchy.

A special case appears when a subclass invokes a non-existing method (such as the message `self unknown` in method `foo` of class `A` in Figure 5). In that case, our implementation is conservative and assumes it is a public send.



■ **Figure 5** Compiling a method with an undefined selector.

4.6 Revisiting requirements

Section 4.1 states the requirements that we expect from our implementations. Below we discuss the *backward compatibility* and *not requiring runtime changes* properties. The *no run-time penalty* property is evaluated in Section 5.

- **Backward compatible:** Having `#PHARO` installed does not modify the behavior of existing applications. Indeed, our implementation only affects classes implementing protected methods and their subclasses. In other words, existing programs not using protected methods are not impacted. Moreover, our implementation takes care of not propagating our code transformation upwards in the hierarchy.
- **Not requiring a new runtime:** Our implementation relies on code rewriting techniques that are performed at compile time based on syntactical information. This is possible because all code rewrites happen on classes with explicit protected annotations, their subclasses, and particularly on message send sites that are syntactically identified as self-sends. Thus, our solution does not require a modified VM.

In addition, as a demonstration of the applicability of `PROTDYN` to other dynamic languages than Pharo, we ported our solution to Python (See Appendix E for more details).

5 Performance evaluation

The main objective of our performance evaluation is to determine if our implementation presents non-negligible performance degradations when present (but not used) and when used in the presence/absence of lookup caches. Our evaluation consists of the comparison of three different benchmark configurations, run on three different virtual machine configurations derived from the Pharo VM v8.1.0-alpha-335-g70b7e3542. Notice that we do not perform comparisons within different VMs, but rather comparisons of different benchmark configurations on top of the same VM.

5.1 Experimental design and methodology

Our performance impact evaluates three different scenarios:

The impact in method lookup. Given that public methods increment the size of the selector namespace and in the worst case they double method dictionary sizes, we measure the performance impact in a token-threaded interpreter implementation of Pharo [14], with lookup caches disabled [10].

The impact in presence of a global lookup cache. Global lookup caches cache method lookup results avoiding the subsequent lookup of the same (*receiver-type, selector*) pair [10]. This scenario uses the same token-threaded interpreter implementation above but with a global lookup cache enabled. The global lookup cache is a hash table with 1024 entries that performs a maximum of three probes per message-send before doing a method lookup. This scenario is two-fold: we compare the impact on run-time performance and compare the impact on cache behavior (*i.e.*, # of hits and misses).

The impact in presence of Polymorphic Inline Caches. Polymorphic inline caches further avoid lookups by localizing a lookup cache on each message send-site, typically implemented with machine code stubs and code patching [21]. This scenario uses a mixed-mode Pharo implementation combining a token-threaded interpreter, a 1024-entry global lookup cache, and a non-optimizing method JIT compiler with polymorphic inline caches.

All our benchmarks are run in three different configurations:

Baseline. #PHARO not installed (and thus not used).

Case 1. #PHARO is installed but not used.

Case 2. #PHARO is installed and used in a *worst-case scenario*, as explained below.

As we want an estimation of the run-time penalty using #PHARO, we did not manually protect methods in the existing benchmark code bases. Instead, we tagged all classes in each benchmark as protected. This leads to all the methods being installed using our double registration mechanism and rewritten with mangled self-sends. This corresponds to the worst-case scenario. Indeed, having some of those methods being protected would reduce the total number of entries in method dictionaries.

Protected Modifier

5.2 Methodology and setup

We designed a benchmark methodology following the guidelines of Georges *et al.* and made the best effort to minimize the system’s noise [17, 3]. We did our best effort to close all non-related applications and services, and we shut down the internet connection. The machine was plugged in for the duration of the benchmarks and there was no user interaction until the benchmark finished.

We ran our benchmarks for #PHARO on a MacBook Pro 17.1 with an Apple M1 processor (8 cores including 4 performance cores and 4 efficiency cores) and 16GB of LPDDR4 RAM, running on macOS 12.0.1.

Notice that our methodology uses a fixed number of in-process iterations to determine steady-state instead of dynamically detecting it because none of our three VM configurations include speculative or profile-guided optimizations. More concretely, our methodology is as follows:

- 200 VM invocations.
- 55 benchmark iterations for each VM invocation.
- the first 5 benchmark iterations of each VM invocation are discarded (done to warm up the caches), leaving 50 measures per invocation.
- Package loading is done beforehand and saved in a snapshot.

For each benchmark, we report the average run times of each VM invocation. Figures plot numbers relative to the baseline configuration instead of absolute numbers for readability. For each benchmark, we will show the distributions with violin graphs of the relative run times for the baseline without #PHARO and for the two other cases. The violin graph is completed with a whisker boxplot showing the median, lower, and upper quartiles. The upper whisker is the minimum value between the max relative run time value and the value of the $upper_quartile + 1.5 * interquartile_range$. The lower whisker is the minimum value between the min relative runtime value and the value of the $lower_quartile - 1.5 * interquartile_range$. Assuming a Gaussian distribution, 99% of the values should be inside the whiskers. Any data points outside the whiskers are shown by black dots.

The average relative run time is also shown by a red dot in the middle of each of the violin shapes. This is the main performance indicator used in the analyses.

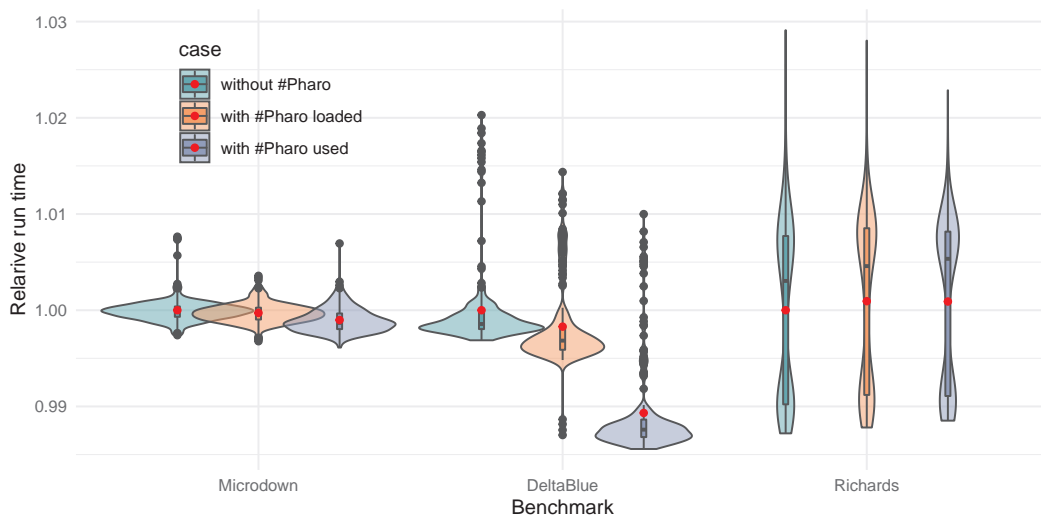
5.3 Selected benchmarks

We selected four different benchmarks, avoiding on principle microbenchmarks because they would not exercise message sends and our protected implementation:

- Microdown Benchmark. Microdown is a new derivative of Markdown. The package defines a parser and document tree. The parser uses a delegation-based approach using many polymorphic calls. We wrote a benchmark implementation that parses the README.md file of the Pharo GitHub repository.
- Delta Blue Benchmark is a benchmark of the DeltaBlue one-way constraint solver. We used the implementation available in the SMark benchmark library [32].

- Richards Benchmark is a benchmark of an OS kernel simulation originally written in BCPL by Martin Richards. We used the benchmark implementation available in the SMark benchmark library [32].
- Bytecode Compiler Benchmark. The Pharo bytecode compiler exercises various compilation aspects: parsing, AST generation, semantic analysis, linear IR generation, and bytecode generation. During the benchmark, we perform several method compilations with source code larger than typical Pharo methods. We used the benchmark implementation available in the SMark benchmark library [32]. Notice that this benchmark exhibits different behavior than the previous three since our protected method implementation introduces modifications in the compiler.

Numerical values for the following benchmarks are available in Appendix F.



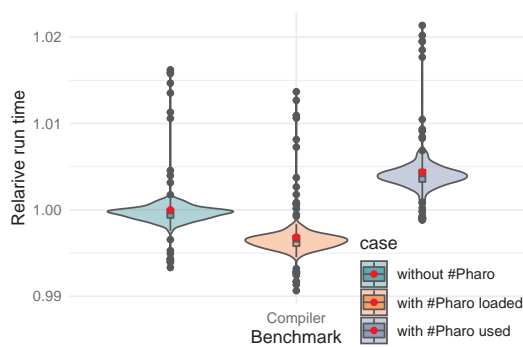
■ **Figure 6** Relative run time performances on the VM without global cache for Microdown, Smark Deltablue and Smark Richards benchmarks. Lower is better. Red dot is the average.

5.4 Experiment 1: Lookup performance

To check the impact of `#PHARO` on lookup performance, we use the VM with no global cache. This means that each time a message is sent the lookup is triggered and it searches in the class and its superclass for the method to activate. Figure 6 shows the distribution of the relative run times of three benchmarks (Microdown, Smark Deltablue, and Smark compiler).

For Microdown and Richards benchmarks, the variation average relative run time is less than 0.1% when `#PHARO` is loaded and when it is used in the worst-case scenario. For the DeltaBlue benchmark, we have an average run time with `#PHARO` used that is at 98.9% of the average run time without `#PHARO`. This corresponds to 9ims less than 8.286s average run time without `#PHARO` (change is less than 1.1%). The run time distribution is consistent for each benchmark.

Protected Modifier




■ **Figure 7** Relative run time performances on the VM without global cache for the Smark Compiler benchmark. Lower is better. Red dot is the average.

Figure 7 shows the performance of the Smark Compiler benchmark. This benchmark has been set aside as the code run is changed by the loading of `#PHARO`. Indeed, loading `#PHARO` will install the compilation hooks needed for the library. The case with `#PHARO` only loaded shows that the average impact at compilation time is a 0.3% speed-up, which corresponds to 13ms less on a 4.22s average run time.

Once the compiler package itself is using `#PHARO`, we have a 0.4% slowdown. This is the biggest slowdown compared to the other benchmarks on the VM without a global cache.

Overall, with this experiment forcing the lookup, we conclude that our prototype does not slow down the lookup process. We think that the small differences that we can observe are a mix of noise in the measurements and outliers that are skewing the average value. If these small performance changes were significant, their impact is usually below 0.5% of the average speed without `#PHARO` and always below 1.1%. *This shows that our prototype is a viable solution and that we do not introduce run-time costs related to the lookup.*



■ **Figure 8** Relative run time performances on the VM with global cache only for Microdown, Smark DeltaBlue and Smark Richards benchmarks. Lower is better. Red dot is the average.

5.5 Experiment 2: Global lookup cache performance

This second experiment focuses on the impact of the global lookup cache on performance, running the same benchmarks as Experiment 1, but using a VM with a global lookup cache.

For this experiment, the benchmark that is slowed down the most is Microdown (see Figure 8). With `#PHARO` only loaded and `#PHARO` used, there is a 1.1% slowdown in the average run time compared to without `#PHARO`. Delta Blue benchmark with `#PHARO` loaded is also slowed down by 0.5%. However, the Delta Blue benchmark with `#PHARO` used is sped up by 0.4%. There is also a speedup on the Richards benchmark, both with `#PHARO` loaded (0.6%) and used (1.7%).

Regarding the compiler Benchmark presented in Figure 9, loading `#PHARO` with the global lookup cache is not slowing the benchmark as we observed in Experiment 1. The speedup using `#PHARO` increases, from 0.3% to over 1% with the global lookup cache.




Figure 9 Relative run time performances on the VM with global cache only for the Smark Compiler benchmark. Lower is better. Red dot is the average.

These changes in relative run time performances can be related to the number of hits from the global lookup cache. See section Section 5.6, where we study the percentages of cache hits and misses on a VM with only a global lookup cache.

With this experiment, we saw that the slowdown of our benchmark is at most 1.1% with a global lookup cache. *This confirms that our prototype performs well in the presence of a global lookup cache.*


5.6 Experiment 3: Lookup cache behavior

Figure 10 shows the percentages of hits and misses for the global lookup cache. Pharo's lookup cache works by probing up to three times in the hash table. A lookup in the table is considered a miss if the third probe fails. The higher the percentage of hits the faster the program will run because cache misses will trigger a costly lookup. These measurements are done after the five warm-up iterations.

In Figure 10, for the Microdown benchmark, we see 5% fewer first probe hits when `#PHARO` is installed compared to when it is not, and again 10% more when it is used. Considering the first and second probe's hits, there are 14% fewer hits on with `#PHARO` installed and 3% more when used. These differences contribute to slowing down the average run time of the Microdown Benchmark seen above.

On the contrary, for the Richards benchmarks with first and second probes, we have 2% more hits with `#PHARO` installed and 11.5% more again with `#PHARO` used.

Protected Modifier



■ **Figure 10** Cache hits and misses.

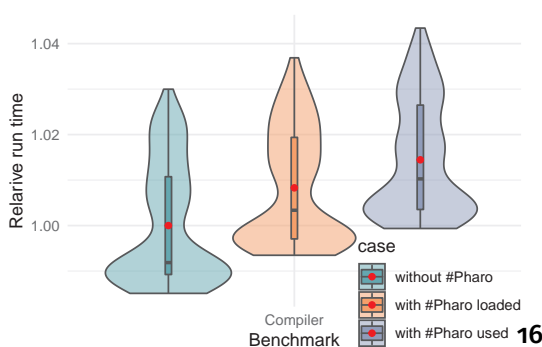
This explains the speedup of the average run time of the Richards Benchmark shown before.

For the Delta Blue benchmark, we see 4% more first probe hits with #PHARO installed and 2% more first and second probe hits. When #PHARO is used we have 2% fewer first probe hits and 2.5% fewer first and second probe hits compared to without #PHARO. It matches the run-time variation on this specific VM.

For the Compiler benchmark, we see more misses: around 6% for all configurations. This may be due to a lot of hash collisions. There is a slight increase in the percentages in misses: +0.07% with #PHARO installed and +0.4% with #PHARO used compared to without #PHARO. However, we doubled the number of selectors in the method dictionary of all the classes of the targeted package, so this is a relatively small increase.

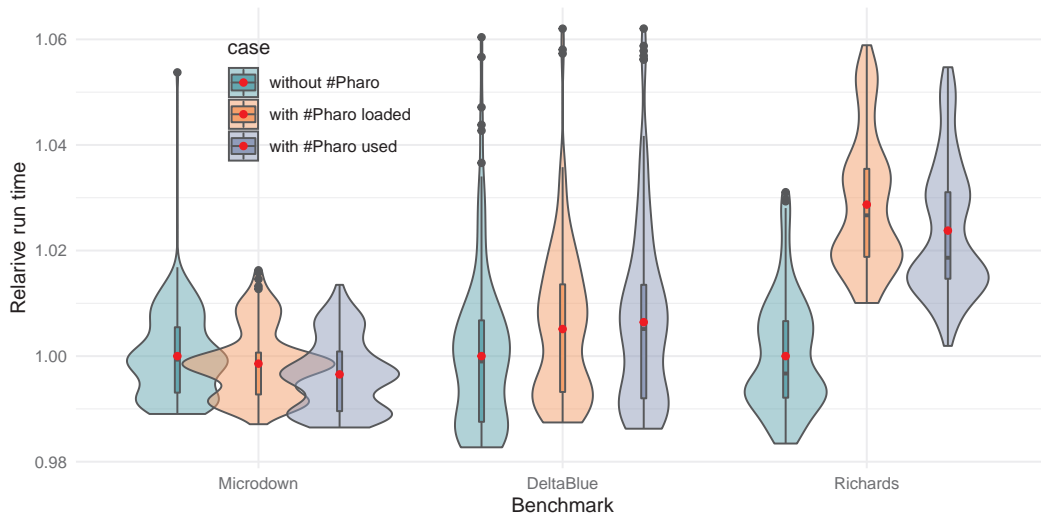
Overall there is little to no variation in the number of misses. Performance variations depend on which probe hits: the faster a probe hits, the faster the benchmark will run. Depending on the benchmark, the percentages of probe hits increase or decrease when #PHARO is installed or used. While we cannot extract general rules on when collisions happen while comparing without #PHARO, with #PHARO loaded and with #PHARO used, we can observe that probe hits are related to the changes in run time performances. Therefore, *the variations in run time performances are linked to global cache hits and #PHARO can have a small positive or negative impact on those.*

5.7 Experiment 4: Polymorphic Inline Cache PIC performance



■ **Figure 12** Relative run time performances on the default VM for the Smark Compiler benchmark. Lower is better. Red dot is the average.

Figure 11 shows the average relative run time distribution for the Microdown, Delta Blue, and Richards benchmarks on a VM with the JIT and the Polymorphic Inline Cache (PIC) enabled. For the Microdown benchmark, we observe speedups of 0.15% and 0.35%



■ **Figure 11** Relative run time performances on the default VM for Microdown, Smark Deltablue, and Smark Richards benchmarks. Lower is better. Red dot is the average.

with `#PHARO` respectively installed and used. The distribution of each case’s relative run times is similar. For the Delta Blue benchmark, the distribution of each case’s relative run times is also similar, but there are slowdowns of 0.51% and 0.64% with `#PHARO` respectively installed and used.

Figure 12 shows the results for the Compiler benchmark. For this benchmark with `#PHARO` installed and used, the average run time is respectively 0.8% and 1.4% slower compared to the version without `#PHARO`.

The Richards Benchmark average relative run time with `#PHARO` installed is the biggest slowdown we encounter with a run time 2.9% longer. With `#PHARO` used it is 2.4% longer. In absolute run time, this corresponds respectively to an additional 12.9ms and 10.6ms compared to 448.2 ms run time without `#PHARO`.

These results reflect the performance of `#PHARO` on the default Pharo VM with the polymorphic inline cache. In the worst-case scenario, we have an average run time that is less than 3% longer on the most optimized VM. However, in most cases, average run time variations with `#PHARO` installed and used are around 1% or less. As we have seen in previous experiments, those variations can be partially due to the global cache performances and are not directly linked to the installation or the use of `#PHARO`.

They are variations in average run times, so we cannot claim that there is **no run-time penalty when not using protected** and the **no run-time penalty when using**

Protected Modifier

protected constraint is met with our implementation. However, we believe that the run time performance variations in both cases are low enough to prove that this is a viable implementation.

6 Implementation decisions

The implementation of #PHARO presents a series of implementation decisions that are worthy to report. Other implementations of PROTDYN may choose differently, so we explore them in this section.

6.1 Lookup mechanism modification

In our implementation, we have avoided modifying the lookup mechanism used in the Pharo VM by using a selector mangling mechanism. We decided to enforce the property **No Requiring a new Runtime**. However, if another implementation does not want to enforce this property (No Requiring a new Runtime), there are alternative implementations that we have explored.

An alternative implementation is to split the lookup into two different operations. One for public methods and the other for protected methods. To identify the methods we consider the following alternatives:

- Marking a method with its visibility, or
- Splitting methods into two different collections (e.g., method dictionaries).

Marking methods with their visibility requires the lookup method to iterate all possible methods in a class and differentiate them when using the public or protected lookup. This approach simplifies the structure of the classes, as it is the same structure as the one used by a runtime that does not support protected methods. It also may affect the performance of the lookup mechanism as it has to filter and iterate more methods than required.

The second alternative (Splitting methods into two different method dictionaries) requires having two different method dictionaries per class. One is to be used by the protected lookup and another for the public lookup. The method dictionary used by the protected lookup mechanism may include all protected methods and public methods. The one used by the public lookup mechanism only includes the public methods. This alternative simplifies the lookup mechanism but requires having two different collections per class.

Having a modified method lookup affects all classes in the system, not only the ones using protected methods. Our implementation only affects the classes using protected methods.

6.2 Visibility checks at run time

In our implementation, when the lookup for a protected method is performed from an object-send site, the protected method will not be found and an error will be raised.

This error is equivalent to the one produced when a message is not implemented in the receiver, in Smalltalk terminology a *Message Not Understood* error.

An alternative implementation could raise an error when a protected method is activated. This strategy is the one used by Ruby. This alternative implementation may intercept the method activation for checking the nature of the receiver and the message send location. However, this alternative requires performing the check for each method activation. We chose to use the lookup solution as it does not require modification of the method activation and it profits from all existing optimizations when performing a lookup, and finally because the check of *Message Not Understood* errors is already implemented and commonly used in Pharo.

	Solution	Visibility	Enforcement	Narrowing
statically	Java	class, subclasses, package	compile time	Forbidden.
	C++	class, subclasses, friends	compile time	Virtual method visibility narrowing is forbidden
	C#	class, subclasses	compile time	Virtual method visibility narrowing is forbidden
dynamically-typed	Smalltalk	all send-sites	N/A	N/A
	Ruby private	class and subclasses using self or implicit receiver	run time	Allowed and visibility is checked at run time
	Ruby protected	class and subclasses	run time	Allowed
	PHP private	class	run time	Forbidden
	PHP protected	class and subclasses	run time	Forbidden
	Python private	class and subclasses	static (name mangling)	
	Python protected	class	convention only	N/A
	JavaScript (private)	class	run-time	N/A
PROTDYN	class and subclasses using self-sends sites	static through name mangling	Forbidden	

■ **Table 1** Accessibility of methods according to visibility modifiers in different languages.

7 Related work

Table 1 summarizes the semantics of the protected modifier for some mainstream languages regarding three aspects: their visibility semantics, visibility mechanism, and narrowing. Visibility semantics vary in each language. In what follows, first, we will focus on Ruby’s method visibility modifiers, as it is one of the few dynamic languages offering them. We will then discuss Python, which uses name mangling to implement its private visibility. Then, we present method encapsulation in Javascript and Java, C#, and C++. Finally, we will discuss work that is related to encapsulation in dynamically-typed languages.

Protected Modifier

7.1 Existing visibility modifiers

Ruby. Ruby is one of the few dynamic languages that offers method modifiers: methods can be qualified as public, protected, or private. Ruby syntactically distinguishes private from public methods [29]. Modifiers can be changed in subclasses: a private method may be made public in subclasses opening further the API, or a public method can be restricted to protected or private leading to errors when not called properly. However, if a subclass uses a stricter modifier, an instance of the such subclass can not be used exactly as an instance of the superclass.

Ruby's protected methods can be invoked by sending a message from the same class where it has been defined or from its descendants. The receiver can be implicit, self, or another instance from the same family. This means that instances of siblings classes can call protected methods defined in a common ancestor on each other. This is different from our approach (which is purely syntactical) and closer to the class-based semantics of Schärli [30].

Ruby's private methods are similar to our protected proposal: in Ruby, private methods can only be called in the class and its subclasses, and only if the receiver is syntactically self or the implicit receiver. In addition, a private method can also be overridden in subclasses. Originally, private methods in Ruby can only be invoked by sending a message to an *implicit* receiver (*i.e.*, no *self*). This restriction had been removed in Ruby 2.7 where private methods can be invoked by sending a message to an *implicit* receiver or on *self*. A call to a private method is not statically bound to the method defined in the class but can be overridden in subclasses.

While the private ruby modifier semantic is similar to #PHARO, the implementation is different: in Ruby, the visibility is checked and enforced dynamically with flags on the methods. On the implementation side, Python (see Section 7.1) is therefore closer to what we are proposing. A semantic difference between Ruby and our model is the fact that super is also a valid receiver of protected sends in our semantics and that we forbid method visibility narrowing.

PHP. PHP supports also three visibility modifiers which are public, protected, and private, however, the last one has a different semantic. In PHP, private methods can only be invoked from the class where it is defined, but they can be invoked on another instance of the same class. This private semantic is class-based, rather than self-send-based like ruby's one. Note that the official description of the language semantics is rather vague. This is why we looked into the PHP Zend Virtual Machine: we found that the visibility mechanism is based on flags set on methods like in Ruby.

Python. Python supports public and private visibility modifiers [27] through name mangling. By default, all attributes and methods are public, and the addition of a double underscore in front of their names marks them as private. Private attributes and methods are mangled by including the name of the class and installed in the defining class, while method code is compiled making a syntactical difference between self-sends and object-sends. Each self-send in the class to a selector identified as private is compiled as a self-send using the mangled name restricting the visibility within the same class but not its subclasses. Protected methods are marked by convention with a simple underscore but are not enforced.

Our solution uses a similar name mangling process on protected methods. However, our name mangling is applied to all self and super-sends, only excluding messages implemented in public-only superclasses to avoid propagation (see Section 4.5). This allows one to change the modifiers in all classes using protected and their subclasses without having recompilation issues.

Javascript. An ECMAScript proposal in stage 4 as of the writing of this paper proposes the inclusion of private fields to enforce encapsulation. Private fields are prefixed with a # and accessible only from inside the class defining it. This is for both properties and methods. Referring to # names outside the scope will result in a syntax error at run time. In contrast, our solution proposes larger visibility while still supporting encapsulation. Moreover, our solution enforces visibility at compile-time using code rewrites.

Java, C#, and C++. In Java, C#, and C++ visibility modifiers are based on the type of the sender and receiver as well as on which package/assemblies they are defined in, they are not based on the identity of each instance. All three support the redefinition of virtual protected methods as public methods in subclasses. In our model, we support the same “opening” of protected methods in subclasses, since it means that from the redefinition of a protected method, object-sends would work as self-send. We decided chose to implement it to offer more flexibility for the developer using classes with protected methods.

7.2 Object encapsulation

Today, most statically-typed object-oriented languages such as Java, C++, and C# provide relatively good support for module encapsulation, and many proposals have been made for augmenting the static type systems of such languages so that they can also express object encapsulation [1, 2, 6, 7, 8, 20, 22, 24, 26]. We report here the work related to dynamically-typed language encapsulation.

In ConstrainedJava [19], the authors extend BeanShell, an extension of Java, to explore dynamic ownership in the context of a dynamically-typed language. Their model for dynamic ownership provides alias protection and encapsulation enforcement by maintaining a dynamic notion of object ownership at run time. It places restrictions on messages sent between objects based on their ownership. Their model further classifies message sends as *internal* (If the target of a message sent is this or is owned by self then the send is internal), or *visible* (sends to other visible objects). In addition, Dynamic Ownership recognizes two kinds of externally independent messages — pure messages that do not access state, and one-way messages that do not return results.

In the programming language MUST[36], methods can be private (visible only in the current class with here), public (visible everywhere), subclass-visible (callable with super), superclass-visible (callable with self), or both latest visibility. This encapsulation is based on a syntactic distinction: message-sends to any object, to self, to super, or to here are treated differently.

Schärli *et al.* proposed encapsulation policies as a way to constrain the interface of an object [30]. With Object-Oriented Encapsulation (OOE), two cases are distinguished: (1) an inheritance perspective where a class can change the way the superclass

Protected Modifier

methods are bound from the subclass perspective and (2) an object perspective where the interface of an object itself may be changed by associating encapsulation policies with object references. From the inheritance perspective, an encapsulation policy associated with a subclass may be defined to change how methods in the superclass are bound. Schärli *et al.* define three different rights to define the encapsulation of a method: the right to override (**o**), to re-implement(**e**), and to call(**c**) a method. OOE semantics is also based on the syntactic distinction of three different messages: super-sends, self-sends, and object-sends. Only self-sends can be early bound.

Adding a protected method mechanism is closely tied to the implementation of encapsulation. The goal of OOE is broader than that of adding protected modifiers as presented in our model. What we want for a protected modifier is something between a **{co}** and an **{o}** policy.

In #PHARO, we make the same distinction between object and self-sends: Object-sends can only access public methods, and self-sends access public or protected methods. This technique has also been chosen in [36, 12] Finally, OOE required changes to the virtual machines to include the new semantics, while our model does not require such changes and can be loaded as a library.

8 Conclusion

In dynamically-typed object-oriented languages, visibility modifiers are not as common or mature. In this paper, we presented `PROTDYN`, a self-send-based visibility model calculated at compile-time for dynamically-typed languages. Our model restricts protected methods activation to self/super-sends and makes them available only from an instance of the class defining them and its subclasses.

We present an implementation of `PROTDYN` in Pharo called #PHARO and a port to Python. Our implementations are designed to be loadable as libraries and add negligible run-time costs. We show that the overhead introduced by our solution based on name-mangling is usually below 1% and behaves nicely when introducing optimizations such as global lookup caches and polymorphic inline caches. This makes it a viable solution to introduce a visibility modifier in dynamic languages.

Acknowledgments We gratefully acknowledge the support of the "Région Hauts-de-France" for the financial support of the PhD of I. Thomas.

References

- [1] Jonathan Aldrich, Valentin Kostadinov, and Craig Chambers. Alias annotations for program understanding. In *OOPSLA '02: Proceedings of the 17th International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, volume 37(11), pages 311–330, New York, NY, USA, 2002. ACM. doi:10.1145/582419.
- [2] Paulo Sérgio Almeida. Balloon types: Controlling sharing of state in data types. In *Proceedings of ECOOP '97*, LNCS, pages 32–59. Springer Verlag, June 1997.

doi:10.1007/BFb0053371.

- [3] Edd Barrett, Carl Friedrich Bolz-Tereick, Rebecca Killick, Vincent Knight, Sarah Mount, and Laurence Tratt. Virtual machine warmup blows hot and cold. In *Proceedings of International Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA '17)*. ACM, October 2017. doi:10.1145/3133876.
- [4] Alexandre Bergel, Stéphane Ducasse, Oscar Nierstrasz, and Roel Wuyts. Stateful traits. In *Advances in Smalltalk — Proceedings of 14th International Smalltalk Conference (ISC'06)*, volume 4406 of LNCS, pages 66–90. Springer, August 2007. doi:10.1007/978-3-540-71836-9_3.
- [5] Andrew P. Black, Stéphane Ducasse, Oscar Nierstrasz, Damien Pollet, Damien Cassou, and Marcus Denker. *Pharo by Example*. Square Bracket Associates, Kehrsatz, Switzerland, 2009. URL: <http://books.pharo.org>.
- [6] Chandrasekhar Boyapati, Alexandru Salcianu, William Beebee, Jr., and Martin Rinard. Ownership types for safe region-based memory management in real-time java. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation*, pages 324–337, New York, NY, USA, 2003. ACM. doi:10.1145/781131.781168.
- [7] John Boyland, James Noble, and William Retert. Capabilities for sharing, a generalisation of uniqueness and read-only. In *Proceedings ECOOP '01*, number 2072 in LNCS, pages 2–27. Springer, June 2001. doi:10.1007/3-540-45337-7.
- [8] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *Proceedings OOPSLA '98*, pages 48–64. ACM Press, 1998. doi:10.1145/286936.286947.
- [9] Serge Demeyer, Theo Dirk Meijler, Oscar Nierstrasz, and Patrick Steyaert. Design guidelines for tailorable frameworks. *Communications of the ACM*, 40(10):60–64, October 1997. URL: <http://scg.unibe.ch/archive/papers/Deme97bDesignGuidelines.pdf>, doi:10.1145/262793.262805.
- [10] L. Peter Deutsch and Allan M. Schiffman. Efficient implementation of the Smalltalk-80 system. In *Proceedings POPL '84*, Salt Lake City, Utah, January 1984. URL: <http://webpages.charter.net/allanms/popl84.pdf>, doi:10.1145/800017.800542.
- [11] Stéphane Ducasse, Eliot Miranda, and Alain Plantec. Pragmas: Literal messages as powerful method annotations. In *International Workshop on Smalltalk Technologies IWST'16*, Prague, Czech Republic, August 2016. doi:10.1145/2991041.2991050.
- [12] Stéphane Ducasse, Roel Wuyts, Alexandre Bergel, and Oscar Nierstrasz. User-changeable visibility: Resolving unanticipated name clashes in traits. In *Proceedings of 22nd International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'07)*, pages 171–190, New York, NY, USA, October 2007. ACM Press. doi:10.1145/1297027.1297040.
- [13] Class field declarations for javascript – consulted on 24 february 2022. <https://github.com/tc39/proposal-class-fields>.
- [14] M. Ertl and David Gregg. The structure and performance of efficient interpreters. *The Journal of Instruction-Level Parallelism*, 5, November 2003.

Protected Modifier

- [15] Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992. doi:10.1016/0304-3975(92)90014-7.
- [16] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 171–183. ACM Press, 1998. URL: <http://www.cs.brown.edu/~sk/Publications/Papers/Published/fkf-classes-mixins/>, doi:10.1145/268946.268961.
- [17] Andy Georges, Dries Buytaert, and Lieven Eeckhout. Statistically rigorous java performance evaluation. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA '07*, pages 57–76, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1297027.1297033.
- [18] Adele Goldberg and David Robson. *Smalltalk 80: the Language and its Implementation*. Addison Wesley, Reading, Mass., May 1983. URL: <http://stephane.ducasse.free.fr/FreeBooks/BlueBook/Bluebook.pdf>.
- [19] Donald Gordon and James Noble. Dynamic ownership in a dynamic language. In Pascal Costanza and Robert Hirschfeld, editors, *DLS '07: Proceedings of the 2007 symposium on Dynamic languages*, pages 41–52, New York, NY, USA, 2007. ACM. doi:10.1145/1297081.
- [20] John Hogg. Islands: aliasing protection in object-oriented languages. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '91), ACM SIGPLAN Notices*, volume 26, pages 271–285, 1991.
- [21] Urs Holzle. *Adaptive Optimization for Self: Reconciling High Performance with Exploratory Programming*. PhD thesis, Stanford University, Stanford, CA, USA, 1994. doi:10.5555/921271.
- [22] Günter Kniesel and Dirk Theisen. JAC — access right based encapsulation for Java. *Software — Practice and Experience*, 31(6):555–576, May 2001.
- [23] Gábor Márton and Zoltán Porkoláb. Selective friends in C++. *Software: Practice and Experience*, 48(8):1493–1519, 2018. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.2587>, doi:10.1002/spe.2587.
- [24] Peter Müller and Arnd Poetzsch-Heffter. Universes: A type system for alias and dependency control. Technical Report 279, Fernuniversität Hagen, 2001. URL: www.informatik.fernuni-hagen.de/pi5/publications.html.
- [25] Oscar Nierstrasz. A survey of object-oriented concepts. In W. Kim and F. Lochovsky, editors, *Object-Oriented Concepts, Databases and Applications*, pages 3–21. ACM Press and Addison Wesley, Reading, Mass., 1989. URL: <http://scg.unibe.ch/archive/osg/Nier89aSurveyOfOOConcepts.pdf>.
- [26] James Noble, Jan Vitek, and John Potter. Flexible alias protection. In Eric Jul, editor, *Proceedings of the 12th European Conference on Object-Oriented Programming (ECOOP '98)*, volume 1445 of *LNCS*, pages 158–185, Brussels, Belgium, July 1998. Springer-Verlag.

- [27] Pep 8 – style guide for python code – consulted on 24 february 2022. <https://www.python.org/dev/peps/pep-0008>.
- [28] Clyde Ruby and Gary T. Leavens. Safely creating correct subclasses without seeing superclass code. In *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '00)*. Association for Computing Machinery, 2000. doi:10.1145/354222.353186.
- [29] Ruby lang 3.0 class module documentation – consulted on 24 february 2022. <https://docs.ruby-lang.org/en/3.0/Module.html>.
- [30] Nathanael Schärli, Andrew P. Black, and Stéphane Ducasse. Object-oriented encapsulation for dynamically typed languages. In *Proceedings of 18th International Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA'04)*, pages 130–149, October 2004. doi:10.1145/1028976.1028988.
- [31] Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Roel Wuyts. Composable encapsulation policies. In *Proceedings of European Conference on Object-Oriented Programming (ECOOP'04)*, volume 3086 of LNCS, pages 26–50. Springer Verlag, June 2004. doi:10.1007/b98195.
- [32] Smark, benchmarking suite for smalltalk – consulted on 26 february 2022. <https://github.com/smarr/SMark>.
- [33] Alan Snyder. Encapsulation and inheritance in object-oriented programming languages. In *Proceedings OOPSLA '86, ACM SIGPLAN Notices*, volume 21, pages 38–45, November 1986. doi:10.1145/28697.28702.
- [34] Patrick Steyaert, Carine Lucas, Kim Mens, and Theo D'Hondt. Reuse contracts: Managing the evolution of reusable assets. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA '96*, pages 268–285. ACM Press, 1996. doi:10.1145/236337.236363.
- [35] David Ungar and Randall B. Smith. Self. In *Conference on History of Programming Languages (HOPL)*, pages 9–1–9–50, New York, NY, USA, 2007. ACM. doi:10.1145/1238844.1238853.
- [36] Mario Wolczko. Encapsulation, delegation and inheritance in object-oriented languages. *IEEE Software Engineering Journal*, 7(2):95–102, March 1992. doi:10.1049/sej.1992.0010.
- [37] Christian Zoller and Axel Schmolitzky. Measuring inappropriate generosity with access modifiers in java systems. In *2012 Joint Conference of the 22nd International Workshop on Software Measurement and the 2012 Seventh International Conference on Software Process and Product Measurement*, pages 43–52, 2012. doi:10.1109/IWSM-MENSURA.2012.15.

A Definitions

Because the focus of this paper is on dynamically-typed languages, we briefly introduce some key vocabulary taken from the Smalltalk terminology.

Protected Modifier

Definition 1. Message. Messages are the key operations in object-oriented languages, also known in other languages as method invocations. A message is composed of a receiver, a selector, and zero or more arguments.

Definition 2. Message receiver. The message receiver is the object targeted by the message.

Definition 3. Message selector. The message selector is an identifier used to choose what method is executed. In dynamically-typed languages, a selector is used as a method signature. In statically-typed languages, the method signature contains a selector and also the types of its arguments and return value.

Definition 4. Method lookup. The method lookup is the process of searching for the right method to execute from a given message. In single-dispatched dynamically-typed object-oriented languages, the method lookup is a function on the receiver object and the selector. It looks up in the receiver's hierarchy the method whose signature is equal to the selector.

Definition 5. Method activation. A method activation is the execution of a method, triggered by a message-send. To activate a method, first, the method-lookup finds the corresponding method, then the method is executed on the message receiver and arguments.

Definition 6. Current method receiver. The current method receiver is the receiver object that led to the current method activation. It is usually denoted with special keywords or pseudo-variables named `self` or `this`.

Definition 7. Message-send site. A message send site is a location in the code where a message is sent, also known in other languages as a call site.

Now that the basic message-send terminology is set up, define the terminology related to method visibility. The following definitions are based on the previous fine-grained distinction between message, method activation, and message-send site.

Definition 8. Method visibility. A method is visible from a message-send site if the method lookup finds this method in the receiver's hierarchy. We say a send-site *can see* a method if the method is visible to the send-site.

Definition 9. Visibility semantics. The visibility semantics of a programming language are the rules that decide which methods are visible from which message-send sites.

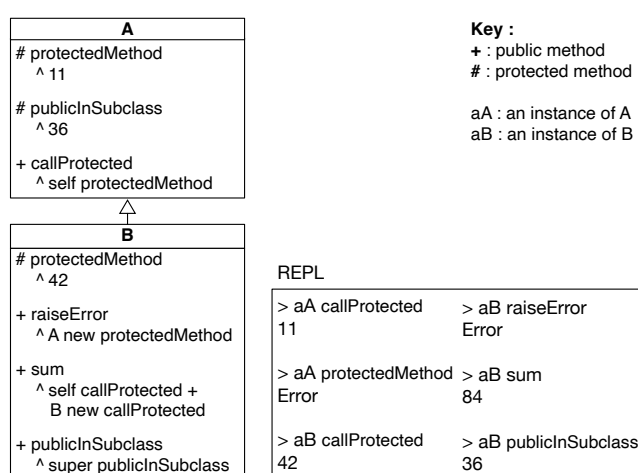
Definition 10. Visibility mechanism. The visibility mechanism of a programming language is the technique used to guarantee that the language visibility semantics are not violated.

Definition 11. self-send site or self-send. A self-send site (self-send for short from now on) is a message-send site where we can syntactically identify the receiver as the current method receiver (*i.e.*, `self` or `this`). Unless they are specifically mentioned, we consider super-send sites as part of self-send sites.

Definition 12. object-send site or object-send. An object-send site (object-send for short from now on) is a message-send site where we cannot syntactically identify the receiver as the current method receiver (*i.e.*, it is not self or this).

Definition 13. Protected and public methods. A protected method is a method explicitly annotated by a developer as *protected*. A public method is a non-protected method.

B Object-send and self-send lookup semantic by example - additional explanation



■ **Figure 13** Message sending is modified to distinguish between object-sends and self-sends: only object-sends can invoke protected methods which can also be overridden and taken into account by default method lookup.

An object-send targetting a protected method produces an error.

- Expression: `aA bar` – an instance of the class A is sent the message `bar`.
- Result: Error
- Explanation: The object-send `aA bar` does not find any public method with the selector `bar` and produces an error. Note that the method is not visible to this message send site even if at run time the receiver of `bar` is an instance of A because the receiver is not syntactically self nor super.

A self-send targetting a protected method activates the method.

- Expression: `aA foo.` – an instance of the class A is sent the message `foo`.
- Result: 11
- Explanation: The object-send `aA foo` finds the public method `A»foo`. From `A»foo` the self-send `self bar` finds and activates the protected method `A»bar`.

Method lookup semantics are not modified by self-sends.

- Expression: `aB foo.` – an instance of the class B is sent the message `foo`.

Protected Modifier

- Result: 42
- Explanation: The object-send `aB foo` finds the public method `A»foo` and activates it on `aB`. From `A»foo` the self-send `self bar` finds and activates the overridden protected method `B»bar` because `aB` the receiver is an instance of `B`.

Protected method visibility is not lexically bound.

- Expression: `aB raiseError`. – an instance of the class `B` is sent the message `raiseError`.
- Result: Error
- Explanation: The object-send `aB raiseError` finds the public method `B»raiseError` and activates it on `aB`. From `B»raiseError` the object-send `A new bar` looks for a public method with selector `bar`, finds none, and produces an error. Notice again that protected methods are not visible to this message send site even if the lexical scope contains a definition of `bar`.

Any message-send targetting a public method activates the method.

- Expression: `aB sum`. – an instance of the class `B` is sent the message `sum`.
- Result: 84
- Explanation: The object-send `aB sum` finds the public method `B»sum` and activates it on `aB`. From `B»sum` (a) the self-send `self foo` finds and activates the superclass' public method `A»foo` and (b) the object-send `B new foo` finds and activates the superclass' public method `A»foo`.

Increasing visibility in subclasses.

- Expression: `aB foobar`. – an instance of the class `B` is sent the message `foobar`.
- Result: 36
- Explanation: The object-send `aB foobar` finds the public method `B»foobar` and activates it on `aB`. From `B»foobar` the self-send `super foobar` finds and activates the superclass' protected method `A»foobar`. This example shows that subclasses are able to redefine and increase of visibility of protected methods in subclasses. Note that restricting visibility is not allowed by construction as explained in the next section.

C SMALLTALKLITE

We present `SMALLTALKLITE`, a Smalltalk-like dynamic language featuring single inheritance, message-passing, field access and update, and **self** and **super** sends. `SMALLTALKLITE` is similar to `CLASSICJAVA`, but removes interfaces and static types. Fields are private in `SMALLTALKLITE`, so only local or inherited fields may be accessed.

We base our approach on the object model used by Flatt *et al.* [16] to give a semantic for mixins for Java-like languages. We adapt the `CLASSICJAVA` model they introduce to develop `SMALLTALKLITE`, a simple calculus that captures the key features of Smalltalk-like dynamic languages.

$$\begin{aligned}
 \epsilon &= v \mid \mathbf{new} \ c \mid x \mid \mathbf{self} \mid \epsilon.f \mid \epsilon.f = \epsilon \\
 &\quad \mid \epsilon.m(\epsilon^*) \mid \mathbf{super}\langle o, c \rangle.m(\epsilon^*) \mid \mathbf{let} \ x = \epsilon \ \mathbf{in} \ \epsilon \\
 E &= [\] \mid o.f = E \mid E.m(\epsilon^*) \mid o.m(v^* \ E \ \epsilon^*) \\
 &\quad \mid \mathbf{super}\langle o, c \rangle.m(v^* \ E \ \epsilon^*) \mid \mathbf{let} \ x = E \ \mathbf{in} \ \epsilon \\
 v, o &= \mathbf{nil} \mid \mathbf{oid}
 \end{aligned}$$

■ **Figure 14** Redex syntax

C.1 SMALLTALKLITE reduction semantics

The syntax of SMALLTALKLITE is shown in Figure 15. SMALLTALKLITE is similar to CLASSICJAVA, while eliding the features related to static typing. We similarly ignore features that are not relevant to a discussion of traits, such as reflection or class-side methods.

To simplify the reduction semantics of SMALLTALKLITE, we adopt an approach similar to that used by Flatt *et al.* [16], namely we annotate field accesses and **super** sends with additional static information that is needed at “run-time”. This extended redex syntax is shown in Figure 14. The figure also specifies the evaluation contexts for the extended redex syntax in Felleisen and Hieb’s notation [15].

Predicates and relations used by the semantic reductions are listed in Figure 17. (The predicates $\text{CLASSES_ONCE}(P)$ *etc* are assumed to be preconditions for valid programs, and are not otherwise explicitly mentioned in the reduction rules.)

$P \vdash \langle \epsilon, \mathcal{S} \rangle \leftrightarrow \langle \epsilon', \mathcal{S}' \rangle$ means that we reduce an expression (redex) ϵ in the context of a (static) program P and a (dynamic) store of objects \mathcal{S} to a new expression ϵ' and (possibly) updated store \mathcal{S}' . A redex ϵ is essentially an expression e in which field names are decorated with their object contexts, *i.e.*, f is translated to $o.f$, and **super** sends are decorated with their object and class contexts. Redexes and their subexpressions reduce to a value, which is either an object identifier or \mathbf{nil} . Subexpressions may be evaluated within an expression context E .

The store consists of a set of mappings from object identifiers $\mathbf{oid} \in \text{dom}(\mathcal{S})$ to tuples $\langle c, \{f \mapsto v\} \rangle$ representing the class c of an object and the set of its field values. The initial value of the store is $\mathcal{S} = \{\}$.

Translation from the main expression to an initial redex is specified out by the $o[[e]]_c$ function (see Figure 16). This binds fields to their enclosing object context and binds **self** to the \mathbf{oid} of the receiver. The initial object context for a program is \mathbf{nil} . (*i.e.*, there are no global fields accessible to the main expression). So if e is the main expression associated with a program P , then $\mathbf{nil}[[e]]_{\text{Object}}$ is the initial redex.

The reductions are summarised in Figure 18.

new c [*new*] reduces to a fresh \mathbf{oid} , bound in the store to an object whose class is c and whose fields are all \mathbf{nil} . A (local) field access [*get*] reduces to the value of the field. Note that it is syntactically impossible to access a field of another object. The

Protected Modifier

```

P = defn*e
defn = class c extends c { f*meth* }
e = new c | x | self | nil
    | f | f=e | e.m(e*)
    | super.m(e*) | let x=e in e
meth = m(x*) { e }
c = a class name | Object
f = a field name
m = a method name
x = a variable name

```

■ **Figure 15** SMALLTALKLITE syntax

```

o[[new c']]c = new c'
o[[x]]c = x
o[[self]]c = o
o[[nil]]c = nil
o[[f]]c = o.f
o[[f=e]]c = o.f=o[[e]]c
o[[e.m(ei*)]]c = o[[e]]c.m(o[[ei]]c*)
o[[super.m(ei*)]]c = super(o, c).m(o[[ei]]c*)
o[[let x=e in e']]c = let x=o[[e]]c in o[[e']]c

```

■ **Figure 16** Translating expressions to redexes

redex notation $o.f$ is only generated in the context of the object o . Field update [set] simply updates the corresponding binding of the field in the store.

When we send a message [send], we must look up the corresponding method body e , starting from the class c of the receiver o . The method body is then evaluated in the context of the receiver o , binding **self** to the receiver's *oid*. Formal parameters to the method are substituted by the actual arguments (see Figure 19). We also pass in the actual class in which the method is found, so that **super** sends have the right context to start their method lookup.

super sends [super] are similar to regular message sends, except that the method lookup must start in the superclass of the class of the method in which the **super** send was declared. When we reduce the **super** send, we must take care to pass on the class c'' of the method in which the **super** method was found, since that method may make further **super** sends. **let in** expressions [let] simply represent local variable bindings.

Errors occur if an expression gets “stuck” and does not reduce to an *oid* or to nil. This may occur if a non-existent variable, field, or method is referenced (for example, when sending any message to nil). For the purpose of this paper, we are not concerned with errors, so we do not introduce any special rules to generate an error value in these cases.

\prec_p	Direct subclass $c \prec_p c' \iff \mathbf{class\ } c \mathbf{\ extends\ } c' \dots \{\dots\} \in P$
\leq_p	Indirect subclass $c \leq_p c' \equiv$ transitive, reflexive closure of \prec_p
\in_p	Field defined in class $f \in_p c \iff \mathbf{class\ } \dots \{\dots f \dots\} \in P$
\in_p	Method defined in class $\langle m, x^*, e \rangle \in_p c \iff \mathbf{class\ } \dots \{\dots m(x^*)\{e\} \dots\} \in P$
\in_p^*	Field defined in c $f \in_p^* c \iff \exists c', c \leq_p c', f \in_p c'$
\in_p^*	Method lookup starting from c $\langle c, m, x^*, e \rangle \in_p^* c' \iff c' = \min\{c'' \mid \langle m, x^*, e \rangle \in_p c'', c \leq_p c''\}$
CLASSESONCE(P)	Each class name is declared only once $\forall c, c', \mathbf{class\ } c \dots \mathbf{class\ } c' \dots$ is in $P \Rightarrow c \neq c'$
FIELDONCEPERCLASS(P)	Field names are unique within a class declaration $\forall f, f', \mathbf{class\ } c \dots \{\dots f \dots f' \dots\}$ is in $P \Rightarrow f \neq f'$
FIELDSUNIQUELYDEFINED(P)	Fields cannot be overridden $f \in_p c, c \leq_p c' \implies f \notin_p c'$
METHODONCEPERCLASS(P)	Method names are unique within a class declaration $\forall m, m', \mathbf{class\ } c \dots \{\dots m(\dots)\{\dots\} \dots m'(\dots)\{\dots\} \dots\}$ is in $P \Rightarrow m \neq m'$
COMPLETECLASSES(P)	Classes that are extended are defined $\text{range}(\prec_p) \subseteq \text{dom}(\prec_p) \cup \{\text{Object}\}$
WELLFOUNDEDCLASSES(P)	Class hierarchy is an order \leq_p is antisymmetric
CLASSMETHODSOK(P)	Method overriding preserves arity $\forall m, m', \langle m, x_1 \dots x_j, e \rangle \in_p c, \langle m, x'_1 \dots x'_k, e' \rangle \in_p c', c \leq_p c' \implies j = k$

■ **Figure 17** Relations and predicates for SMALLTALKLITE

D PROTECTEDLITE: PROTDYN semantics

To specify the semantics of self-sends versus object-sends, we define PROTECTEDLITE, an extension of SMALLTALKLITE [4]. SMALLTALKLITE is a dynamic language calculus featuring single inheritance, message-passing, field access and updates, and self/super sends. The syntax used in the calculus is presented in Figure 20.

SMALLTALKLITE is heavily inspired by CLASSICJAVA defined by Flatt *et al.* [16]. We do not consider it to be a contribution of this article. We repeated the full description of SMALLTALKLITE in the appendix to help the reader.

Each class in PROTECTEDLITE has a list of protected methods after the public ones (see Figure 20). The MethodOncePerClass predicate presented in Figure 21 specifies that two methods shall not have the same name, even if they have different modifiers.

The OverridingPublicMethod and OverridingProtectedMethod predicates guarantee that public methods can be only overridden by public methods and that protected methods can be overridden by public or protected methods. These predicates guarantee per construction that is not possible to restrict the visibility of a public method.

The next four predicates express the lookup mechanism used in PROTECTEDLITE. Predicates \in_p^* and $\widehat{\in}_p^*$ express the lookup of public and protected methods respectively. The lookup mechanism finds the closest superclass (c') of the receiver class (c)

Protected Modifier

$P \vdash \langle E[\mathbf{new} \ c], \mathcal{S} \rangle \hookrightarrow \langle E[oid], \mathcal{S}[oid \mapsto \langle c, \{f \mapsto \text{nil} \mid \forall f, f \in_p^* c\} \rangle] \rangle$	[new] where $oid \notin \text{dom}(\mathcal{S})$
$P \vdash \langle E[o.f], \mathcal{S} \rangle \hookrightarrow \langle E[v], \mathcal{S} \rangle$	[get] where $\mathcal{S}(o) = \langle c, \mathcal{F} \rangle$ and $\mathcal{F}(f) = v$
$P \vdash \langle E[o.f=v], \mathcal{S} \rangle \hookrightarrow \langle E[v], \mathcal{S}[o \mapsto \langle c, \mathcal{F}[f \mapsto v] \rangle] \rangle$	[set] where $\mathcal{S}(o) = \langle c, \mathcal{F} \rangle$
$P \vdash \langle E[o.m(v^*)], \mathcal{S} \rangle \hookrightarrow \langle E[o[[e[v^*/x^*]]]_{c'}], \mathcal{S} \rangle$	[send] where $\mathcal{S}[o] = \langle c, \mathcal{F} \rangle$ and $\langle c, m, x^*, e \rangle \in_p^* c'$
$P \vdash \langle E[\mathbf{super}(o, c).m(v^*)], \mathcal{S} \rangle \hookrightarrow \langle E[o[[e[v^*/x^*]]]_{c''}], \mathcal{S} \rangle$	[super] where $c \prec_p c'$ and $\langle c', m, x^*, e \rangle \in_p^* c''$ and $c' \leq_p c''$
$P \vdash \langle E[\mathbf{let} \ x=v \ \mathbf{in} \ \epsilon], \mathcal{S} \rangle \hookrightarrow \langle E[\epsilon[v/x]], \mathcal{S} \rangle$	[let]

■ **Figure 18** Reductions for SMALLTALKLITE

new $c [v/x]$	=	new c
$x [v/x]$	=	v
$x' [v/x]$	=	x'
self $[v/x]$	=	self
nil $[v/x]$	=	nil
$f [v/x]$	=	f
$f=e [v/x]$	=	$f=e[v/x]$
$e.m(e_i^*) [v/x]$	=	$e[v/x].m(e_i^*[v/x])$
super . $m(e_i^*) [v/x]$	=	super . $m(e_i^*[v/x])$
let $x=e$ in $e' [v/x]$	=	let $x=e[v/x]$ in e'
let $x'=e$ in $e' [v/x]$	=	let $x'=e[v/x]$ in $e'[v/x]$

■ **Figure 19** Variable substitution

P	=	$\text{defn}^* e$	meth	=	$m(x^*) \{ e \}$
defn	=	class c extends $c' \{ f^* \text{meth}^* \text{protectmeth}^* \}$	protectmeth	=	$\#m(x^*) \{ e \}$
e	=	new $c \mid x \mid \mathbf{self} \mid \text{nil}$	c	=	a class name \mid Object
		$f \mid (f=e) \mid e.m(e^*)$	f	=	a field name
		super . $m(e^*) \mid \mathbf{let} \ x=e \ \mathbf{in} \ e$	m	=	a method name
			x	=	a variable name

■ **Figure 20** Protected Pharo syntax.

that contains a method with the given selector (m'). The main difference between both mechanisms is that the public lookup only searches in public methods and the protected lookup only in protected ones.

The public and protected lookup mechanisms are then used in Figure 22 to define the self-send, super-send, and object-send. When a message is sent to an object (without using self or super), we look up the method body, starting from the class of

METHODONCEPERCLASS(P)	Method names are unique within a class declaration $\forall m, m', \#m, \#m'$ $\mathbf{class} \ c \cdots \{ \cdots m(\cdots) \{ \cdots \} \cdots m'(\cdots) \{ \cdots \} \cdots$ $\#m(\cdots) \{ \cdots \} \cdots \#m'(\cdots) \{ \cdots \} \}$ is in P $\Rightarrow \text{Set}(m, m', \#m, \#m') \text{size} = 4$
OVERRIDINGPROTECTEDMETHOD(P)	Protected methods can be overridden by public or protected methods $\forall m, \langle m, x^*, e \rangle \in_p c, \langle m, x^*, e \rangle \in_p c', c \leq_p c'$ or $\forall m, \langle m, x^*, e \rangle \in_p c, \langle m, x^*, e \rangle \in_p c', c \leq_p c'$
OVERRIDINGPUBLICMETHOD(P)	Public methods can be overridden only by public methods $\forall m, \langle m, x^*, e \rangle \in_p c, \langle m, x^*, e \rangle \in_p c', c \leq_p c'$
\in_p	Public method defined in class $\langle m, x^*, e \rangle \in_p c \iff \mathbf{class} \cdots \{ \cdots m(x^*) \{ e \} \cdots \} \in P$
$\widehat{\in}_p$	Protected method defined in class $\langle m, x^*, e \rangle \widehat{\in}_p c \iff \mathbf{class} \cdots \{ \cdots \#m(x^*) \{ e \} \cdots \} \in P$
\in_p^*	Public method lookup starting from c $\langle c, m, x^*, e \rangle \in_p^* c' \iff c' = \min\{c'' \mid \langle m, x^*, e \rangle \in_p c'', c \leq_p c''\}$
$\widehat{\in}_p^*$	Protected method lookup starting from c $\langle c, m, x^*, e \rangle \widehat{\in}_p^* c' \iff c' = \min\{c'' \mid \langle \#m, x^*, e \rangle \widehat{\in}_p c'', c \leq_p c''\}$

■ **Figure 21** Relations and predicates for PROTECTEDLITE.

$P \vdash$	$\langle E[o.m(v^*)], \mathcal{S} \rangle \hookrightarrow \langle E[o[[e[v^*/x^*]]]_{c'}], \mathcal{S} \rangle$ where $\mathcal{S}[o] = \langle c, \mathcal{F} \rangle$ and $\langle c, m, x^*, e \rangle \in_p^* c'$	[object-send]
$P \vdash$	$\langle E[\mathbf{self}(o, c).m(v^*)], \mathcal{S} \rangle \hookrightarrow \langle E[o[[e[v^*/x^*]]]_{c'}], \mathcal{S} \rangle$ where $\mathcal{S}[o] = \langle c, \mathcal{F} \rangle$ and $\langle c, m, x^*, e \rangle \widehat{\in}_p^* c'$ or $\langle c, m, x^*, e \rangle \in_p^* c'$	[self-send]
$P \vdash$	$\langle E[\mathbf{super}(o, c).m(v^*)], \mathcal{S} \rangle \hookrightarrow \langle E[o[[e[v^*/x^*]]]_{c''}], \mathcal{S} \rangle$ where $c \prec_p c'$, and $\langle c', m, x^*, e \rangle \widehat{\in}_p^* c''$ or $\langle c', m, x^*, e \rangle \in_p^* c''$, and $c' \leq_p c''$	[super-send]

■ **Figure 22** Message passing reductions for PROTECTEDLITE.

the object c and only looking at public methods. For both self-send and super-send, we use both lookups.

Super-sends use the same mechanism as self-sends, but the lookup starts from the superclass (c') of the class defining the method.

Selector mangling for self-sends sites. As we saw in Section 4.4, during compilation, all self-sends are mangled. Figure 23 formalizes this transformation. The expression $[m'/m]_{self}$ renames all the self-sends to m to m' , where m' is the mangled selector obtained from the ρ_{\square} hiding function. The scope of this hiding function is global.

Protected Modifier

$\mathbf{new} \ c \ [m'/m]_{self}$	$=$	$\mathbf{new} \ c$	<i>instance creation</i>
$\mathbf{let} \ x = e \ \mathbf{in} \ e' \ [m'/m]_{self}$	$=$	$\mathbf{let} \ x = e[m'/m]_{self} \ \mathbf{in} \ e'[m'/m]_{self}$	
$x \ [m'/m]_{self}$	$=$	x	<i>variable access</i>
$f \ [m'/m]_{self}$	$=$	f	<i>field access</i>
$(f=e) \ [m'/m]_{self}$	$=$	$(f=e[m'/m]_{self})$	<i>field assignment</i>
$\mathbf{nil} \ [m'/m]_{self}$	$=$	\mathbf{nil}	
$\mathbf{self}.m(e_i^*) \ [m'/m]_{self}$	$=$	$\mathbf{self}.m'(e_i^*[m'/m]_{self})$	<i>self/super sends</i>
$\mathbf{self}.n(e_i^*) \ [m'/m]_{self}$	$=$	$\mathbf{self}.n(e_i^*[m'/m]_{self}), \ \text{if } n \neq m$	
$\mathbf{super}.m(e_i^*) \ [m'/m]_{self}$	$=$	$\mathbf{super}.m'(e_i^*[m'/m]_{self})$	<i>can call protected methods</i>
$\mathbf{super}.n(e_i^*) \ [m'/m]_{self}$	$=$	$\mathbf{super}.n(e_i^*[m'/m]_{self}), \ \text{if } n \neq m$	
$e.m(e_i^*) \ [m'/m]_{self}$	$=$	$e[m'/m]_{self}.m(e_i^*[m'/m]_{self})$	<i>object-sends only call public methods</i>

- **Figure 23** Self and super-send call site renaming (also called selector mangling). Method names of object-sends are not renamed to protected method name mangling, while self and super sends are.

E Applicability outside Pharo: PROTDYN Python port

Python does not provide the same facilities as Pharo regarding the modification of the internal compiler directly from the language itself. We use Python decorators to mark methods as `@protected`, as illustrated in listing 2 with the `bar` method, and two hooks of Python's Meta-Object-Protocol to detect class creation.

The protected decorator registers a callback method — using the `__set_name__` hook¹ — that is automatically called when the owning class object is created. When this callback is triggered, the class is recompiled as well as its superclasses. The recompilation of a class is made of two major steps: the modification of all self-send sites and the installation of the non-protected methods under both the original selector and the mangled selector in the class.

¹ <https://docs.python.org/3/reference/datamodel.html#creating-the-class-object>

```
1 from protected import protected
2 # ...
3 class B(A):
4     @protected
5     def bar(self):
6         return 42
7
8     def raiseError(self):
9         return A().bar()
10
11     def sum(self):
12         return self.foo() + B().foo()
13
14     def foobar(self):
15         return super().foobar()
16
```

■ **Listing 2** Excerpt of the example in Figure 2 reimplemented in Python using a `@protected` decorator as protected modifier.

Finally, we register a callback on each class at the top of the inheritance tree to customize subclass creation using the `__subclass_init__` hook.² The customization of the subclass creation triggers the recompilation process for each new subclass created, even if those subclasses do not define protected methods. *Note:* the registration of this dedicated callback for each top class ensures that the subclasses customization process supports multiple inheritance.

Limitations: The recompilation process triggered by the `__set_name__` and `__subclass_init__` hooks only deals with methods that are declared in their classes. Our current implementation does not deal with dynamically added methods (monkey patching) on a class that defines protected methods in its inheritance tree. To address this issue, the implementation can rely on the change of the normal Python metaclass of each class at the top of the inheritance tree for a special Python metaclass that registers a callback on attribute injection in the class.

F Numerical results for benchmarks

The following table is showing the results of the three runtime performance experiments described in Section 5.3

² https://docs.python.org/3/reference/datamodel.html#object.__init_subclass__

Protected Modifier

#Experiment	Benchmarks	Without protected	Loaded	Used
1	Microdown	1880.8ms ± 0.4 ms	1880.3ms ± 0.3ms	1878.9ms ± 0.4ms
	Delta Blue	8286ms ± 5ms	8272ms ± 5ms	8197ms ± 6ms
	Richards	3519ms ± 4ms	3522ms ± 4ms	3522ms ± 4ms
	Compiler	4225ms ± 2ms	4211ms ± 2ms	4243ms ± 2ms
2	Microdown	788.2ms ± 0.7ms	797.1ms ± 0.6ms	797.2ms ± 0.6ms
	Delta Blue	4197ms ± 2ms	4219ms ± 5ms	4180ms ± 2ms
	Richards	2148ms ± 2ms	2134ms ± 2ms	2111ms ± 2ms
	Compiler	2891ms ± 2ms	2889ms ± 2ms	2921ms ± 2ms
4	Microdown	441.4ms ± 0.5ms	440.8ms ± 0.4ms	439.9ms ± 0.4ms
	Delta Blue	1275ms ± 3ms	1281ms ± 2ms	1283ms ± 3ms
	Richards	448.2ms ± 0.7ms	461.1ms ± 0.8ms	458.9ms ± 0.8ms
	Compiler	906.6ms ± 1.6ms	914.2ms ± 1.6ms	919.7ms ± 1.6ms

■ **Table 2** Average run times and standard interval with a confidence of (95%) for the different benchmarks: without using the protected modifier library, with the protected modifier library loaded and with the protected modifier library used