



Small-Amp: Test amplification in a dynamically typed language

Mehrdad Abdi¹ · Henrique Rocha² · Serge Demeyer¹ · Alexandre Bergel³

Accepted: 13 April 2022

© The Author(s) 2022

Abstract

Some test amplification tools extend a manually created test suite with additional test cases to increase the code coverage. The technique is effective, in the sense that it suggests strong and understandable test cases, generally adopted by software engineers. Unfortunately, the current state-of-the-art for test amplification heavily relies on program analysis techniques which benefit a lot from explicit type declarations present in statically typed languages. In dynamically typed languages, such type declarations are not available and as a consequence test amplification has yet to find its way to programming languages like Smalltalk, Python, Ruby and Javascript. We propose to exploit profiling information —readily obtainable by executing the associated test suite— to infer the necessary type information creating special test inputs with corresponding assertions. We evaluated this approach on 52 selected test classes from 13 mature projects in the Pharo ecosystem containing approximately 400 test methods. We show the improvement in killing new mutants and mutation coverage at least in 28 out of 52 test classes ($\approx 53\%$). Moreover, these generated tests are understandable by humans: 8 out of 11 pull-requests submitted were merged into the main code base ($\approx 72\%$). These results are comparable to the state-of-the-art, hence we conclude that test amplification is feasible for dynamically typed languages.

Communicated by: Carlo A. Furia

✉ Mehrdad Abdi
mehrdad.abdi@uantwerpen.be

Henrique Rocha
henrique.rocha@gmail.com

Serge Demeyer
serge.demeyer@uantwerpen.be

Alexandre Bergel
abergel@dcc.uchile.cl

¹ University of Antwerp – Flanders Make vzw, Antwerp, Belgium

² Loyola University Maryland, Baltimore, MD, USA

³ RelationalAI (USA) and Department of Computer Science, University of Chile (Chile), Santiago, Chile

Keywords Test amplification · Pharo smalltalk

1 Introduction

Modern software projects contain a considerable amount of hand-written tests which assure that the code does not regress when the system under test evolves. Indeed, several researchers reported that test code is sometimes larger than the production code under test (Daniel et al. 2009; Tillmann and Schulte 2006; Zaidman et al. 2011). More recently, during a large scale attempt to assess the quality of test code, Athanasiou et al. reported six systems where test code takes more than 50% of the complete codebase (Athanasiou et al. 2014). Moreover, Stack Overflow posts mention that test to code ratios between 3:1 and 2:1 are quite common (Agibalov 2015).

Test amplification is a field of research which exploits the presence of these manually written tests to strengthen existing test suites (Danglot et al. 2019a). The main motivation of test amplification is based on the observation that manually written test cases mainly exercise to the default scenarios and seldom cover corner cases. Nevertheless, experience has shown that strong test suites must cover those corner cases in order to effectively reveal failures (Li and Offutt 2016). Test amplification therefore automatically transforms test-cases in order to exercise the boundary conditions of the system under test.

Danglot et al. conducted a literature survey on test amplification, identifying a range of papers that take an existing test suite as the seed value for generating additional tests (Fraser and Arcuri 2012; Rojas et al. 2016; Yoo and Harman 2012). This culminated in a tool named DSPOT which represents the state-of-the-art in the field (Baudry et al. 2015; Danglot et al. 2019b). In these papers, the authors demonstrate that DSPOT is effective, in the sense that the tool is able to automatically improve 26 test classes (out of 40) by triggering new behaviors and adding valuable assertions. Moreover, test cases generated with DSPOT are well perceived by practitioners — 13 (out of 19) pull requests with amplified test have been incorporated in the main branch of existing open source projects (Danglot et al. 2019b).

Unfortunately, the current state-of-the-art for test amplification heavily relies on program analysis techniques which benefit a lot from explicit type declarations present in statically typed languages. Not surprisingly, previous research has been confined to statically typed programming languages including Java, C, C++, C#, Eiffel (Danglot et al. 2019a). In dynamically typed languages, performing static analysis is difficult since source code does not embed type annotation when defining variable. As a consequence test amplification has yet to find its way to dynamically-typed programming languages including Smalltalk, Python, Ruby, Javascript, etc.

In this paper, we demonstrate that test amplification is feasible for dynamically typed languages by exploiting profiling information readily available from executing the test suite. As a proof of concept, we present SMALL-AMP which amplifies test cases for the dynamically typed language Pharo (Black et al. 2010; Bergel et al. 2013); a variant of Smalltalk (Goldberg and Robson 1983). We argue that Pharo is a good vehicle for such a feasibility study, because it is purely object-oriented and it comes with a powerful program analysis infrastructure based on metalinks (Costiou et al. 2020). Pharo uses a minimal computation model, based on object and message passing, thus reducing possibilities to experiences biases due to some particular and singular language constructions. Moreover, Pharo has a growing and active community with several open source projects welcoming pull requests

from outsiders. Consequently, we replicate the experimental set-up of DSPOT (Danglot et al. 2019b) by including a quantitative and qualitative analysis of the improved test suite.

This paper is an extension of a previous paper presenting the proof-of-concept to the Pharo community (Abdi et al. 2019b). As such, we make the following contributions:

- *Small-Amp*, a test amplification algorithm and tool, implemented in Pharo Smalltalk. To the best of our knowledge this is the first test amplification tool for a dynamically typed language.
- Demonstrating the use of *dynamic type profiling* as a substitute for type declarations within a system under test.
- *Quantitative evaluation* of our test amplification for the Pharo dynamic programming language on 13 mature projects with good testing and maintenance practices. We repeated the experiment three times. For 28 out of 52 test classes we see an improvement in killing new mutants and consequently the mutation score. Our evaluation shows that generated test methods are focused (i.e. they do not overwhelm the developer) and all amplification steps are necessary to obtain strong and understandable tests.
- *Qualitative evaluation* of our approach by submitting pull requests containing amplified tests on 11 active projects. 8 of them ($\approx 72\%$) were accepted and successfully merged into the main branch.
- We contribute to *open science* by releasing our tool as an open-source package under the MIT license (<https://github.com/mabdi/small-amp>). The experimental data is publicly available as a replication package (<https://github.com/mabdi/SmallAmp-evaluations>).

The remainder of this paper is organised as follows. Section 2, provides the necessary background information on test amplification and the Pharo ecosystem. Sections 3 and 4 explain the inner workings of SMALL-AMP, including the use of dynamic profiling as a substitute for static type information. Section 5 discusses the quantitative and qualitative evaluation performed on 13 mature open source projects; a replication of what is reported by Danglot et al. (2019a). Section 6 enumerates the threats to validity while Section 7 discusses related work and Section 8 lists limitations and future work. Section 9 summarizes our contributions and concludes our paper.

2 Background

2.1 Test amplification

In their survey paper, Danglot et al. define test amplification as follows:

Test amplification consists of exploiting the knowledge of a large number of test cases, in which developers embed meaningful input data and expected properties in the form of oracles, in order to enhance these manually written tests with respect to an engineering goal (e.g., improve coverage of changes or increase the accuracy of fault localization). Danglot et al. (2019a)

Test amplification is a not replacement for other test generation techniques and should be considered as a complementary solution. The main difference between test generation and test amplification is the use of an existing test suite. Most work on test generation accept only the program under test or formal specifications and ignore the original test suite which is written by an expert.

```

1  SmallBank >> withdraw: amount
2      balance >= amount
3      ifTrue: [
4          balance := balance - amount.
5          ^ true ].
6      ^ false
7
8  SmallBankTest >> testWithdraw
9      | b |
10     b := SmallBank new.
11     b deposit: 100.
12     self assert: (b balance = 100).
13     b withdraw: 30.
14     self assert: (b balance = 70).
15
16 SmallBankTest >> testWithdrawAll
17     | b success |
18     b := SmallBank new.
19     b deposit: 30.
20     self assert: (b balance = 30).
21     success := b withdraw: 30.
22     self assert: success.
23     self assert: (b balance = 0).
24
25 SmallBankTest >> testWithdrawOnZero
26     | b success |
27     b := SmallBank new.
28     success := b withdraw: 30.
29     self deny: success.
30     self assert: (b balance = 0).

```

Listing 1 testWithdraw amplified into testWithdrawOnAll and testWithdrawOnZero

A typical test amplification tool is based on two complementary steps.

- (i) *Input amplification*. The existing test code is altered in order to force previously untested paths. This involves changing the set-up of the object under test, providing parameters that represent boundary conditions. Additional calls to state-changing methods of the public interface are injected as well.
- (ii) *Assertion amplification*. Extra assert statements are added to verify the expected output of the previously untested path. The system under test is then used as an oracle: while executing the test the algorithm inspects the state of the object under test and asserts the corresponding values.

The input amplification step is typically governed by a series of *amplification operators*. These operators represent syntactical changes to the test code that are likely to force new paths in the system under test. To verify that this is indeed the case, the amplification tool compares the (mutation) coverage before and after the amplification operator. It is beyond the scope of this paper to explain the details of mutation coverage; we refer the interested reader to the survey by Papadakis et al. (2019).

We illustrate the input and assertion amplification steps via an example based on SmallBank¹ and its test class SmallBankTest in Listing 1. In this example testWithdraw is the original test method while testWithdrawAll and testWithdrawOnZero are two new test methods derived from it. In the testWithdrawAll, the input amplification has changed the literal value of 100 with 30 (line 19), and the assertion-amplification step regenerated the assertions on the balance (line 20) and added a missing assertion on the status of the operation (line 22). The testWithdrawAll test method thus verifies the boundary condition of withdrawing by an amount equal to the balance. In the testWithdrawOnZero, on the other hand, an input amplifier has removed the call to the deposit: method in line 11. This test method now verifies the boundary condition that calling a withdraw: with an amount more than zero when the balance is zero is not allowed. This is illustrated by the extra assertions in line 29 and 30.

¹ Available at: <https://github.com/mabdi/smalltalk-SmallBank>

2.2 Pharo

Pharo [<http://www.pharo.org/>] is a pure object oriented language based on Smalltalk (Black et al. 2010; Bergel et al. 2013). It is dynamically typed; i.e. there are no type declarations for variables, parameters, nor return values statically, but dynamically, the environment enforces that all objects to have a type and only respond to messages part of the interface. It includes a run-time engine and an integrated development environment with code browsers and live debugging. Pharo users work in a live environment called *Pharo image* where writing code and executing it is tied seamlessly together.

Invoking a method in Pharo is called *message sending*. As a pure language, every action in Pharo is achieved by sending messages to objects. There are no predefined operators, like `+` or `-`, nor control structures like `if` or `while`. Instead, a Pharo program sends the message `#+` or `#-` to a number object, a `#ifTrue:ifFalse:` message to a boolean object, or the message `#whileTrue:` to a boolean returning block object. Any message can be sent to any object. In case the message is not part of the object interface, instead of a compile-time syntax error, the system raises a `MessageNotUnderstood` exception in runtime. Thus, when transforming test code, a test amplification tool should be attentive to not create faulty test codes.

Like Java, all ordinary classes inherit from the class `Object` and every class can add instance variables and methods. Unlike Java, all instance variables are private and all methods are public. Pharo encourages programmers to write short methods with intention revealing names so that the code becomes self explanatory.

Protocols Pharo, and Smalltalk in general, features *protocols* to organize the methods defined in classes. The notion of protocol is a tag of a method and it acts like a metadata provided by the integrated development environment. As such, classifying a method under a particular protocol has no impact on the behavior.

Since all instance variables are private in Pharo, in order to make them accessible by the external world, accessor methods should be provided which are typically grouped into the protocol `accessing`. In a similar vein, all methods used to set the content of an object upon initialization are grouped into the protocol `instance creation`. Long lived classes who evolve over time, use the `deprecated` protocol, signalling that these methods will be removed from the public interface in the near future. And while all methods are public, Pharo uses the protocol `private` to mark methods which are not expected to be used from the outside. However, as we mentioned earlier, protocols are a tag and Pharo does not block an access to a private method.

The most similar concepts to protocols in other languages are naming conventions, annotations and also access modifiers. For instance, a Java equivalent for methods in `accessing` protocol is following a naming convention like `setVar()` and `getVar()`. In a similar vein, Java uses `@Deprecated` annotation to identify the deprecated methods. An equivalent for methods in *private* protocol in Python is the naming convention of using underscore before the name of private methods, but Java uses *access modifiers* for this purpose.

2.3 Coding Conventions in Dynamically Typed Languages

In this section, we describe typical coding conventions that are used by programmers to compensate for the lack of type declarations. When we transform code (like we do when amplifying tests), special care must be taken to adhere to such coding conventions otherwise the code will look artificial and will decrease chances to be adopted by test engineers. Our

perspective comes from Pharo / Smalltalk (as documented in Ducasse 2019), but similar coding conventions must be adhered to when amplifying tests in Python, Ruby or Javascript.

Untyped Parameters In dynamically typed languages, when defining a method which accepts a parameter, the type of the parameter is not specified. However, it is a convention to name the parameter after the class one expects or the role it takes. This is illustrated by the code snippet in Listing 2. Line 1 specifies that this is a method `drawOn:` defined on the class `Morph` which expects one parameter. The parameter itself is represented by an untyped variable `aCanvas` however the name of the variable suggests that the method expects an instance of the class `Canvas`, or one of its subclasses. Line 7 on the other hand specifies that the method `withdraw:` expects one parameter and its role is to be an amount. There is no clue on the type of the parameter (integer, longinteger, float, ...); all we can infer from looking at the code is that we should be allowed to pass it as an argument when invoking the messages `>=` (line 8) and `-` (line 10) on `balance`.

⇒ *When passing a parameter to a method, a test amplification tool has no guaranteed way of knowing the expected type. The name of the parameter only hints at the expected type, hence during input assertion special care must be taken.*

No Return Types In dynamically typed languages, there is no explicit declaration of the return type of a method. In Pharo, all computation is expressed with objects sending messages and a message sends always returns an object. By default a method returns the receiving object, which is the equivalent of the `void` return type in Java. However a program can explicitly return another value using `^` followed by an expression.

This is illustrated in Listing 3, showing the methods `printOn:` (displays the receiver on a given stream) and `printString` (which returns a string representation of the receiver). `printOn:` is the equivalent of a `void` call thus returns the receiving object; however the method is declared on `Object` so the receiver object can be anything. `printString` on the other hand returns the result of sending the message `contents` to `aStream`. The exact type of what is returned is difficult to infer via static code analysis. Smalltalk programmers would assume that the return type is a `String` because of the intention revealing name of the method. However, there is no guarantee that this is indeed the case. Thus, when a test amplification tool manipulates the result of a method, it cannot easily infer the type of what is returned.

⇒ *The lack of explicit return types makes it hard to manipulate the result of a method call while ensuring that no `MessageNotUnderstood` exceptions will be thrown.*

Different Return Types In addition to the lack of return type declarations, it is also possible to write a method that can return different types of values. For example, in Listing 4 the

<pre> 1 Morph >> drawOn: aCanvas 2 aCanvas fillRectangle: self bounds 3 fillStyle: self fillStyle 4 borderstyle: self BorderStyle. 5 6 </pre>	<pre> 7 SmallBank >> withdraw: amount 8 balance >= amount 9 ifTrue: [10 balance := balance - amount. 11 ^ true]. 12 ^ false </pre>
--	---

Listing 2 Examples of naming conventions for untyped parameters

```

1  Object >> printOn: aStream
2  | title |
3  title := self class name.
4  aStream nextPutAll:
5      (title first isVowel
6          ifTrue: [ 'an ' ]
7          ifFalse: [ 'a ' ] );
8      nextPutAll: title
9
10 Object>>printString
11 | aStream |
12 aStream := WriteStream on:
13     (String new: 16).
14 self printOn: aStream.
15 ^aStream contents

```

Listing 3 Example of a void method (left) or function method (right)

method `someMethod`: can return an instance of the classes `Integer`, `Boolean` or `Object` (the default return value is `self`).

As a result, removing the return operator (a common mutation operator) will not cause a syntax error yet may cause a change in the return type of a methods. For example, in Listing 4 the method `width` (lines 5 and 6), if the return operator is removed in the mutation testing, the type of the return value will be converted from a number to a `Shape` object.

⇒ *Methods in dynamically typed languages can return various types. Test amplification tools must be aware that a small change in the code may lead to changes in the returned type. Consequently, assertions verifying the result of a method call must be adapted.*

Accessor Methods In Pharo, all instance variables are private and only accessible by the object itself. If one wants to manipulate the internal state of an object one should implement a method for it, as illustrated in Listing 5 which shows the setter method `x:` and the getter method `x`. In Pharo, such accessor methods are typically collected in the protocol accessing and are a convenient way for programmers to look for ways to read or write the internal state of an object.

Such accessor methods are especially relevant for all test generation algorithms (Fischer et al. 2020). For test amplification in particular, the setter methods are necessary in the input amplification step to force the object into a state corresponding to a boundary condition. The getter methods are necessary in the assertion amplification step to verify whether the object is in the appropriate state. However, there is no explicit declaration for the type of the parameter passed to the setter method `x:` nor for the type to be returned by the getter method `x`.

⇒ *When manipulating the state of an object one cannot rely on type declarations to infer which parameter to pass to a setter method and which result to expect from a getter method.*

Pass-by-reference In dynamic languages including Pharo, when sending messages, all arguments are passed by reference. This may imply that sometimes the state is changed and sometimes it is not. This is illustrated by the method `r` in Listing 6, which returns the radius

```

1  Example >> someMethod: anInt
2  anInt = 1 ifTrue: [ ^ 1 ]
3  anInt = 0 ifTrue: [ ^ false ]
4
5  Shape >> width
6  ^ width

```

Listing 4 Examples of a changing the return type

```

1  Point >> x
2      ^x
3  Point >> x: anInteger
4      x := anInteger

```

Listing 5 Example of a getter (left) and a setter method (right)

in polar coordinates. This involves some calculation (the invocation of `dotProduct:`) which passes the receiver object as a reference. There is no “pass-by-value” type declaration for `dotProduct:`, so one cannot know whether the internal state is changed or not. If `dotProduct:` does not alter the internal state it may be used as a pure accessor method during assertion amplification anywhere in the test. However, if the accessor method does change the internal state the order in which the accessor methods are called has an effect on the outcome of the test.

⇒ *The pass-by-reference parameter passing makes it difficult to distinguish pure accessor methods. Pure accessor methods can be inserted anywhere during assertion amplification, for accessor methods changing the internal state one must take into account the calling order.*

Cascading Listing 7 shows the archetypical `Hello World` example. Line 2 specifies that this is a method `helloWorld` defined on a class `HelloWorld`. Line 4 and 6 each sends the message `cr` (a message without any parameters) to the global variable `Transcript` which emits a carriage return on the console. Line 5 sends the message `show:` with as parameter the string ‘hello world’ to the global variable `Transcript` which writes out the expected message.

However, a Pharo programmer would never write this piece of code like that. When a series of messages is being sent to the same receiver, this can be expressed more succinctly as a cascade. The receiver is specified just once, and the sequence of messages is separated by semi-colons as illustrated on lines 7–10.

Instance Creation Cascading is frequently used when creating instances of a class as illustrated by the `createBorder` example in the left of Listing 8. In line 2 it creates a new `SimpleBorder` object and then initialises the object with color blue (line 3) and width 2 (line 4). During input amplification we need to change the internal state of the object under test, hence it is tempting to inject extra calls in such a cascade. However, because we cannot distinguish between state-changing and state-accessing methods, we risk injecting errors. The code snippet to the right illustrates that injecting an extra `isComplex` call (a call to a state-accessing method) at the end of the cascade erroneously returns a boolean instead of an instance of `SimpleBorder`. This will eventually result in a run-time type error via a `messageNotUnderstood` exception when the program tries to use the result of `createBorderErroneous`.

⇒ *When injecting additional calls during instance creation, one runs the risk of returning an inappropriate value.*

```

1  Point >> r
2      ^(self dotProduct: self) sqrt

```

Listing 6 Is `r` a pure accessor method that does not alter the internal state?

<pre> 1 HelloWorld >> helloWorld 2 Transcript cr. 3 Transcript show: 'hello world'. 4 Transcript cr. 5 </pre>	<pre> 6 HelloWorld >> helloWorldCascading 7 Transcript 8 cr; 9 show: 'hello world'; 10 cr. </pre>
---	---

Listing 7 A sequence of messages sent to the same receiver object (left) is written as a cascade (right)

Like most dynamically typed languages, Pharo has a lot of coding conventions. When transforming code (for instance, when amplifying tests) we must adhere to these conventions. However, the lack of explicit type information hinders static analysis, needed to identify relevant code constructs.

3 SMALL-AMP Design

In this section, we explain the design of the SMALL-AMP which is an adaption and extension of DSPOT (Baudry et al. 2015; Danglot et al. 2019b) for the Pharo ecosystem. DSPOT is an opensource² test amplification tool to amplify tests for Java programs. Our SMALL-AMP implementation is also publicly available³ on GitHub.

3.1 Main Algorithm

The main amplification algorithm is presented in Algorithm 1 and represents a search-based test amplification algorithm. The algorithm accepts a class under test (CUT) and its related test class (TC) and returns the set of amplified test methods (ATM). In addition, the algorithm needs a set of input amplification operators (AMPS) and is governed by a series of hyperparameters:

- $N_{\text{iteration}}$ – This parameter specifies the number of iterations and shows the maximum number of transformations on a test input. The default value for this parameter is 3.
- $N_{\text{maxInputs}}$ – This parameter specifies the maximum number of generated test inputs that algorithm keeps. It discards other test inputs. The default value for this parameter is 10.

Initially, the code of CUT and TC is instrumented to allow for dynamic profiling. The test class is executed, all required information is collected and then the instrumentation is removed again. This extra information including the type information allows us to perform input amplification more efficiently and circumvent the lack of type information in the source code (line 2). We discuss about the profiling in Section 4.1.

The main loop of the algorithm amplifies all test methods one by one (line 3). V is the set of test inputs, thus test methods without assertion statements. In the beginning, V has only one element which is obtained from removing assertion statements in the original test method (line 4). U is the set of generated test methods which are generated by adding new assertion statements to the elements in V (lines 5). Then the coverage is calculated using the

²<https://github.com/STAMP-project/DSPot>

³<https://github.com/mabdi/small-amp>

<pre> 1 TestBorder >> createBorder 2 ^ SimpleBorder new 3 color: Color blue; 4 width: 2. "returns self" 5 </pre>	<pre> 6 TestBorder >> createBorderErroneous 7 ^ SimpleBorder new 8 color: Color blue; 9 width: 2; 10 isComplex. "returns a boolean" </pre>
--	--

Listing 8 Injecting extra statements may result in type errors

generated test methods accumulated in U and the tests increasing the coverage are added to the final result. SMALL-AMP uses mutation score as a coverage criteria (line 6).

Algorithm 1 SMALL-AMP amplification algorithm.

```

input : CUT: class-under-test
input : TC: original test class
input : AMPS: a set of input amplification operators
input : hyperparameters  $\{N_{\text{iteration}}, N_{\text{maxInputs}}\}$ 
output: ATM: set of amplified test methods
1   $ATM \leftarrow \{\}$ ;
2   $extraInfo \leftarrow \text{profileCollect}(CUT, TC)$ ;
3  for each  $t \in TC$  do
4     $V \leftarrow \{\text{removeAssertions}(t)\}$ ;
5     $U \leftarrow \text{amplifyAssertions}(V)$ ;
6     $ATM \leftarrow ATM \cup \{x \in U \mid x \text{ improves mutation score}\}$ ;
7    for  $i \leftarrow 0$  to  $N_{\text{iteration}}$  do
8       $TMP \leftarrow \{\}$ ;
9      for each  $amp \in AMPS$  do
10        $TMP \leftarrow TMP \cup \text{amplifyInputs}(amp, V, extraInfo)$ ;
11     end
12      $V \leftarrow \text{reduce}(TMP, N_{\text{maxInputs}})$ ;
13      $U \leftarrow \text{amplifyAssertions}(V)$ ;
14      $ATM \leftarrow ATM \cup \{x \in U \mid x \text{ improves mutation score}\}$ ;
15   end
16 end
17  $ATM \leftarrow \text{improveReadability}(ATM)$ ;
18 return  $ATM$ 

```

In the inner loop of the algorithm (lines 7 to 15), SMALL-AMP generates additional tests by repeating the following steps $N_{\text{iteration}}$ times:

1. SMALL-AMP applies different input amplification operators on V (the current test inputs) to create new variants of test methods accumulated in the variable TMP (line 10). We discuss input amplification in Section 3.2.
2. SMALL-AMP reduces TMP by keeping only $N_{\text{maxInputs}}$ of current inputs and discarding the rest (line 12). We discuss input reduction in Section 4.2.
3. SMALL-AMP injects assertions on the remaining test inputs in V and stores the result in U (line 13). We discuss assertion amplification in Section 3.3.
4. SMALL-AMP selects all test methods in U that increase mutation score and adds them to the final result ATM (line 14). We discuss about test selection in Section 3.4.

After both loops have terminated, **SMALL-AMP** applies a set of post-processing steps to increase the readability of the generated tests (line 17). We discuss these steps in Section 4.3. Algorithm 1 is heavily inspired by **DSPOT**, but not entirely the same. In other words, we have added a pre-process step (line 2) to collect the necessary information about CUT and TC before entering the main loop. We also have added a post-processing step (line 17) to make the output more readable. We discuss about extras to **DSPOT** algorithm in Section 4.

3.2 Input Amplification

During input amplification, existing test code is altered to force previously untested paths. Input amplification involves changing the set-up of the object under test, passing arguments which represent boundary conditions. Additional calls to state-changing methods of the public interface are injected as well. Such changes are bound to fail the original assertions of TC, therefore **SMALL-AMP** removes all assertions from a test t in TC.

The test code itself is transformed via a series of *Input Amplification Operators*. These change the code in such a way that they are likely to force untested paths and cover boundary conditions. Input amplification operators are based on the genetic operators introduced in Evolutionary Test Classes (Tonella 2004). Below we explain the *Input Amplification Operators* adopted from **DSPOT**.

Amplifying Literals This input amplifier scans the test input source code to find literal tokens (numbers, booleans, strings). Then it transforms the literal to a new literal based on its type according to Table 1. For example, test input shown in Listing 9 is transformed into `testVectorGreaterL` by manipulating the second element from the literal array.

Amplifying Method Calls This input amplifier scans the test input source code to find the method invocations on an object. Then it transforms the source code by *duplicating* or *removing* the method invocations. It also *adds new method invocations* on the objects. If the method requires new values as arguments, the amplifier creates new objects. For primitive parameters, a random value is chosen from the profiled values. For object parameters, the default constructor is used i.e it creates a new instance by sending `#new` message to the class. **SMALL-AMP** ignores *private* and *deprecated* methods (regarding to their protocol) when it adds a new method call. The type information required do safely apply these transformations is obtained in the profiling step explained in Section 4.1 – p. 13.

Table 1 Transformations in literal amplification

Type	Transformation
Numbers	0,
	increased and decreased values (+1 and −1),
	doubled and halved values (×2 and ÷2),
	negated value (× −1)
Booleans	replacing with an existing number from the test body
	negate via <i>not</i>
Strings	add a new random character to a random position
	remove a character randomly
	change a character randomly
	replace by a random string in the same size

<pre> 1 testVectorGreater 2 u w 3 u := #(-1 0 1) asPMVector. 4 w := u > 0. </pre>	<pre> 5 testVectorGreater_L 6 u w 7 u := #(-1 1 1) asPMVector. 8 w := u > 0. </pre>
--	--

Listing 9 Example Literals Amplification Operator (line 3 vs. line 7)

3.3 Assertion Amplification

During the assertion amplification step, we inject assertion statements which verify the state of the object under test. The object under test is then used as an oracle: while executing the test the algorithm inspects the state of the object under test and asserts the corresponding values. The assertion amplification step is based on Regression Oracle Checking (Xie 2006).

Note that assertion amplification is applied twice during the amplification algorithm (Algorithm 1, in lines 5 and 13). There are two reasons for this seemingly redundant design. (1) We assure that the original test method is assertion amplified as well. Since the test inputs are reduced in line 12, there is a possibility that the original test method is discarded and never reaches the assertion amplification in line 13. (2) We can run only assertion amplification by setting the value of $N_{iteration} = 0$. This way no new tests will be generated, but existing tests may become stronger because they check more conditions.

Observing State Changes Via Object Serialisation SMALL-AMP manipulates the test code and surrounds each statement with a series of what we call “observer meta-statements” (see Listing 10). Such meta-statements include a surrounding block to capture possible exceptions (lines 19–20 and 24–25) and calls to observer methods to capture the state of the receiver (line 17 and line 18) and the return value (line 18 and line 23). When necessary, temporary variables are added to capture intermediate return values (tmp1 on line 21 and line 23).

After manipulating the test method, SMALL-AMP runs the test to capture the values by the observer methods. SMALL-AMP serializes objects by capturing the values from its accessor methods. If the return value of an accessor method is another object, it recursively repeats the object serialization up to $N_{serialization}$ times. $N_{serialization}$ is a configurable value (default value is 3). The output of this step of assertion amplification is a set of trace logs which reflect the object states.

<pre> 1 testDeposit 2 b 3 b := SmallBank for: 'JDoe'. 4 b deposit: 100. 5 6 7 8 9 10 11 12 13 </pre>	<pre> 14 testDeposit_instrumented 15 b tmp1 16 [b := SmallBank for: 'JDoe'. 17 self observe: SmallBank. 18 self observeRetVal: b. 19] on: Error do: [:ex 20 self observeException: ex]. 21 [tmp1:= b deposit: 100. 22 self observe: b. 23 self observeRetVal: tmp1. 24] on: Error do: [:ex 25 self observeException: ex] </pre>
--	---

Listing 10 Injection of observer meta-statements

Identifying Accessor Methods SMALL-AMP relies on the Pharo/Smalltalk coding conventions and therefore selects methods if they belong to protocols `#accessing` or `#testing` or when their name is identical to one of the instance variables. From the selected methods, all methods lacking an explicit return statement and all methods in the protocols `#private` or `#deprecated` are rejected and the remaining are considered as accessors.

Preventing Flaky Tests Via Trace Logs A flaky test is a test that may occasionally succeed (green) or fail (red). This may happen if the test is asserting a non-deterministic value. SMALL-AMP tries to detect non-deterministic values before making assertions on them. The assertion amplification module, repeats collecting the trace logs for $N_{\text{flakiness}}$ (default value is 10) times. Then it compares the observed values. If a value is not identical between all collected logs, SMALL-AMP marks it as non-deterministic.

Recursive Assertion Generation Based on the type of the observed value, zero, one or more assert statements are generated. If the type is a variant of collection or an object, which include other internal values, the assertion generator uses a recursive method to build valid assertion statements. For non-deterministic values, the value is not asserted and only its type is asserted. The output of the assertion amplification step is a passing (green) test with extra assertions.

Intended Values Versus Actual Values During assertion amplification, the assertion statements should the expected value which is deduced from an oracle. We assume that the current implementation of the program is correct, and therefore we deduce the oracle from the current state of the object under test. However, when there is a defect in the method under test, the generated assertions would verify against an incorrect oracle. This is an inherent limitation for both DSPOT and SMALL-AMP, inherited from Regression Oracle Checking (Xie 2006).

Example Listing 11 shows an example of a trace log collected by line 22 from Listing 10 (left) and its recursive assertion statements (right). In this example, we point out that the method `timestamp` is an accessor method in `SmallBank` class which returns a

```

1  b:
2  type → SmallBank,
3  accessors:
4    balance:
5      flaky → false,
6      value → 100
7  timestamp:
8    flaky → true,
9    value → 1624
10 user:
11 type → SmallBankUser,
12 accessors:
13   name:
14     flaky → false,
15     value → 'JDoe'

16
17 testDeposit_withAssertions
18 | b tmp1 |
19 b := SmallBank for: 'JDoe'.
20 "... "
21 tmp1 := b deposit: 100.
22 self assert: b class equals: SmallBank.
23 self assert: b balance equals: 100.
24 self assert: b timestamp class equals:
25   Integer. "flaky"
26 self assert: b user class equals:
27   SmallBankUser.
28 self assert: b user name equals: 'JDoe'
29 "... "
```

Listing 11 An example of a trace log and its assertion statements

timestamp value. Since this value differs in different executions, it has been marked as a flaky value (line 8) hence only its type is asserted (line 24).

3.4 Test Selection – Prefer Focussed Tests

During each iteration of the inner loop (lines 7 to 15 in Algorithm 1 – p. 11) **SMALL-AMP** generates $N_{\text{maxInputs}}$ new tests with their corresponding assertions. In the test selection step (lines 6 and 14) the algorithm selects those tests which kill mutants not killed by other tests.

First of all, **SMALL-AMP** performs a mutation testing analysis on CUT and TC and creates a list of live and uncovered mutants. Then **SMALL-AMP** selects those test methods from U (the set of amplified test methods) which increase the mutation score, thus killing a previously live or uncovered mutant. If multiple tests are killing the same mutant, the shortest test is chosen. If there are multiple short tests, the test with the least changes is chosen. In the **DSPOT** paper, a similar heuristic is chosen, which the authors refer to as *Focused Test Cases Selecting*.

4 SMALL-AMP Extras Compared to DSPOT

While the design of **SMALL-AMP** was inspired by **DSPOT**, the lack of explicit type information forced us to make major changes but also permitted us to make improvements. This section describes additional and diverging aspects of **SMALL-AMP** compared to **DSPOT**.

4.1 Dynamic Profiling to Collect Type Information

At the very beginning of the main algorithm (Algorithm 1 line 2), dynamic type profiling is done only once by executing the original test methods and observing the actual type information of variables.

In dynamically typed languages like Pharo, type annotations are not provided in the source code. So, performing static analysis which depend on types are challenging. In the context of **SMALL-AMP**, the most important step that relies on static code analysis is input amplification. The other steps are either based on dynamic analysis like assertion amplification, or depend on a third-party library such as selection based on mutation-testing.

In input-amplification, we can group operators into two classes as:

1. *Type sensitive operators*. These operators heavily depend on the type information and without type information they are ineffective or impossible. An important type sensitive input amplifier in **SMALL-AMP** is method call addition. The types of variables defined in a test method must be inferred when adding a valid method call. In addition, it needs the type information of parameters in the newly called method.
2. *Type insensitive operators*. These are all operators that are still applicable without the type information. An example is the operators amplifying literals. These operators are easy to adapt to a dynamic language because literals are distinguishable from a token representation of the source code.

To obtain accurate type information we rely on the presence of manually written tests, which should be representative for the normal behaviour of the program under test. We exploit profiling tools (commonly available in modern program environments) to extract accurate type information from the variables present in the program. The profiler is configured to attach hooks to the relevant elements in the code. When these important code

elements are executed, the hooks are triggered, the profiler reads the information from the program state and logs it.

In **SMALL-AMP**, we rely on two distinct profilers:

- A *Method-proxy* profiler, which collects the type of parameters in Class-Under-Test methods.
- A *Metalink* based profiler which collects the type of variables in the test methods

To apply test amplification to other dynamically typed languages one needs comparable profiling technology. Some languages provide reflexive facilities that can be exploited. Python metaclasses for example allow one to transparently hook into the code proxy objects similar to the method-proxies adopted in **SMALL-AMP**. If such reflexive facilities are not available, one can resort to the debugger APIs to inspect values of variables at run-time.

Profiling by Method-Proxies For gathering the type of parameters in methods, **SMALL-AMP** uses method proxies (Ducasse 1999; Peck et al. 2015). Proxies are methods wrapping the methods in the class under test and trigger instead of the original methods. They first log the arguments and then pass the control to the original method (Listing 12).

The main drawback of the Method-proxy profiler is that when a method is not covered by the test class, it will not be profiled. **SMALL-AMP** reports the list of such uncovered methods as one of its outputs. Using this report, a developer can decide to add new tests for uncovered methods, make them private (using an adequate protocol / method tag), or remove them.

Profiling by Metalinks Pharo provides Metalinks as a fine-grained behavioral reflection solution (Denker et al. 2007; Costiou et al. 2020). For collecting the type of variables in the test method, **SMALL-AMP** uses Metalinks.

A metalink contains an action to perform which is defined by providing a meta-object, a selector, and also a control. Metalinks can be installed on one or more nodes in the abstract syntax tree. Listing 13 shows how metalink is defined and installed on all variable nodes in the test method.

Line 1 to 5 shows how Metalink is initialized. It says that after execution the AST node containing this link, the method `logNode:context:object:` will be called with the following arguments:

- `node`: The static representation of the AST node. It is used to get information such as name and the position in the code.
- `context`: The context of execution including dynamic values of the variables and stack. It is used to access to the values of temporary variables.
- `object`: The state of the object on which the metalink is installed (in this case the test class). It is used to access to the values of instance variables.

In lines 7 to 10, all variable nodes in the test method are selected and then the link is installed on them. After installing the metalinks, the test method is executed. When the execution passes each variable node, the metalink is triggered and the logger method is called. The logger method extracts the type information from the context, logs them and

```
1 ProfilingProxy >>> run: aSelector with: anArray in: aReceiver
2   self logCalled: aSelector withArguments: anArray inType: aReceiver.
3   ^ aReceiver withArgs: anArray executeMethod: method.
```

Listing 12 Wrapper method to log the types of the parameters

```

1  link := MetaLink new
2    metaObject: self;
3    control: #after;
4    selector: #'logNode:context:object:';
5    arguments: #(node context object).
6
7  nodes := testMethod ast allChildren
8    select: #isVariable.
9  nodes do: [:node |
10    node link: link ]

```

Listing 13 Defining a metalink to log the variable node type after execution

returns. Then, the execution on the test method continues until the end or another metalink is triggered.

How the Collected Data is Used The collected data from each profiler is stored as a dictionary object mapping the identifier of the profiled data to its type and a list of sample values (only for primitive types). In SMALL-AMP, there are two dictionaries, for the type of method parameters and the type of variable nodes. During the input amplification, when type information is needed, the corresponding dictionary is consulted.

4.2 Test Input Reduction

The input amplification step quickly produces a large number of new test inputs with the inner loop of Algorithm 1 – p. 11 – lines 7 to 15. For instance, if the number of inputs in the first iteration is $|v|$, this number in the second iteration grows to $|v| \times |v|$, and in iteration i reaches $|v|^i$. We refer to this problem as *test-input explosion*. Since the number of test inputs grows exponentially, either the number of transformations ($N_{iteration}$) needs to be chosen as small values for being feasible to try all generated test input, or we need to reduce the number of inputs by using a heuristic to select a limited number of them.

SMALL-AMP uses a random selection heuristic which maximises diversity in order to select a maximum number ($N_{maxInputs}$) of test inputs. This selection is different from selection by mutation score (Section 3.4); we name it *reduction*.

SMALL-AMP reduction considers two techniques:

- *a competitive selection*. a portion of test inputs (by default half of $N_{maxInputs}$) are selected completely randomly from the output of all input amplifiers.
- *a balanced selection* in the remaining portion, SMALL-AMP assures that all input amplifiers are contributed by selecting from their outputs regarding an assigned weight. In SMALL-AMP, all input amplifiers are assigned a weight (it is 1 by default for all amplifiers). This maintains a diversity in the selected test inputs.

Why Diversity is Important? Each input-amplifier algorithm performs transformations based on different considerations. As a result, the number of generated tests is different for input amplifiers. If the test inputs are selected purely random, the result will be dominated by generated tests from amplifiers generating more outputs. Therefore, we need to have a balance between the outputs from each amplifier.

As an example, we compare the number of new test inputs from a *statement-removal amplifier* and a *statement-addition amplifier*. The former has a $O(S)$ complexity where S is the number of statements in the test method. It means that if the number of lines in the test is increased, the number of new test inputs generated by this input-amplifier shows a linear increase. However, the latter has a $O(S * M)$ complexity where M is the number of methods in the class under test. It means that the increase in the outputs depends on not only the numbers of statements in the test, but also the number of methods in the class-under-test.

Now, if we select a number of generated test methods randomly, the outputs from the latter operator is more likely to be selected; so the result will be dominated by the result from the second input-amplifier.

4.3 Improving Readability Via Post-Processing

In order to make the generated tests more readable, SMALL-AMP adds a few steps after finishing the main loop of the algorithm (line 17 in Algorithm 1 – p. 11). These steps do not have any effect on the mutation score of the amplified test suite; they only make the test cases more readable for SMALL-AMP users.

Assertion Reduction As described in Section 3.3, SMALL-AMP generates all possible assertions for all observation points. Consequently, the generated test methods easily include hundreds of new assertions most of which appear redundant. The assertion reducer is a post-processing step that discards all assertions that do not affect the mutation score.

Each amplified test method encompasses the identifier of all newly killed mutants. SMALL-AMP surrounds all assertion statements by exception handling blocks to catch exceptions, especially `AssertionFailure` raised from the assertion statements. Then, the mutation testing framework is run using the newly killed mutants only. When an `AssertionFailure` is caught, the identifier of the assertion is logged as important. Finally, SMALL-AMP keeps only important assertions and remove all other assertion statements.

In some cases, an assertion may call an *impure* accessor methods, i.e. an accessor method that alters the internal state of the object. When such assertions are removed, some of the next assertions may fail. SMALL-AMP runs each test method after removing unnecessary assertions, to confirm that they remain green and the mutants are still killed by the test. If the confirmation failed, the assertion reduction is not successful and all removed assertions are reinserted.

Comply with Coding Conventions Before processing a test method, SMALL-AMP breaks complex statements (chains of method invocations and cascades) into an explicit sequence of message sends to permit observing state changes (see Listing 10 – p. 14). This is necessary to observe state changes during assertion amplification. In this post-process step, SMALL-AMP cleans up all unused temporary variables, and chooses a better name for the remaining variables based on the type of the variable. If possible, it reconstructs message chains and cascades to make the source code more readable and conform to Pharo coding conventions.

5 Evaluation

To evaluate SMALL-AMP, we replicated the experimental protocol introduced for DSPOT (Danglot et al. 2019b). We adopted a qualitative experiment by sending pull-requests in GitHub for evaluating whether the generated tests are relevant to the developers or not (RQ1). Next, we use a quantitative experiment to evaluate the effectiveness of SMALL-AMP (RQ2, RQ3 and RQ4). The order of RQ1 to RQ4 is exactly the same order in Danglot et al. (2019b) to facilitate the comparing and it does not reflect the importance of the research questions. In RQ5, we make a detailed comparison of our results versus the ones in the original experiment. Finally, in RQ6, we report the time cost of the running SMALL-AMP, with

special attention to the performance penalty induced by the additional steps (profiling and oracle reduction).

- RQ1 Pull Requests.** *Would developers be ready to permanently accept amplified test cases into the test repository?* We create pull-request on mature and active open-source projects in the Pharo ecosystem. We propose the improvement as a pull request on GitHub, comprising improvements on an existing test (typically due to assertion amplification) or new tests (typically the result of input plus assertion amplification). We interpret the statements where extra mutants are killed to provide a manual motivation on why this pull request is an improvement. The main contributors then review, discuss and decide to merge, reject or ignore the pull request. The ratio of accepted pull requests gives an indication of whether developers would permanently accept amplified test cases into the test repository. More importantly, the discussions raised during the review of the pull request provides qualitative evidence on the perceived value of the amplified tests.
- RQ2 Focus.** *To what extent are improved test methods considered as focused?* We assess whether the amplified tests don't overwhelm developers, by assessing how many extra mutants the amplified tests kill. Ideally, the amplified test method kills only a few extra mutants as then we consider the test *focussed* (cfr. Section 3.4 – p. 16). We present and discuss the proportion of focused tests out of all proposed amplified tests. An amplified test case is considered focus if, compared to the original, at least 50% of the newly killed mutants are located in a single method.
- RQ3 Mutation Coverage.** *To what extent do improved test classes kill more mutants than developer-written test classes?* We assess whether the amplified tests cover corner cases by using a proxy — the improvement in mutation score via the mutation testing tool MUTALK (Wilkinson et al. 2009). We first run MUTALK on the original class under test (CUT) as tested by the test class (TC) to compute the original mutation score. We distinguish between strong tests and weaker tests, by splitting the set of test classes in half after sorting according to the mutation score. Next, we amplify the test class and compute the new mutation score. We report the relative improvement (in percentage).
- RQ4 Amplification Steps.** *What is the contribution of input amplification and assertion amplification (the main steps in the test amplification algorithm) to the effectiveness of the generated tests?* Here as well we use mutation score as a proxy for the added value of both the input and assertion amplification step and here as well we distinguish between strong and weak test classes. Therefore, we compare the relative improvement (in percentage) of assertion amplification against the relative improvement of input and assertion amplification combined. We report separately which amplification operators have the most impact, paying special attention to the ones which are sensitive to type information.
- RQ5 Comparison.** *How does SMALL-AMP compare against DSPOT?* To analyse the differences in result between SMALL-AMP and DSPOT, we compare the qualitative and quantitative results reported in the DSPOT paper against the results we obtained for RQ1 to RQ4.
- RQ6 Time Costs.** *What is the time cost of running SMALL-AMP, including its steps?* To study the applicability of SMALL-AMP, we analyse the time cost of all runs in the quantitative analysis. We compare the relative time cost of each step, paying special attention to the extra overhead of profiling and oracle reduction.

5.1 Dataset and Metrics

Selecting a Dataset Firstly, we collected some candidate projects under test from different sources: (1) We looked at the projects used in a recent paper focusing on testing in Pharo (Delplanque et al. 2019). (2) We looked at the projects introduced in “Innovation Technology Awards” section of ESUG conference from year 2014. (3) We used GitHub API to find the Pharo projects hosted in GitHub with more than 10 forks and 20 stars.

Then we applied a set of inclusion and exclusion criteria. Our projects needs to be hosted in GitHub and written in Pharo. They should include a test suite written in sUnit, and can run in Pharo 8 (stable version). For not being overwhelmed with resolving dependencies, they need to support installation with Metacello and not depend on system level packages like databases or a special installation service. We discarded all libraries that are part of the Pharo system such as collections, or compiler.

Based on the mentioned criteria, we selected randomly 20 projects. Then, we rejected projects having less than 4 green test classes with known class under test and mutation coverage less than 100.

Similar to the experimental protocol in DSPOT (Danglot et al. 2019b), we select randomly 4 test classes, 2 high mutation coverage and 2 low, for each project. If a project lacks at least 2 test classes having high (or low) mutation score, we select from lower (higher) covered classes instead. As result, we have 52 test classes, 27 of them considered strong (high mutation score) and 25 considered weaker (lower mutation score).

Table 2 shows the descriptive statistics of the selected projects with a short description, area of usage, number of test classes and test methods and their version based on git commit id, and selected test classes (a superscript ^{*h*} is used to indicate a test class with high mutation coverage, and ^{*l*} is used to indicate low mutation coverage).

Detecting the Class Under Test SMALL-AMP needs a test class and its class-under-test as inputs. Finding a mapping between a test and a class can be challenging. As the default mapping heuristic, we rely on the pattern used by Pharo IDE to detect a test method for a class. The Pharo code browser finds a unit test for a class as follows: it adds the postfix "Test " to the name of the class. If there is such class loaded in the system that is a subclass of TestCase it is considered as the unit test class. If this heuristic is not followed in a project, one can explicitly define the class-under-test by overriding a hook method in test classes.

Metrics We adopt the same metrics used in the experimental protocol in Danglot et al. (2019b):

- **All killed mutants** ($\#Mutants.killed$): The absolute number of mutants killed by a test class in a given class under test.
- **Mutation score** ($\%M.Score$): The ratio (in percentage) of killed mutants over the number of all mutants injected in the class under test.

$$\%M.Score = 100 \times \frac{\#Mutants.killed}{\#Mutants.All}$$

- **Newly killed mutants** ($\#Mutants.killed_{new}$): The number of all new mutants that are killed in an amplified version of the test class.

$$\#Mutants.killed_{new} = \#Mutants.killed_{amplified} - \#Mutants.killed_{original}$$

Table 2 (continued)

Project	commit Id	#test classes	#test methods	Description (Based on Github page)	Test classes in the experiment
pharo-launcher	f1ce748	38	216	Manager for pharo images	PhLAboutCommandTest ^b , PhLCopyImageCommandTest ^b , PhLDirectoryBasedImageRepositoryTest ⁱ , PhLLocalTemplateTest ⁱ
PolyMath	205dfc	57	605	Scientific Computing with Pharo	PMBinomialGeneratorTest ^b , PMBernoulliGeneratorTest ^b , PMFixpointTest ⁱ , PMExponentialDistributionTest ⁱ
Roassal3	69b5645	18	152	Visualization Engine	RSLabelGeneratorTest ^b , RSUMLCClassBuilderTest ^b , RSDraggableCanvasTest ⁱ , RSAThensRendererTest ⁱ
Seaside	3038b49	55	919	Framework for developing sophisticated web applications	WAKeyGeneratorTest ^b , WAXmlCanvasTest ⁱ , WAXmlErrorHandlerTest ⁱ , WACookieTest ^b
Telescope	a4e128a	11	75	Engine for efficiently creating meaningful visualizations	TLExpandCollapseNodesActionTest ^b , TLHideActionTest ^b , TLDistributionMapTest ⁱ , TLLegendTest ⁱ
zinc	ee0d071	27	292	HTTP Components to deal with the HTTP networking protocol	ZnEasyTest ^b , ZnMessageBenchmarkTest ^b , ZnStatusLineTest ⁱ , ZnBivalentWriteStreamTest ⁱ

Table 3 Pull requests submitted on GitHub

Project	Status	Pull request url
PolyMath	Merged	https://github.com/PolyMathOrg/PolyMath/pull/178
Pharo-Launcher	Merged	https://github.com/pharo-project/pharo-launcher/pull/500
DataFrame	Merged	https://github.com/PolyMathOrg/DataFrame/pull/132
Bloc	Merged	https://github.com/feenkcom/Bloc/pull/7
GraphQL	Merged	https://github.com/OBJECTSEMANTICS/GraphQL/pull/12
Zinc	Merged	https://github.com/svenvc/zinc/pull/58
DiscordSt	Merged	https://github.com/JurajKubelka/DiscordSt/pull/75
MaterialDesignLite	Merged	https://github.com/DuneSt/MaterialDesignLite/pull/308
PetitParser2	Open	https://github.com/kursjan/petitparser2/pull/64
OpenPonk	Open	https://github.com/OpenPonk/openponk/pull/35
Telescope	Open	https://github.com/TelescopeSt/Telescope/pull/162

- **Increase killed** ($\%Inc.killed$): The ratio (in percentage) of all newly killed mutants over the number of all killed mutants.

$$\%Inc.killed = 100 \times \frac{\#Mutants.killed_{new}}{\#Mutants.killed_{original}}$$

5.2 RQ1 — Pull Requests

In this experiment, we choose an amplified test method for each project and send a pull-request in GitHub. Before the experiment, we sent a pilot pull-request to learn how developers deal with external contributions. Firstly, we explain the pilot pull-request and then all pull-requests are described one by one. Table 3 demonstrates the status as well as the url of each pull-request per project.

5.2.1 Pull-Requests Preparation

Each pull-request contains a single amplified test method⁴. In order to attract the developers' interest, we try to select a test method testing an important class/method. We select the class under test by scanning their name and relating the names to the context of the project. For example, we know the project *Zinc* is a HTTP component, so the class *ZnRequest* should be a core class. We run the tool on the selected test class, and then scan the generated test methods to select one of them. In selecting an amplified test method, we still consider the vocabularies in the name of the original test method. We also prioritize the tests with more mutants killed.

In addition, we need to explain why a test is valuable in each pull-request. Since some developers may not be familiar with the concept of mutation testing, we need to understand the test in advance and explain it in simpler words.

We inspect at the selected test method and try to understand the effect of the killed mutant and come up with an easy to understand explanation. Examples of explanations are: increasing branch coverage (*PolyMath*), raising an exception (*pharo-project*,

⁴Exception for the project *MaterialDesignLite* where we 2 very similar test methods in a single pull request

DataFrame), covering new state revealing methods (Bloc, Zinc), reducing technical debt (GraphQL).

After selecting an amplified test method, we perform small corrections on the generated code, as a normal Pharo developer would do when see an auto-generated code. These corrections include choosing more meaningful names for the test method, variables and string constants, or deleting the superfluous lines, and adding comments for small hints.

All preparation steps are performed by the first author and are reviewed by the second and third authors. In the time of experiment, the familiarity of the first author about the projects was only studying parts of provided readme description in GitHub. So, he was totally unfamiliar with the projects, he had not contributed to any of the projects, and had never reviewed their code. In fact this shows although there might be more interesting tests for experts, a normal Pharo developer with limited knowledge about the projects is able to review the output and detect some useful test methods that are merged to the projects. The preparation of the tests was quite straightforward and normally did not take more than one hour for each project.

5.2.2 Pilot Pull-Request

Initially, we sent a pull-request⁵ to Seaside project containing the suggestion for adding a set of new lines into an existing test method. The main goal of this pull-request was to learn more about how developers deal with pull-requests from strangers.

We consider the fact that Seaside project is a framework for web application development and we scanned the name of classes and selected WRequestTest because we expected this test class is related to a core class-under-test which interacts with Http requests. Then, we amplified the test class and selected the test method with the most mutants killed.

The selected new test method was able to kill 6 new mutants and was the result of a cooperation between assertion amplification (Section 3.2) and input amplification (Section 3.3). We merged the parts of amplified test into the original test method (#testPostFields). The test is shown in Fig. 1. Lines 5 to 10 are produced by assertion amplification on the original test method (#testPostFields). Line 15 is added by the *method-call-adder* input-amplifier.

We wrote a description for the pull-request trying to explain why this test is useful. We also expressed that the test method is the output of a tool, because it is important to inform developers in advance that they are participating in an experiment.

After a few days the test was merged by one of the project's developers. Moreover, the developer left a valuable comment containing the following points:

- **The suggestions do not fit this test method:** The developer said “*I expected the test-Postfields unit test method to focus on testing the postFields*”. We agree with his remark. If the suggested changes do not have a semantic relation to the original test method, it should be moved to another test or a new one. We considered this advice in the subsequent pull-requests.
- **Usefulness of the result to refactoring the tests:** The developer also stated “*the result of the test amplification makes me evaluate the existing unit tests and refactor them to improve the test coverage and test factorization*”. This shows that even if the immediate results of test amplification are not tidy enough, they still help refactor existing tests.

⁵<https://github.com/SeasideSt/Seaside/pull/1215>

```

2      testPostFields
3      | request headers |
4      request := WRequest method: 'POST' uri: '/foo?bar=1'.
5  +      self
6  +          deny: request isXmlHttpRequest;
7  +          assert: request headers class equals: WAHeaderFields;
8  +          assert: request remoteAddress isNil;
9  +          assert: request isPost;
10 +          assert: request sslSessionId isNil.
11      headers := Dictionary new.
12      headers at: 'content-type' put: WAMimeType formUrlencoded greaseString.
13      request setHeaders: headers.
14      request setBody: 'baz=2&bar=3'.
15 +      self should: [ request bodyDecoded ] raise: WAILlegalStateException.
16      request setPostFields: (WRequestFields new at: 'baz' put: '2'; at: 'bar'
17      put: '3'; yourself).

```

Fig. 1 Improvements on an existing test method submitted to SEASIDE

5.2.3 Pull-Request Details

In the following parts we describe the details on the pull-requests on each project.

PolyMath We sent a pull-request⁶ to this project containing the suggestion for adding a new test method in the test class `PMVectorTest`. The suggested test method is shown in Fig. 2.

This test method is testing the call of the method `#householder` on two different vectors. Before this test, the method `#householder` was not covered in this test class.

The *method-call-adder* input amplifier adds calls to an existing method in the public interface of a class to the test to force the object under test in a new state. We merged two of them in a new test method that execute two different branches in the test method (based on the condition $x \leq 0$). The former vector (line 75 in Fig. 2) forces the `ifTrue` branch and the latter vector (line 80 in Fig. 2) forces `ifFalse` branch. Note that the comment text (line 75) is added manually to increase the readability of the test.

The original test method included two assertions to confirm the type of the returned value of the method (`self assert: w class equals: Array`). The developers asked us to omit these assertion statements. We changed the pull request accordingly and it was merged immediately.

Pharo-Launcher We sent a pull-request⁷ to this project containing the suggestion for adding a new test method in the test class `PhLImportImageCommandTest`. The suggested test method is shown in Fig. 3.

This test is produced from the original test method of `testCanImportAnImage` which verifies an image can be imported using a valid filename. **SMALL-AMP** applies a literal mutation on the file name (`'foo.image'` changed to `'fo.image'`) that results

⁶<https://github.com/PolyMathOrg/PolyMath/pull/178>

⁷<https://github.com/pharo-project/pharo-launcher/pull/500>


```

73 + PMVectorTest >> testHouseholder [
74 +   | u w |
75 +   u := #(-1 0 1) asPMVector. "`x <= 0` when x = -1"
76 +   w := u householder.
77 +   self
78 +       assert: (w at: 1) equals: 1.7071067811865475;
79 +       assert: (w at: 2) asArray equals: #(1.0 -0.0 -0.4142135623730951).
80 +   u := #(1.00001 2.00007) asPMVector. "`x <= 0` when x = 1.00001"
81 +   w := u householder.
82 +   self
83 +       assert: (w at: 1) equals: 0.5527953485259909;
84 +       assert: (w at: 2) asArray equals: #(1.0 -1.6180158992689828).
85 + ]

```

Fig. 2 A new test method submitted to POLYMATH

an invalid filename and consequently raising a `FileDoesNotExistException` error. While preparing the test for the pull-request, we modified the name of the test method and the file to be more meaningful.

The pull-request was merged in the same day with this comment: *“Indeed, the test you are adding has a value. Good job SmallAmp”*.

DataFrame We sent a pull-request⁸ to this project containing the suggestion for adding a new test method in the test class `DataFrameTest`. The suggested test method is shown in Fig. 4.

The variable `df` is an instance variable that has been initialized in the `#setUp` method. It includes a tabular data mixed from numbers and texts. The initial amplified test method was generated by adding the method `#range` as the first statement in one of the original test methods. We recognized the remaining statements as superfluous lines and removed all of them. We also added a comment including the exception description.

This test method makes it explicit that calling the method `#range` on a `DataFrame` object containing non-numerical columns throws an exception. With this new test it becomes an explicit part of the contract for `DataFrame`.

The pull-request was merged after a few weeks. A developer of the project commented: *“Small-amp seems to be a very valuable tool!”*

Bloc We sent a pull-request⁹ to this project containing the suggestion for adding a new assertion in an existing test method in the test class `BlKeyboardProcessorTest`. The suggested test method is shown in Fig. 5.

By calling the state revealing method `#keyStrokesAllowed`, the assertion verifies the correctness of the object state after an `#processKeyDown:` event. This test is the result of combining assertion-amplification with oracle-reduction. Normally, the assertions-amplification step generates lots of assertions, and the oracle-reduction module removes all assertion statements that do not kill any mutant. So, the test code did not need any special preparation and we only need to provide a comment to explain the test method.

⁸<https://github.com/PolyMathOrg/DataFrame/pull/132>

⁹<https://github.com/feenkcom/Bloc/pull/7>

```

38 + PhLImportImageCommandTest >> testImportNonExistingImage [
39 +   | command |
40 +   command := PhLImportImageCommand new.
41 +   command context: presenter.
42 +   presenter := presenter
43 +   requestAnswer: presenter fileSystem / 'tmp' /
44 +   'does_not_exists.image'.
45 +   self should: [ command execute ] raise: FileDoesNotExistException
46 + ]

```

Fig. 3 A new test method sent in a pull-request to the project Pharo-Launcher

The pull-request was also merged after a few weeks with a positive comment.

GraphQL We sent a pull-request¹⁰ to this project containing the suggestion for adding a new test method in the test class `QLSSchemaNodeTest`. The suggested test method is shown in Fig. 6.

This test method verifies the return value of directives in an schema object. The returned value is generated in the method `QLSSchemaNode >> initializeDefaultDirectives` and contains technical debt. This test method guards against future evolutions which may break assumptions made by clients. We selected a meaningful name for the test and wrote a comment text. We also added back some of the assertions removed by oracle-reduction step. The pull-request was merged after a few days.

Zinc We sent a pull-request¹¹ to this project containing the suggestion for adding a new test method in the test class `ZnRequestTest`. The suggested test method is shown in Fig. 7.

This test method calls the method `#setAcceptEncodingGzip` on an request object. Then calls another method `#acceptsEncodingGzip` to verify the change. Both of these methods were not covered in this test class before this test method.

This method is built by cooperating three module of *SMALL-AMP*. First, *method-call-adder* input amplifier adds a new method call. Then *assertion amplification* inserts a set of new assertions after the added method call. And finally, after the main amplification loop is finished, the oracle-reduction step rejects all superfluous assertion statements. This test method did not need much preparation and we only selected a meaningful name for it. The pull-request was merged in the same day.

DiscordSt We sent a pull-request¹² to this project containing the suggestion for adding a new test method in the test class `DSEmbedImageTest`. The suggested test method is shown in Fig. 8.

The method covers the method `#extent` which was not covered in the test class before. This test method did not need much preparation and we only selected a meaningful name for it. The pull-request was merged after a few days.

¹⁰<https://github.com/OBJECTSEMANTICS/GraphQL/pull/12>

¹¹<https://github.com/svenvc/zinc/pull/58>

¹²<https://github.com/JurajKubelka/DiscordSt/pull/75>

```

2874 + DataFrameTest >> testRangeError [
2875 +   self should: [ df range ] raise: MessageNotUnderstood "Instance of
    Character did not understand #Barcelona"
2876 + ]

```

Fig. 4 A new test method sent in a pull-request to the project DataFrame

```

29   BKeyboardProcessorTest >> testProcessKeyDown [
30     | eventA |
31
32     eventA := BKeyDownEvent new.
33     eventA key: BKeyboardKey a.
34
35     processor processKeyDown: eventA.
36 +   self assert: processor keystrokesAllowed.
37   self assert: (processor buffer hasEvent: BKeyboardKey a)
38 ]

```

Fig. 5 A new assertion suggested in a pull-request to the project Bloc

```

292 + GraphQLSchemaNodeTest >> testDirectives [
293 +   schema := self
294 +     parseSchema:
295 +       'type A {
296 +           id: InternalCount
297 +           isB: BooleanType
298 +           size: Int
299 +           idA: ID_A
300 +           values: [ Int ! ]
301 +           params (name: StringName, prom:
    FloatingPoint, key: String): [Int]
302 +         }'.
303 +
304 +   self assert: schema directives class equals: Array.
305 +   self assert: schema directives size equals: 2.
306 +   self
307 +     assert: (schema directives at: 1) class
308 +       equals: GraphQLDirectiveNode.
309 +   self assert: (schema directives at: 1) name equals: 'include'.

```

Fig. 6 A new test method suggested in a pull-request to the project GraphQL

```

1 + tests
2 + testAcceptsEncodingGzip
3 +   | request |
4 +   request := ZnRequest new.
5 +   request setAcceptEncodingGzip.
6 +   self assert: request acceptsEncodingGzip ⊖

```

Fig. 7 A new test method suggested in a pull-request to the project Zinc

```

1 + tests
2 + testExtent
3 +     object := self newObjectToTest.
4 +     object width: 41.
5 +     object height: 42.
6 +     self assert: object extent equals: 41 @ 42

```

Fig. 8 A new test method suggested in a pull-request to the project DiscordSt

MaterialDesignLite We sent a pull-request¹³ to this project containing the suggestion for adding two new test methods in the test class MDLCalendarTest. The suggested test methods are shown in Fig. 9.

Both of test methods are similar and are created by adding a new method call to the test input. The tests are created for the *Calendar* widget and verify correctness of #selectPreviousYears and #selectNextYears methods. In these test methods, the oracle-reduction step removed most of the assertions and it only preserved the first assertion killing the mutant: self assert: calendar yearsInterval fourth equals: 2006. We replaced the assertions with more human readable assertions (asserting first and last of the interval). The pull-request was merged the day after.

PetitParser2 We sent a pull-request¹⁴ to this project containing the suggestion for adding a new test method in the test class PP2NoopVisitorTest. The suggested test method is shown in Fig. 10.

The test method tests the value of currentContext in result object. This test method resulted from assertion amplification combined with oracle-reduction. The test had two assertions: self deny: visitor isRoot and self assert: visitor currentContext class equals: PP2NoopContext. We added back some of removed assertions relating to currentContext and also removed the self deny: visitor isRoot to make the test more focused. The pull-request is not merged up to the date of writing (July 5, 2022).

OpenPonk We sent a pull-request¹⁵ to this project containing the suggestion for adding a set of new lines in an existing test method in the test class OPDiagramTest. The suggested test method is shown in Fig. 11.

The original test method is presented in Listing 14.

SMALL-AMP has broken the statement at line 4 in Listing 14 (the result is visible in lines 104 and 108 in Fig. 11) and then added a series of assertions. Since this test is dedicated to test model, we kept all assertions reflecting the state of model and removed other assertions. So, the assertions in lines 104 to 107 verify the state of a freshly initialized OPDiagram object (where model is nil), and the assertions in lines 109 to 112 verify the public API through the accessor methods. The pull-request is not merged up to the date of writing this paper.

¹³<https://github.com/DuneSt/MaterialDesignLite/pull/308>

¹⁴<https://github.com/kursjan/petitparser2>

¹⁵<https://github.com/OpenPonk/openponk>

```

104 + MDLCalendarTest >> testSelectNextYears [
105 +     calendar currentDate: (Date year: 2016 month: 4 day: 10).
106 +     calendar selectNextYears.
107 +     self assert: calendar yearsInterval first equals: 2021.
108 +     self assert: calendar yearsInterval last equals: 2029
109 + ]

120 + MDLCalendarTest >> testSelectPreviousYears [
121 +     calendar currentDate: (Date year: 2016 month: 4 day: 10).
122 +     calendar selectPreviousYears.
123 +     self assert: calendar yearsInterval first equals: 2003.
124 +     self assert: calendar yearsInterval last equals: 2011
125 + ]

```

Fig. 9 Test methods sent in a pull-request to the project MaterialDesignLite

```

20 + PP2NoopVisitorTest >> testCurrentContext [
21 +     parser := $a asPParser.
22 +     result := visitor visit: parser.
23 +     self assert: result currentContext class equals: PP2NoopContext.
24 +     self assert: result currentContext properties isNil.
25 +     self assert: result currentContext node isNil.
26 +     self assert: result currentContext propertiesCopy isNil
27 + ]

```

Fig. 10 A test method sent in a pull-request to the project PetitParser2

```

101 OPDiagramTest >> testModel [
102     | model project |
103     model := OPTestContainerModel new.
104 +     view := OPDiagram new.
105 +     self assert: view modelType equals: 'UndefinedObject'.
106 +     self assert: view model isNil.
107 +     self assert: view modelName equals: 'UndefinedObject'.
108 +     view model: model.
109 +     self assert: view modelType equals: 'OPTestContainerModel'.
110 +     self assert: view model class equals: OPTestContainerModel.
111 +     self assert: view model name equals: 'container'.
112 +     self assert: view modelName equals: 'container'.
113     self assert: view model equals: model
114 ]

```

Fig. 11 Changes on an existing test method - OpenPonk

Listing 14 Original test method - OpenPonk

```

1 OPDiagramTest >> testModel [
2     | model project |
3     model := OPTestContainerModel new.
4     view := OPDiagram new model: model.
5     self assert: view model equals: model
6 ]

```

```

32 + TLNodeCreationStrategyTest >> testCopyAsSimpleStrategy [
33 +   | aTLNodeCreationStrategy |
34 +   aTLNodeCreationStrategy := strategy copyAsSimpleStrategy.
35 +   self
36 +       assert: aTLNodeCreationStrategy class
37 +       equals: TLNodeCreationStrategy.
38 +   self assert: aTLNodeCreationStrategy childrenStrategy isNil.
39 +   self assert: aTLNodeCreationStrategy compositeProperty isNil.
40 +   self assert: aTLNodeCreationStrategy childrenSortingStrategy isNil.
41 +   self assert: aTLNodeCreationStrategy compositeChildrenLayout isNil.
42 +   self assert: aTLNodeCreationStrategy nodeStyle isNil
43 + ]

```

Fig. 12 A test method suggested in a pull-request to the project Telescope

Telescope We sent a pull-request¹⁶ to this project containing the suggestion for adding a new test method in the test class `TLNodeCreationStrategyTest`. The suggested test method is shown in Fig. 12.

The test method verifies the state of the returned object from calling `copyAsSimpleStrategy`. This method is never covered in the test class. It also contain technical debt. The call to `copyAsSimpleStrategy` is added by method-call-addition amplifier and the state of the returned value is asserted via assertion-amplification. We kept all assertions related to the returned value, and removed all other superfluous lines to make the test more readable. The pull-request is not merged up to the date of writing this paper.

Answer to RQ1: We submitted 11 pull requests through GitHub to propose amplified test methods to developers. In 8 cases, our request was accepted by the developers and the test has been merged to the code base. In the three remaining cases our pull request was ignored. Moreover, we received qualitative feedback from developers acknowledging the relevance of amplified test methods.

5.3 RQ2 — Focus

We use the results in Table 4 for answering the next research questions. These tables present the result of test amplification on the all selected classes selected in our dataset. In Table 4, the first 104 rows represent test amplification for test classes with high mutation coverage, while the remaining of the rows show the test classes with poor mutation coverage. SMALL-AMP algorithm (Algorithm 1) has a stochastic nature, especially test input reduction (Section 4.2) which heavily depends on randomness. Therefore, we ran the algorithm three times on each test class to observe the effect of randomness on the results.

In addition, we ran the algorithm another time by disabling the profiling and the type sensitive operator for investigating the effectiveness of type profiler (denoted by ○ ○ ○).

¹⁶<https://github.com/TelescopeSt/Telescope>

Table 4 The result of test amplification by SMALL-AMP on the 52 test classes. (Tests with high coverage)

Id	Class	# Test methods original	# loc CUT methods original	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# Newly mutant A-amp	% Increase killed only A-amp	# Newly killed type aided	% Increase killed type aided	Time (h:m:s)
1	DataFrameJsonWriterTest	4	51	50	1	1	6	2	33.33	1	16.67	0	0.00	0:00:11
2					1	1	6	2	33.33	1	16.67	0	0.00	0:00:11
3					1	1	6	2	33.33	1	16.67	0	0.00	0:00:11
4	ooo				1	1	6	2	33.33	1	16.67	0	0.00	0:00:09
5	DataFrameCsvReaderTest	4	62	85	0	0	17	0	0.00	0	0	0	0.00	0:00:08
6					0	0	17	0	0.00	0	0	0	0.00	0:00:07
7					0	0	17	0	0.00	0	0	0	0.00	0:00:07
8	ooo				0	0	17	0	0.00	0	0	0	0.00	0:00:07
9	DataFrameJsonReaderTest	9	80	72.73	2	2	16	4	25.00	1	6.25	0	0.00	0:08:43
10					2	2	16	4	25.00	1	6.25	0	0.00	0:08:57
11					2	2	16	4	25.00	1	6.25	0	0.00	0:09:11
12	ooo				2	2	16	4	25.00	1	6.25	0	0.00	0:05:14
13	DataFrameTypeDetectorTest	15	279	94.12	0	0	64	0	0.00	0	0	0	0.00	0:02:02
14					0	0	64	0	0.00	0	0	0	0.00	0:02:06
15					0	0	64	0	0.00	0	0	0	0.00	0:02:06
16	ooo				0	0	64	0	0.00	0	0	0	0.00	0:01:48
17	ZnMessageBenchmarkTest	2	187	80	0	0	24	0	0.00	0	0	0	0.00	0:01:23
18					0	0	24	0	0.00	0	0	0	0.00	0:01:21
19					0	0	24	0	0.00	0	0	0	0.00	0:01:24
20	ooo				0	0	24	0	0.00	0	0	0	0.00	0:01:05

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# mutant A-amp	% Increase killed only A-amp	# Newly killed type aided	% Increase killed type aided	Time (h:m:s)
21	ZnEasyTest	10	65	66.67	0	0	10	0	0.00	0	0	0	0.00	0:46:25
22					0	0	10	0	0.00	0	0	0	0.00	0:44:20
23					0	0	10	0	0.00	0	0	0	0.00	0:46:35
24	ooo				0	0	10	0	0.00	0	0	0	0.00	0:28:16
25	GQLRequestGrammarTest	29	142	96.55	0	0	56	0	0.00	0	0	0	0.00	0:04:07
26					0	0	56	0	0.00	0	0	0	0.00	0:04:09
27					0	0	56	0	0.00	0	0	0	0.00	0:04:02
28	ooo				0	0	56	0	0.00	0	0	0	0.00	0:04:10
29	GQLSingleAnonymousQuery EvaluatorTest	22	56	52.27	2	2	23	2	8.70	0	0	0	0.00	0:05:21
30					2	2	23	2	8.70	0	0	0	0.00	0:05:26
31					2	2	23	2	8.70	0	0	0	0.00	0:05:15
32	ooo				1	1	23	1	4.35	0	0	0	0.00	0:06:02
33	GQLSchemaGrammarTest	57	40	93.2	7	7	96	7	7.29	0	0	6	6.25	2:16:15
34					7	7	96	7	7.29	0	0	6	6.25	2:22:55
35					7	7	96	7	7.29	0	0	6	6.25	2:34:02
36	ooo				0	0	96	0	0.00	0	0	0	0.00	0:11:01
37	WebGrammarTest	14	28	90	2	2	18	2	11.11	0	0	2	11.11	0:46:51
38					2	2	18	2	11.11	0	0	2	11.11	0:43:21
39					2	2	18	2	11.11	0	0	2	11.11	0:42:40
40	ooo				0	0	18	0	0.00	0	0	0	0.00	0:16:10

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# Newly mutant A-amp	% Increase killed only A-amp	# Newly killed type aided	% Increase killed type aided	Time (h:m:s)
41	PP2BufferStreamTest	16	188	83.87	0	0	78	0	0.00	0	0	0	0.00	0:00:34
42					0	0	78	0	0.00	0	0	0	0.00	0:00:35
43					0	0	78	0	0.00	0	0	0	0.00	0:00:36
44	ooo				0	0	78	0	0.00	0	0	0	0.00	0:00:37
45	MDLPanelSwitcherButtonTest	6	100	53.85	0	0	7	0	0.00	0	0	0	0.00	0:00:46
46					0	0	7	0	0.00	0	0	0	0.00	0:00:46
47					0	0	7	0	0.00	0	0	0	0.00	0:00:47
48	ooo				0	0	7	0	0.00	0	0	0	0.00	0:00:47
49	MDLPaginationComponentTest	11	199	71.19	0	0	42	0	0.00	0	0	0	0.00	0:00:27
50					0	0	42	0	0.00	0	0	0	0.00	0:00:27
51					0	0	42	0	0.00	0	0	0	0.00	0:00:26
52	ooo				0	0	42	0	0.00	0	0	0	0.00	0:00:27
53	RSUMLCClassBuilderTest	2	23	50	1	1	1	1	100.00	1	100	0	0.00	0:05:00
54					1	1	1	1	100.00	1	100	0	0.00	0:05:34
55					1	1	1	1	100.00	1	100	0	0.00	0:05:23
56	ooo				1	1	1	1	100.00	1	100	0	0.00	0:03:30
57	RSLabelGeneratorTest	2	353	82.38	0	0	173	0	0.00	0	0	0	0.00	0:11:19
58					0	0	173	0	0.00	0	0	0	0.00	0:11:07
59					0	0	173	0	0.00	0	0	0	0.00	0:12:21
60	ooo				0	0	173	0	0.00	0	0	0	0.00	0:12:55
61	OPNullSerializerTest	3	251	50	0	0	1	0	0.00	0	0	0	0.00	0:00:00

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# mutant A-amp	% Increase killed only A-amp	# Newly killed type aided	% Increase killed type aided	Time (h:m:s)
62					0	0	1	0	0.00	0	0	0	0.00	0:00:00
63					0	0	1	0	0.00	0	0	0	0.00	0:00:00
64	ooo				0	0	1	0	0.00	0	0	0	0.00	0:00:00
65	OPRElementsConstraintTest	2	42	50	1	1	1	1	100.00	1	100	0	0.00	0:00:03
66					1	1	1	1	100.00	1	100	0	0.00	0:00:03
67					1	1	1	1	100.00	1	100	0	0.00	0:00:03
68	ooo				1	1	1	1	100.00	1	100	0	0.00	0:00:03
69	PMBernoulliGeneratorTest	4	113	76.47	2	2	13	3	23.08	0	0	0	0.00	0:00:06
70					3	3	13	2	15.38	0	0	0	0.00	0:00:07
71					2	2	13	2	15.38	0	0	0	0.00	0:00:06
72	ooo				3	3	13	3	23.08	0	0	0	0.00	0:00:05
73	PMBinomialGeneratorTest	3	164	83.33	1	1	10	1	10.00	1	10	0	0.00	0:01:01
74					1	1	10	1	10.00	1	10	0	0.00	0:01:14
75					1	1	10	1	10.00	1	10	0	0.00	0:01:18
76	ooo				1	1	10	1	10.00	1	10	0	0.00	0:01:31
77	TLHideActionTest	3	468	55.56	2	2	5	0	0.00	0	0	0	0.00	0:00:30
78					3	3	5	3	60.00	0	0	1	20.00	0:00:33
79					1	1	5	2	40.00	0	0	0	0.00	0:00:26
80	ooo				1	1	5	2	40.00	0	0	0	0.00	0:00:14
81	TLExpandCollapseNodes ActionTest	3	8	66.67	2	2	14	2	14.29	1	7.14	0	0.00	0:01:28
82					2	2	14	2	14.29	0	0	0	0.00	0:00:58

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# New mutant A-amp	% Increase killed only A-amp	Newly killed type		% Increase killed type aided	Time (h:m:s)
												aided	aided		
83	DSDetectChannelCommandTest	5	51	50	2	2	14	2	14.29	0	0	0	0	0.00	0:01:16
84					1	1	14	1	7.14	0	0	0	0	0.00	0:00:31
85					0	0	3	0	0.00	0	0	0	0	0.00	0:00:40
86					0	0	3	0	0.00	0	0	0	0	0.00	0:00:41
87					0	0	3	0	0.00	0	0	0	0	0.00	0:00:39
88	DSUserTest	14	62	56.52	0	0	3	0	0.00	0	0	0	0	0.00	0:00:30
89					1	1	13	4	30.77	1	7.69	0	0	0.00	0:13:53
90					1	1	13	4	30.77	1	7.69	0	0	0.00	0:17:13
91					1	1	13	4	30.77	1	7.69	0	0	0.00	0:14:45
92					1	1	13	4	30.77	1	7.69	0	0	0.00	0:10:15
93	PhLAAboutCommandTest	1	111	75	0	0	3	0	0.00	0	0	0	0	0.00	0:00:01
94					0	0	3	0	0.00	0	0	0	0	0.00	0:00:01
95					0	0	3	0	0.00	0	0	0	0	0.00	0:00:01
96					0	0	3	0	0.00	0	0	0	0	0.00	0:00:01
97					0	0	4	0	0.00	0	0	0	0	0.00	0:00:02
98	PhLCopyImageCommandTest	1	25	66.67	0	0	4	0	0.00	0	0	0	0	0.00	0:00:02
99					0	0	4	0	0.00	0	0	0	0	0.00	0:00:01
100					0	0	4	0	0.00	0	0	0	0	0.00	0:00:02
101					1	1	33	4	12.12	1	3.03	0	0	0.00	0:02:33
102					2	2	33	4	12.12	1	3.03	1	3.03	3.03	0:02:39
103	WACookieTest	16	34	50	2	2	33	4	12.12	1	3.03	1	3.03	3.03	0:02:30
104					1	1	33	4	12.12	1	3.03	0	0	0.00	0:01:45

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# mutant A-amp	% Increase only killed A-amp	# Newly killed aided	% Increase killed type aided	Time (h:m:s)
105	BIShortcutTest	1	113	20	1	1	2	4	200.00	1	50	0	0.00	0:00:04
106					1	1	2	4	200.00	1	50	0	0.00	0:00:04
107					1	1	2	4	200.00	1	50	0	0.00	0:00:04
108	ooo				1	1	2	4	200.00	1	50	0	0.00	0:00:02
109	BICompulsoryCombinationTest	4	271	11.76	1	1	4	2	50.00	0	0	1	25.00	0:00:34
110					1	1	4	2	50.00	0	0	1	25.00	0:00:33
111					1	1	4	2	50.00	0	0	1	25.00	0:00:36
112	ooo				0	0	4	0	0.00	0	0	0	0.00	0:00:23
113	BILayoutExactResizerTest	8	36	40.54	3	3	15	3	20.00	1	6.67	2	13.33	0:01:08
114					3	3	15	3	20.00	1	6.67	2	13.33	0:01:07
115					3	3	15	3	20.00	1	6.67	2	13.33	0:01:06
116	ooo				2	2	15	2	13.33	2	13.33	0	0.00	0:00:54
117	BIInsetsTest	10	30	44.14	14	12	64	49	76.56	2	3.12	3	4.69	0:03:10
118					18	14	64	46	71.88	2	3.12	11	17.19	0:03:08
119					18	18	64	50	78.13	2	3.12	4	6.25	0:03:13
120	ooo				19	19	64	50	78.13	2	3.12	0	0.00	0:02:01
121	ZnBivalentWriteStreamTest	2	39	44.44	0	0	8	0	0.00	0	0	0	0.00	0:00:04
122					0	0	8	0	0.00	0	0	0	0.00	0:00:04
123					0	0	8	0	0.00	0	0	0	0.00	0:00:04
124	ooo				0	0	8	0	0.00	0	0	0	0.00	0:00:04
125	ZnStatusLineTest	7	220	28.57	1	1	16	1	6.25	0	0	1	6.25	0:00:41
126					1	1	16	1	6.25	0	0	1	6.25	0:00:43

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# Newly mutant A-amp	% Increase killed only A-amp	# Newly killed type		% Increase killed type aided	Time (h:m:s)
												aided	aided		
127					1	1	16	1	6.25	0	0	1		6.25	0:00:44
128	ooo				0	0	16	0	0.00	0	0	0		0.00	0:00:39
129	QQLArgumentsTest	16	83	47.06	0	0	8	0	0.00	0	0	0		0.00	0:31:10
130					0	0	8	0	0.00	0	0	0		0.00	0:30:56
131					0	0	8	0	0.00	0	0	0		0.00	0:32:28
132	ooo				0	0	8	0	0.00	0	0	0		0.00	0:27:17
133	PP2ParsingGuardTest	3	29	47.06	1	1	8	2	25.00	1	12.5	0		0.00	0:00:08
134					1	1	8	2	25.00	1	12.5	0		0.00	0:00:07
135					1	1	8	2	25.00	1	12.5	0		0.00	0:00:07
136	ooo				1	1	8	2	25.00	1	12.5	0		0.00	0:00:07
137	PP2BenchmarkTest	3	24	15.91	1	1	7	6	85.71	1	14.29	0		0.00	0:09:15
138					1	1	7	6	85.71	1	14.29	0		0.00	0:08:48
139					1	1	7	6	85.71	1	14.29	0		0.00	0:08:31
140	ooo				1	1	7	6	85.71	1	14.29	0		0.00	0:08:51
141	MDLDialogTest	4	103	14.29	0	0	1	0	0.00	0	0	0		0.00	0:00:00
142					0	0	1	0	0.00	0	0	0		0.00	0:00:00
143					0	0	1	0	0.00	0	0	0		0.00	0:00:00
144	ooo				0	0	1	0	0.00	0	0	0		0.00	0:00:00
145	MDLNestedListTest	11	558	41.04	1	1	55	8	14.55	1	1.82	0		0.00	0:00:14
146					1	1	55	8	14.55	1	1.82	0		0.00	0:00:15
147					1	1	55	8	14.55	1	1.82	0		0.00	0:00:15
148	ooo				1	1	55	8	14.55	1	1.82	0		0.00	0:00:14

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT original	% Mut. score	# New test methods	# Focused methods	# Killed mutants original	# Newly killed amplified	% Increase killed amplified	# mutant A-amp	% Increase killed only A-amp	# Newly killed type aided	% Increase killed type aided	Time (h:m:s)
149	RSAthensRendererTest	1	19	9.88	1	1	48	1	2.08	1	2.08	0	0.00	0:01:08
150					1	1	48	1	2.08	1	2.08	0	0.00	0:01:10
151					1	1	48	1	2.08	1	2.08	0	0.00	0:01:15
152	ooo				1	1	48	1	2.08	1	2.08	0	0.00	0:01:04
153	RSDraggableCanvasTest	6	84	35.71	0	0	10	0	0.00	0	0	0	0.00	0:00:34
154					0	0	10	0	0.00	0	0	0	0.00	0:00:35
155					0	0	10	0	0.00	0	0	0	0.00	0:00:36
156	ooo				0	0	10	0	0.00	0	0	0	0.00	0:00:32
157	OPNavigatorAdaptersTest	3	161	24	0	0	6	0	0.00	0	0	0	0.00	0:00:26
158					0	0	6	0	0.00	0	0	0	0.00	0:00:24
159					0	0	6	0	0.00	0	0	0	0.00	0:00:23
160	ooo				0	0	6	0	0.00	0	0	0	0.00	0:00:26
161	OPPProjectTest	2	40	28	0	0	7	0	0.00	0	0	0	0.00	0:00:03
162					0	0	7	0	0.00	0	0	0	0.00	0:00:03
163					0	0	7	0	0.00	0	0	0	0.00	0:00:03
164	ooo				0	0	7	0	0.00	0	0	0	0.00	0:00:04
165	PMExponentialDistributionTest	1	25	17.07	3	3	7	3	42.86	0	0	2	28.57	0:00:23
166					2	2	7	2	28.57	0	0	1	14.29	0:00:15
167					3	3	7	3	42.86	0	0	2	28.57	0:00:11
168	ooo				1	1	7	1	14.29	0	0	0	0.00	0:00:05
169	PMFixpointTest	6	18	39.47	5	5	30	12	40.00	3	10	0	0.00	0:03:01
170					4	4	30	9	30.00	3	10	0	0.00	0:03:03

Table 4 (continued)

Id	Class	# Test methods original	# loc CUT	% Mut. score	# New test methods	# Focused methods	# Killed mutants		% Increase killed	# Newly mutant A-amp		% Increase killed only A-amp	# Newly killed type aided		% Increase killed type aided	Time (h:m:s)
							original	amplified		amplified	A-amp		aided	aided		
171					5	5	30	12	40.00	3	10	10	0	0	0.00	0:04:35
172	ooo				7	7	30	8	26.67	3	10	10	0	0	0.00	0:02:11
173	TLDistributionMapTest	1	168	37.04	3	3	10	4	40.00	1	10	10	2	2	20.00	0:00:26
174					3	3	10	4	40.00	1	10	10	2	2	20.00	0:00:26
175					2	2	10	3	30.00	1	10	10	1	1	10.00	0:00:28
176	ooo				1	1	10	2	20.00	1	10	10	0	0	0.00	0:00:14
177	TLLegendTest	2	96	0	0	0	0	0	-	0	-	-	0	0	-	0:00:02
178					0	0	0	0	-	0	-	-	0	0	-	0:00:02
179					0	0	0	0	-	0	-	-	0	0	-	0:00:02
180	ooo				0	0	0	0	-	0	-	-	0	0	-	0:00:02
181	DSSendUserTextMessageItemTest	5	62	33.33	2	2	4	3	75.00	1	25	25	0	0	0.00	0:00:34
182					2	2	4	3	75.00	1	25	25	0	0	0.00	0:00:36
183					2	2	4	3	75.00	1	25	25	0	0	0.00	0:00:34
184	ooo				2	2	4	3	75.00	2	50	50	0	0	0.00	0:00:21
185	DSEmbedTest	11	33	19.67	7	7	12	20	166.67	1	8.33	8.33	6	6	50.00	0:01:58
186					7	7	12	20	166.67	1	8.33	8.33	6	6	50.00	0:02:02
187					8	8	12	20	166.67	1	8.33	8.33	6	6	50.00	0:02:16
188	ooo				2	2	12	12	100.00	2	16.67	16.67	0	0	0.00	0:01:15
189	PhLLocalTemplateTest	11	81	12	1	1	3	2	66.67	1	33.33	33.33	0	0	0.00	0:00:33
190					1	1	3	2	66.67	1	33.33	33.33	0	0	0.00	0:00:32
191					1	1	3	2	66.67	1	33.33	33.33	0	0	0.00	0:00:33
192	ooo				1	1	3	2	66.67	1	33.33	33.33	0	0	0.00	0:00:39

Table 4 (continued)

Id	Class	# Test methods original	# loc	% Mut. score	# New test methods	# Focused methods	# Killed		# Newly killed		% Increase		# Newly killed only		% Increase		% Increase Time (h:m:s)
							mutants original	killed	mutants amplified	killed	mutant A-amp	killed A-amp	mutant A-amp	killed A-amp			
193	PhLDirectoryBasedImageRepositoryTest	14	156	39.8	1	1	39	3	7.69	1	2.56	0	0.00	0	0.00	0:01:21	
194					1	1	39	3	7.69	1	2.56	0	0.00	0	0.00	0:01:17	
195					1	1	39	3	7.69	1	2.56	0	0.00	0	0.00	0:01:19	
196	ooo				1	1	39	3	7.69	1	2.56	0	0.00	0	0.00	0:01:30	
197	WAKeyGeneratorTest	1	51	33.33	0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:12	
198					0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:09	
199					0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:09	
200	ooo				0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:10	
201	WAXmlCanvasTest	1	142	33.33	0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:00	
202					0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:00	
203					0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:01	
204	ooo				0	0	1	0	0.00	0	0	0	0.00	0	0.00	0:00:01	
205	WAErroHandlerTest	11	182	38.46	0	0	5	0	0.00	0	0	0	0.00	0	0.00	0:01:15	
206					0	0	5	0	0.00	0	0	0	0.00	0	0.00	0:01:13	
207					0	0	5	0	0.00	0	0	0	0.00	0	0.00	0:01:17	
208	ooo				0	0	5	0	0.00	0	0	0	0.00	0	0.00	0:01:19	

The columns in this table indicate:

- *Id*: Used as a reference for the row in the table.
- *Class*: The name of the test class to be amplified.
- *# Test methods original*: The number of test method in the test class before test amplification.
- *# loc CUT*: The number of lines in the class under test.
- *% Mut. score*: The mutation score (percentage) of the test class before test amplification.
- *# New test methods*: The absolute number of newly generated test methods after test amplification.
- *# Focused methods*: The absolute number of focused methods in the generated test methods.
- *# Killed mutants original*: The absolute number of killed mutants by the test class before test amplification.
- *# Newly killed amplified*: The absolute number of newly killed mutants by the test class after test amplification.
- *% Increase killed amplified*: The increase (in percentage) of killed mutants by the test class after test amplification.
- *# Newly mutant A-amp*: The absolute number of newly killed mutants only by Assertion amplification ($N_{iteration} = 0$ in Algorithm 1).
- *% Increase killed only A-amp*: The relative increase (in percentage) of killed mutants only by assertion amplification.
- *# Newly killed type aided*: The absolute number of newly killed mutants by type sensitive input amplifiers.
- *% Increase killed type aided*: The relative increase (in percentage) of increase killed mutants by type sensitive input amplifiers.
- *Time*: The duration of test amplification process in the hours-minutes-seconds (h:m:s) format.

RQ2: To what extent are improved test methods considered as focused? For answering this research question, we use values in the column *# Focused methods*. We use the same definition for focused methods as DSPOT:

Focus is defined as where at least 50% of the newly killed mutants are located in a single method. Danglot et al. (2019b)

Generating focused tests is important because analysing a focused test is easier (most mutants reside in the same method under test) hence should take less review time from developers. For calculating this value, we use generated annotations by SMALL-AMP on the newly generated test methods which show the details of the killed mutants by the method.

We see that almost all amplified tests are focussed. Only on two cases (BlInsetsTest, #117 and #118) we see that some generated methods are not focused.

Answer to RQ2: We see at least one focused test method in all amplified cases. This illustrates that amplified tests seldom overwhelm developers, hence should save valuable developer time.

5.4 RQ3 — Mutation Coverage

RQ3: To what extent do improved test classes kill more mutants than developer-written test classes? In 86/156 cases (55.12%), SMALL-AMP successfully amplified an existing test class. The distribution of the number of killed mutants, and increase kill are presented in Figs. 13 and 14 for all test classes, highly covered as well as poorly covered ones. The number of newly killed mutants in these classes (column # *Newly killed amplified*) varies from 1 up to 50 mutants (case #119). In the executions amplifying test classes having high coverage, SMALL-AMP is able to amplify 38 out of 78 (48.7%), and for the test classes having poor coverage this number is 48 out of 78 (61.5%). Therefore, we see more amplification in the classes with poor coverage. The relative increase in mutation score (column % *Increase killed amplified*) varies from 2.08% (cases 149-151) up to 200% (cases 105-107). It is also observable regarding to these metrics that amplification on classes with poor coverage is more successful.

Surprisingly, despite running MUTALK with its all mutation operators, the mutation testing framework did not manage to create any mutant for the class TLLegendTest (cases 177-179). MUTALK mutation operators work statically and only a limited set of well-known transformations are provided in the tool.

The Effect of Randomness In this section we report the effect of randomness on the results. Based on the Algorithm 1, the main randomness happens during the input-amplification (Line 10) and oracle reduction (Line 12) steps. Therefore, we expect to see the minimum difference in the results generated by assertion amplification (Line 5). The column 11 (# *Newly mutant A-amp*) shows the absolute number of killed mutants only by assertion amplification. These values are identical in executions for all classes except the case 81 (TLExpandCollapseNodesActionTest). The reason for this exception is that a specific mutant may be killed by input amplification in a test method, and if it is not killed, it will be killed by assertion amplification in the next test method. Based on the information presented in Table 4, regardless of time, the same result are achieved from different executions in 43 out of 52 test classes.

For a deeper investigation, we randomly select 5 test classes from the cases that are affected by randomness, and 5 test classes from the cases without an observable change. Then, we run SMALL-AMP on these classes for 10 times (10 class \times 10 times = 100 runs). Table 5 shows the results of this experiment. In the column with title *X10*, we report the number of newly killed mutants for each run in order, which is the most important metric for amplification. In the column *X3*, we also echo the values from column # *Newly killed amplified* in Table 4 to make the comparison easier. The next two columns compare the *Median* and *Average* values in these two columns. The first 5 rows in this table are the cases affected by randomness, and the next 5 rows are cases without an observable effect.

When we compare the values in columns *X10* with *X3*, we still do not see any visible effect from randomness in the rows 6 to 10. In the first 5 rows, we see the median values and




Fig. 13 The distributions of the number of killed mutants




Fig. 14 The distributions of the increase kills

average values in both experiments are similar. In three cases (rows 2, 3 and 4), the values of X3 did not achieve the maximum number of killed mutants seen in X10. In one case (row 1), we see some runs lacking any improvements in X10 while all of its runs in X3 shows a successful amplification. To sum up, we see that the randomness has an effect on the result, but the impact is minimal and does not invalidate the findings. In addition, repeating the analysis 3 times is justifiable since running 10 times adds little extra information for a large increase in processing time.

Answer to RQ3: SMALL-AMP successfully amplified 86 test classes out of 156 cases (55.12%). Even for stronger test classes, SMALL-AMP improved the mutation coverage in 38 out of 78 cases. In test classes with poor coverage, test amplification becomes even more effective: SMALL-AMP increased the mutation coverage in 48 out of 78 cases and the absolute and relative increase in mutation score was higher.

5.5 RQ4 — Amplification Steps

RQ4: What is the contribution of input amplification and assertion amplification (the main steps in the test amplification algorithm) to the effectiveness of the generated tests? As we reported in Section 5.4, in 86/156 cases (55.12%), the improvements are achieved from input-amplification and assertion-amplification cooperation. In this research

Table 5 The result of running SMALL-AMP on 10 test class for 10 times

Test class		X10	X3	Median (X10, X3)	Average (X10, X3)
1	PMBernoulliGeneratorTest	3, 3, 0, 0, 2, 2, 2, 3, 2, 3	3, 2, 2	2.0, 2.0	2.0, 2.33
2	TLHideActionTest	0, 4, 1, 2, 2, 4, 2, 1, 2, 2	0, 3, 2	2.0, 2.0	2.0, 1.67
3	TLExpandCollapseNodesActionTest	2, 2, 2, 2, 2, 2, 3, 2, 2, 2	2, 2, 2	2.0, 2.0	2.1, 2.0
4	PMExponentialDistributionTest	3, 2, 1, 4, 3, 2, 2, 3, 3, 2	3, 2, 3	2.5, 3.0	2.5, 2.67
5	TLDistributionMapTest	3, 4, 3, 4, 4, 4, 4, 3, 3, 4	4, 4, 3	4.0, 4.0	3.6, 3.67
6	PMBinomialGeneratorTest	1, 1, 1, 1, 1, 1, 1, 1, 1, 1	1, 1, 1	1.0, 1.0	1.0, 1.0
7	DSSendUserTextMessageItemTest	3, 3, 3, 3, 3, 3, 3, 3, 3, 3	3, 3, 3	3.0, 3.0	3.0, 3.0
8	GQLSchemaGrammarTest	7, 7, 7, 7, 7, 7, 7, 7, 7, 7	7, 7, 7	7.0, 7.0	7.0, 7.0
9	BICompulsoryCombinationTest	2, 2, 2, 2, 2, 2, 2, 2, 2, 2	2, 2, 2	2.0, 2.0	2.0, 2.0
10	ZnMessageBenchmarkTest	0, 0, 0, 0, 0, 0, 0, 0, 0, 0	0, 0, 0	0.0, 0.0	0.0, 0.0

question, we study the results generated only by assertion-amplification, and also generated by the type sensitive operators.

Contribution of the Assertion-Amplification Step In this section, we filter all amplified test methods that are generated only by assertion amplification. In other words, we only account the improvements from all amplified test methods that are selected from the first assertion amplification (Line 5 in Algorithm 1). The column 11 (*# Newly mutant A-amp*) shows the absolute number of killed mutants only by assertion amplification; column 12 (*% Increase killed only A-amp*) also shows the relative increase. The reported results in Table 4 shows that in 61/156 cases (39.1%) at least 1 mutant is killed only by the assertion amplification step. Improvements in four classes (cases 53-55; 65-67; 73-75; 149-151) achieved only by adding new assertions.

Contribution of the Type Sensitive Input Amplifiers Here, we filter all amplified test methods that in at least one of its transformations, a type sensitive input amplifier (in our case *method-call-adder*) is used. While type-sensitive operators benefit the information provided by dynamic profiler step (Section 4.1), the contribution of these operators is important because it can show the effectiveness of dynamic profiling.

Column 13 (*# Newly killed type aided*) shows the absolute number of newly killed mutants by the type sensitive input amplifiers. Column 14 (*% Increase killed type aided*) also shows the relative increase. We see that in 30/156 cases (19.2%) the type sensitive input amplifiers contribute to the result. Especially for 2 test classes (*WebGrammarTest* rows 37-39, and *ZnStatusLineTest* rows 125-127), SMALL-AMP was able to amplify the tests only by the type sensitive input operators.

To assess the impact of type profiling, we quantified the effect of the steps that rely on type profiling. We therefore extended the analysis with an additional evaluation step where we disabled the type profiler in the algorithm, as well as the type sensitive input amplifier (method addition amplifier) and ran the tool on all test classes. The results for this experiment are mentioned in the forth row for each test class in Table 4 (denoted by $\circ \circ \circ$). We focus on cases in which type-sensitive input amplifiers improved the coverage in at least two of three runs (10 test classes, cases starting with 33, 37, 101, 109, 113, 117, 125, 165, 173, 185). In 8 cases we see that disabling the profiling and also the type sensitive operators decrease the improvements and only in two cases we see no difference (case 101) or a slight improvements (case 117).

Answer to RQ4: Our experiments demonstrate that amplifying the tests only using assertion amplification is less efficient than in combination with input amplification. Moreover, the extra information generated by dynamic type profiling helps input amplification in killing more mutants.

5.6 RQ5 — Comparison

RQ5: How does Small-Amp compare against DSpot? With this research question, we compare our results from quantitative and qualitative studies with corresponding results from DSPOT reported in Danglot et al. (2019b).

Table 6 shows how the results from SMALL-AMP and DSPOT are comparable. SMALL-AMP is validated against 52 test classes. It has successfully amplified 28, 29 and 29 of them ($\approx 55\%$), while DSPOT has been validated against 40 test classes of which 26 cases were improved (65%). The most notable differences between the results from SMALL-AMP and DSPOT are the number of mutants in two ecosystems and consequently the number of newly generated test methods (denoted by * in the Table 6). These differences can be attributed to the use of two different mutation testing frameworks in two different languages. SMALL-AMP uses MUTALK which has notably fewer mutation operators than the DSPOT counterpart PITEST. To reduce the effect of the mutation testing framework, we calculate the relative increase in killed mutants within the two ecosystem as follows:

$$\%Total.Inc.killed = 100 \times \frac{\#Total.Mutants.killed_{new}}{\#Total.Mutants.killed_{original}}$$

This value is shown in the row number 12 in the Table 6 for two experiments. It is 14.03% in total for SMALL-AMP, and 20.26% for DSPOT.

We have also submitted 11 pull-requests by using SMALL-AMP outputs and 8 of them were merged by developers ($\approx 72\%$), while Danglot et al. submitted 19 pull requests derived from DSPOT output and 13 of them merged successfully (68%).

Finally, it is worth mentioning that SMALL-AMP is configured as $N_{maxInput} = 10$ which means the reduce step (Algorithm 1, line 12) select 10 test-input in each iteration. This value for DSPOT is not reported in their paper. Increasing this hyperparameter may improve the result, but it also may increase the execution time significantly.

Answer to RQ5: The results from SMALL-AMP and DSPOT in two different ecosystems are comparable. SMALL-AMP and DSPOT have been successful in amplifying respectively 55% and 65% of their input test classes. We also see $\approx 72\%$ and 68% merged pull requests in the tests derived from SMALL-AMP and DSPOT outputs.

Table 6 Comparing results in SMALL-AMP and DSPOT

Id	Metric	SMALL-AMP			DSPOT
		#1	#2	#3	
1	Projects	13	13	13	10
2	Test classes	52	52	52	40
3	Test methods	403	403	403	220
4*	New generated test methods	71	76	75	471
5	# Amplified test classes	28	29	29	26
6	% Amplified test classes	53%	55%	55%	65%
7*	# Total killed mutants by original	1102	1102	1102	7980
8*	# Total newly killed mutants	156	151	157	1617
9	% Total increase killed	14.15%	13.70%	14.24%	20.26%

5.7 RQ6 — Time Costs

RQ6: What is the time cost of running Small-Amp, including its steps? Time costs is an important factor when we study the practicality of test amplification tools. In this research question, we report and compare the execution time of SMALL-AMP and the relative cost of each step. Figures 15 and 16 illustrate a series of box-plots derived from the recorded execution time during the experiments in Table 4. In these figures, *Init.* refers to all initializing steps, includes the dynamic profiling to collect type information (Section 4.1). Here we also calculate an initial mutation score for the original test suite. *IAmp* refers to the input amplification step. This step loops over all input amplification operators (Section 3.2) and afterwards reduces to $N_{\max\text{Inputs}}$ of current inputs and discarding the rest (Section 4.2). *AAmp* represents the assertion amplification step (Section 3.3), while *Sel.* selects all test methods that increase the mutation score (Section 3.4). *Read.* concerns the post-processing steps to increase the readability of the generated tests, in particular the oracle reduction (Section 4.3). Finally, *Tot.* shows the entire execution time for a class.

Figure 15 shows how the execution time is distributed for total execution time and also for each step in seconds. The horizontal axis presents the number of seconds in logarithmic scale. The diagram includes also the values for the median and the upper whisker.

Regarding total execution time (*Tot.*), half of the executions finished in less than 36.7 seconds (the median value). Furthermore, the diagram shows that the majority of amplification (upper whisker) finished in less than 334.4 seconds (5 minutes and 34 seconds). However, we see 25 outliers that refer to the instances that finished in more than 335 seconds. If we set a fixed time budget, for instance a 10 minutes budget for each class, the test amplification process for these classes will not terminate properly. This shows the importance of considering time budget management in test amplification tools.

The median value for other steps are: initializing 1 second, input amplification 4.8 seconds, assertion amplification 13.7 seconds, selection by mutation testing 10.1 seconds and post-processing steps 0.7 seconds.

Figure 16 illustrates the relative proportion of time test amplification dedicates to each step. So, the execution time for each step is divided to the total amplification time to compare the steps relatively.




Fig. 15 The distribution of absolute execution time (in seconds)




Fig. 16 The relative distributions of the time-cost (percentage)

The profiling step and the oracle reductions steps are the fastest steps. The median value for each of these steps are respectively 3.2% and 3.4%. Therefore, we can say profiling and oracle reduction steps do not add much time overhead to the overall process. Next, the input amplification step takes about 11.3%. A considerable portion of execution time is spent during assertion amplification (median 38.8%, upper whisker 95%). The median execution time related to selection step, in which mutation testing is ran, is 24.5%. We suspect the execution time for mutation testing would be more if MUTALK would generate more mutants.

Answer to RQ6: The majority of classes in our dataset has been amplified in less than 5 minutes and 34 seconds. However, in some cases the execution takes longer with a maximum of 2 hours and 34 minutes. Therefore, a time budget management mechanism is needed in the test amplification tools. In the execution time for each steps, we see that the profiling and oracle reductions steps (the extra steps compared to the original DSPOT algorithm) do not add much overhead to the overall process.

6 Threats to Validity

As in any empirical research, we identify factors that may jeopardise the validity of our results and the actions we took to reduce or alleviate the risk. Consistent with the guidelines provided by Wohlin et al. (2000), we organise them into four categories.

Construct Validity Do we measure what was intended? For RQ1 (Pull Requests), we manually selected test methods which we considered valuable additions to the project. And we provided a motivation for the pull request based on a human interpretation of the extra mutants killed. Thus, the percentage of accepted pull requests is a flattered result. If we would have submitted all amplified test methods the results would be far lower. We consider this risk as minimal, because SMALL-AMP at the current stage should never be seen as a fully automated code synthesizer tool but rather as a recommender system supporting the human-in-the-loop.

For RQ2 to RQ4 we heavily rely on mutation coverage as a proxy for the corner cases the amplified tests are supposed to cover. There is an ongoing debate in the mutation testing community of whether mutation operators can serve as proxies for actual faults. Today, there is no alternative so we settled with mutation coverage. But if ever another measure for test effectiveness comes along we need to revise the results.

Internal Validity Are there unknown factors which might affect the outcome of the analyses? For RQ1 (Pull Requests), we don't have any knowledge about the policy the projects had concerning pull requests submitted by outsiders. In the Pharo community, most developers know one another and are likely to trust contributions. However, for our study it was the first author who submitted the pull-requests and at that point in time he was a newcomer in the community. For the three pull requests which were ignored, we don't know whether this newcomer submission played a role.

External Validity To what extent is it possible to generalise the findings? We demonstrated that test amplification is feasible, even for dynamically typed language. We have constructed a proof-by-construction for the Pharo/Smalltalk ecosystem. However, we cannot make any claims regarding other dynamically typed languages (Python, Ruby, Javascript, ...). We are quite confident that type profiling is the key to make test amplification successful in such a context. However, coding conventions are equally important and this may jeopardise the kind of input amplification operators that work.

Reliability (a.k.a. Conclusion Validity) Is the result dependent on the tools? As mentioned earlier, we heavily rely on mutation coverage as calculated by MUTALK. MUTALK lacks several mutation operators compared to the PITEST tool used in the DSPOT paper. This implies that SMALL-AMP will generate fewer test cases and that the newly killed mutants will also be generally lower. We mitigated this threat by always reporting the absolute number combined with the relative increase (in percentage).

The other threat to conclusion validity is the impact of randomness. SMALL-AMP works based on applying random transformations on the tests. Most importantly, the actual amplified tests surviving the Input Reduction step (see Section 4.2 – p. 18) may vary from one run to another. We applied the tool on different classes in different projects from various domains, and achieved amplified tests in all circumstances. In addition to the variety in the projects, we ran the tool three times on each test class. Thus, the impact of randomness should be small at best.

7 Related Work

Test amplification systems can vary based on the engineering goal. In addition to the amplification of the code coverage (Thummalapenta et al. 2011; Yoo and Harman 2012) or mutation score (Baudry et al. 2005; Patrick and Jia 2017), researchers have used test amplification for other goals like fault detection (Milani Fard et al. 2014; Pezze et al. 2013), oracle improvement (Fraser and Zeller 2011; Xie 2006), fault localization (Rößler et al. 2012; Xuan et al. 2015), and incompatible environments detection (Chauvel et al. 2019).

A test amplification system may use search based techniques (Baudry et al. 2005; Yoo and Harman 2012; Rojas et al. 2016; Xu et al. 2010) or symbolic and concolic execution techniques (Xu and Rothermel 2009; Yoshida et al. 2016; Thummalapenta et al. 2011). The

results of the amplification can be added to existing test suite (Baudry et al. 2005; Yoo and Harman 2012) or just modifying the current tests (Xie 2006). They also may consider only new changes (Xu et al. 2010; Xu and Rothermel 2009) rather than working on the snapshot of the entire project.

Our work is pretty close to DSPOT (Baudry et al. 2015; Danglot et al. 2019b) where SMALL-AMP is an adaption of DSPOT into a dynamic language. This work, same as DSPOT, also can be categorised under genetic improvement (Petke et al. 2018) where it takes advantage of existing test suite as well as an automated search algorithm in order to find an improved version of test code.

Brandt and Zaidman (Brandt and Zaidman 2021) use a lighter version of DSPOT to increase the instruction coverage. They also provide an IDE plugin to make the developers interplay with the test amplification tool possible. (Nijkamp et al. 2021) and (Oosterbroek et al. 2021) address the readability of the amplified tests by choosing proper names and removing redundant statements.

Dynamically-Typed Languages As we argued earlier, test case generation is well studied in statically typed languages (Pacheco et al. 2007; Fraser and Arcuri 2013; Panichella et al. 2018) but there are only a few academic works that target dynamically typed languages. We list the ones we were aware of below.

Lukasczyk et al. (2020) introduce automated unit test generation for Python, and the tool PYNGUIN which works based on techniques used in statically typed languages more specifically EVOSUITE (Fraser and Arcuri 2013) and RANDOOP (Pacheco et al. 2007). PYNGUIN circumvents the lack of type information by assuming that the system under test contains type information added by developers in terms of type annotations.

Mirshokraie et al. (2015) created a tool names JSEFT, which generate unit tests for javascript functions and events by a record and replay technique. JSEFT relies on web crawling to collect traces from javascript executions. They create test methods by replaying the executions and adding new oracles using a mutation-based process (Fraser and Zeller 2012).

Wibowo et al. (2015) use a genetic algorithm to generate unit test code for the Lua scripting language. The algorithm starts from a random initialized population and then evolve by crossover and mutation operators. The tool only generates assertions for primitive data type values. SMALL-AMP on the other hand used recursive assertion generation to deal with non-primitive types.

Mairhofer et al. (2011) introduce RUTEG, a test generation for Ruby based on genetic algorithms. For each method under test, RUTEG statically processes the parse tree and collects arguments and the list of methods invoked on each argument. Then they use predefined and customized data generators to generate a part of code that is valid based on data collected from the parse tree. RUTEG does not improve an existing test suite, but rather generates the whole test suite itself.

Mutation Testing Using mutation testing as an actionable target for strengthening an existing test suite is used at large scale at Google (Petrovic et al. 2021; Petrović and Ivanković 2018) and Facebook (Beller et al. 2021). These papers are close in spirit to DSPOT and SMALL-AMP as both create new test cases to increase mutation coverage. However, these works use professional developers to generate new tests (manual test amplification) while DSPOT and SMALL-AMP illustrate that a recommender system is feasible.

8 Future Work

In this section, we present some open problems and the ways how SMALL-AMP and test amplification tools can be improved in the future.

Test Amplification Ecosystem In the current implementation of SMALL-AMP, we run the tool on the whole project. This way of using the tool has some drawbacks: since developers should run it manually, they need a knowledge about the concept, the process and also the tool interface; it may take a long time to amplify all classes in the project; the tool will reamplify some parts of the project on each execution, which will increase the execution time; and finally, developers need to deal with the output manually, they need to understand it, polish the interesting tests, and merge them manually in the code base. Imposing such extra work on developers is likely to make them lose their interest in using the tool often.

In the future, we will integrate the SMALL-AMP with a build system, for instance GitHub Actions. The build systems will run the tool automatically on the specified events such as on each push, or pull-request or periodically. Additionally, SMALL-AMP will amplify only the recent changes on each run. It means that only the mutants in the changed parts will be generated which will reduce the cost of amplification significantly. In this case, the amplification will be run in the project level instead of running class-by-class. So, finding an exact relation between the production classes and the test classes will also lose its importance: all test methods covering a changed part can be included in the amplification.

Furthermore, we will build a web-based test editor dashboard to visualize the test amplification outputs, and also a GitHub-Bot to synchronize the Build system output, code base and the dashboard. The developer can use this dashboard to assess the outputs, and also customize the amplified test. The tests after the polishment will be reevaluated automatically, and if it is green, it can be pushed to the code base. So, developers don't need to overwhelm themselves with tedious tasks and can benefit from test amplification in an ecosystem automated by bots.

Extended Use-Cases for Dynamic Profiling Dynamic profiling is more than merely a type inference solution. It can be generalized to collect various information about unit-under-test based on dynamically running existing test suite. In some cases, statically typed languages also can benefit the profiling mechanism. We enumerate two of these use-cases:

- Pure methods detection: An impure method is a method that looks like an accessor but calling it causes a change in the state of the object (Sălcianu and Rinard 2005). In the scope of SMALL-AMP, identifying pure/impure methods is important during oracle reduction. A new profiler can be implemented using the method proxies to serialize the object state before and after method invocations. If the state is changed, we can infer that the method is impure.
- Providing information for advanced input amplifiers: In this paper we proposed a basic algorithm for test-input reduction (Section 4.2). By addressing the test-input explosion problem, test amplification tools can benefit from wider range of input amplifiers. Advanced input-amplifiers can exploit the profiling step to collect useful information dynamically and increase their knowledge about the program under test. For example, a profiler can collect all literal values from the covered methods and use them in literal values amplification operator. Another example can be object transplantation between test methods. A profiler can collect patterns of how objects are created and manipulated and this information can be used in an input-amplifier.

Test Method Models, Best Practices and Structured Strings Unit tests in object-oriented languages ideally are structured as a sequence of statements that instantiate an object, manipulates its state, and asserts expected values. However, not every test fits this model in real projects. Developers use helper methods, customized assertions, structured strings, some optimizations like grouping the tests or parallelizing them. For instance, a developer may write tests in XML files and load each file in a test method, so these tests heavily depend on parsing structured strings.

If a test does not fit with the ideal test model, the current algorithm of test amplification still can be used, but it may be less successful in producing strong results. We leave identifying and adopting best practices of test methods as an open problem. Additionally, mutating a structured string by understanding its syntax can also be interesting future work.

Using Patterns to Guide Test Amplification As an important future work, we suggest using heuristics to guide the amplification algorithm. Large scale manual test amplification histories like the work at Google (Petrovic et al. 2021; Petrović and Ivanković 2018) can be analyzed to answer questions like: How often do developers write new test methods? Are new test methods similar to an existing test? If they update an existing test method, what is the relation of the updated test method and the mutant to be killed? What transformations are applied to the test?

Answering these questions leads us to find some patterns in how real developers kill mutants. These patterns can help the tool to prioritize some test methods and input amplifiers for killing a particular mutant. Recent advances in deep learning or other program synthesis techniques are promising and can be helpful in making test amplification tools more intelligent (Abdi et al. 2019a).

Reducing the Mutation Testing Burden Test amplification generates tests to optimize mutation coverage, however calculating the mutation score is a time-consuming process. During test amplification, this mutation score is calculated multiple times for each test method: in lines 6, and 14 of Algorithm 1 – p. 11. However, for a test to kill a mutant, it first must reach the statement, then infect the program state, propagate to the output where it must be revealed by an assertion (Lu et al. 2020; Vera-Pérez et al. 2019).

We can optimize the calculation of the mutation score by using a hierarchical coverage measurement. For example, we can first run a code coverage tool, then we can run an extreme transformation (Vera-Pérez et al. 2018), afterwards we only mutate the covered parts.

Another technique for reducing the mutation testing burden is to use mutation operators that are learned from known common bugs like MUTATIONMONKEY (Beller et al. 2021).

Using Multiple Type-Inference Mechanisms The main drawback of using an existing test suite for dynamic profiling is if a method is not covered in the test, we can not collect its type information, hence can not add calls to such methods during input amplification. Static type inference (Pluquet et al. 2009) or live typing (Wilkinson 2019) techniques can be helpful to empower SMALL-AMP to generate method calls to such uncovered methods.

Involving Readability Metrics Based on a previous study (Grano et al. 2020), the most important aspects for developers in assessing the quality of a test suite are readability and maintainability. Although the code coverage metrics are necessary for identifying the poor test suites, they are limited in distinguishing high-quality tests based on how practitioners perceive the test quality. Since the goal of test amplification is to recommend new test cases ready to be merged into the code base, considering readability and maintainability metrics in the test amplification appears to be a critical next step.

9 Conclusion

In this paper, we introduce SMALL-AMP, an approach for test amplification in the dynamically typed language Pharo/Smalltalk. The main algorithm of SMALL-AMP is adapted from DSPOT, a test amplification technique designed for Java programs. In order to mitigate the lack of type information, we exploit profiling information, readily available by running the test suite. We demonstrate that test amplification is feasible for dynamically typed languages, by replicating the experimental set-up of DSPOT, including a qualitative and quantitative analysis of the improved test suite.

In our qualitative analysis, we submitted pull-requests of an amplified test by SMALL-AMP to the GitHub repositories of the projects in our dataset. From 11 pull-requests we submitted, 8 were merged ($\approx 72\%$). The developers' comments on the pull-requests illustrate how valuable they perceive the new tests created by SMALL-AMP. The results from our quantitative study show that SMALL-AMP succeeds to amplify 28 test classes out of 52, approximately 53% of target classes, in 13 projects from our dataset. The majority of the generated tests are focused, and all test amplification steps (including type profiling step) play a critical role. The results from SMALL-AMP and the results from DSPOT are comparable. We see $\approx 72\%$ merged pull-requests and 53% successfully amplified test classes in SMALL-AMP, while for DSPOT these values are 68% and 65%. We also see that the value of total increase killed between two tools in two different ecosystems are similar (14% in SMALL-AMP and 20% in DSPOT).

In conclusion, the results of experiments show that by using a profiling step and collecting type information, we can successfully adopt a test amplification approach in a dynamically typed language.

Acknowledgements This work is supported by (a) the Fonds de la Recherche Scientifique-FNRS and the Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO) under EOS Project 30446992 SECO-ASSIST (b) Flanders Make vzw, the strategic research centre for the manufacturing industry. Bergel is very grateful to Lam Research and Fondecyt Regular 1200067 for partially sponsoring this work.

This work originated during a sabbatical leave by Serge Demeyer in the RMOD lab. Stéphane Ducasse, Marcus Denker and Julien Delplanque helped a lot with the intricacies of Pharo/Smalltalk; Pavel Krivanek helped with MUTALK. Finally, we thank the developers reviewing and discussing our pull requests, they really helped us in improving SMALL-AMP.

Funding This work is supported financially by two public funding organisation in Belgium and Chile. There is no direct or indirect industrial support for the research reported here.

Declarations

Conflict of Interests We confirm that there are no conflicts of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit.

References

- Abdi M, Rocha H, Demeyer S (2019a) Adopting program synthesis for test amplification. In: Proceedings of the 18th Belgium-Netherlands software evolution workshop, Brussels, Belgium. published at <http://ceur-ws.org>. <http://ceur-ws.org/Vol-2605/11.pdf>
- Abdi M, Rocha H, Demeyer S (2019b) Test amplification in the pharo smalltalk ecosystem. In: Proceedings of the 14th edition of the international workshop on smalltalk technologies, IWST, vol 19, pp 1–7
- Agibalov A (2015) What is a normal “functional lines of code” to “test lines of code” ratio? Teslarati. [on line] <https://softwareengineering.stackexchange.com/questions/156883/> — last Accessed In April 2021
- Athanasiou D, Nugroho A, Visser J, Zaidman A (2014) Test code quality and its relation to issue handling performance. *IEEE Trans Softw Eng* 40(11):1100–1125. <https://doi.org/10.1109/TSE.2014.2342227>
- Baudry B, Allier S, Rodríguez-Cancio M, Monperrus M (2015) Dspot: Test amplification for automatic assessment of computational diversity. CoRR: a computing research repository. arXiv:1503.05807
- Baudry B, Fleurey F, Jézéquel J-M, Le Traon Y (2005) From genetic to bacteriological algorithms for mutation-based testing. *Software Testing, Verification and Reliability* 15(2):73–96
- Beller M, Wong C-P, Bader J, Scott A, Machalica M, Chandra S, Meijer E (2021) What it would take to use mutation testing in industry—a study at facebook. In: 2021 IEEE/ACM 43rd international conference on software engineering: software engineering in practice (ICSE-SEIP), pp 268–277
- Bergel A, Cassou D, Ducasse S, Laval J (2013) Deep into pharo. Square Bracket Associates. <http://books.pharo.org/deep-into-pharo/>
- Black AP, Nierstrasz O, Ducasse S, Pollet D (2010) Pharo by example. Lulu. com
- Brandt C, Zaidman A (2021) Developer-centric test amplification the interplay between automatic generation and human exploration. arXiv:2108.12249
- Chauvel F, Morin B, Garcia-Ceja E (2019) Amplifying integration tests with camp. In: 2019 IEEE 30th international symposium on software reliability engineering (ISSRE), pp 283–291
- Costiou S, Aranega V, Denker M (2020) Sub-method, partial behavioral reflection with reflectivity: Looking back on 10 years of use. *The Art, Science, and Engineering of Programming*, 4(3)
- Danglot B, Vera-Perez O, Yu Z, Zaidman A, Monperrus M, Baudry B (2019a) A snowballing literature study on test amplification. *J Syst Softw* 157:110398. <https://doi.org/10.1016/j.jss.2019.110398>. <http://www.sciencedirect.com/science/article/pii/S0164121219301736>
- Danglot B, Vera-Pérez OL, Baudry B, Monperrus M (2019b) Automatic test improvement with dspot: a study with ten mature open-source projects. *Empirical Software Engineering*. Springer, Berlin
- Daniel B, Jagannath V, Dig D, Marinov D (2009) Reassert: Suggesting repairs for broken unit tests. In: Proceedings ASE 2009 (international conference on automated software engineering). IEEE CS, pp 433–444
- Delplanque J, Ducasse S, Polito G, Black AP, Etien A (2019) Rotten green tests. In: 41th international conference on software engineering, ICSE '19. IEEE, pp 500–511. <https://hal.inria.fr/hal-02002346>
- Denker M, Ducasse S, Lienhard A, Marschall P (2007) Sub-method reflection. *Journal of Object Technology* 6:275–295. <https://doi.org/10.5381/jot.2007.6.9.a14>. http://www.jot.fm/contents/issue_2007_10_paper14.html
- Ducasse S (1999) Evaluating message passing control techniques in smalltalk. *Journal of Object Oriented Programming* 12:39–50
- Ducasse S (2019) Pharo with style. Square Bracket Associates
- Fischer S, Haslinger EN, Zimmermann M, Thaller H (2020) An empirical evaluation for object initialization of member variables in unit testing. In: Proceedings VST 2020 (IEEE workshop on validation, analysis and evolution of software tests), pp 8–11
- Fraser G, Arcuri A (2012) The seed is strong: Seeding strategies in search-based software testing. In: 2012 IEEE fifth international conference on software testing, verification and validation, pp 121–130
- Fraser G, Arcuri A (2013) Whole test suite generation. *IEEE Trans Softw Eng* 39(2):276–291. <https://doi.org/10.1109/TSE.2012.14>
- Fraser G, Zeller A (2012) Mutation-driven generation of unit tests and oracles. *IEEE Trans Softw Eng* 38(2):278–292. <https://doi.org/10.1109/TSE.2011.93>
- Fraser G, Zeller A (2011) Generating parameterized unit tests. In: Proceedings of the 2011 international symposium on software testing and analysis, pp 364–374
- Goldberg A, Robson D (1983) Smalltalk-80: The language and its implementation. Addison-Wesley Longman Publishing Co Inc, USA
- Grano G, De Iaco C, Palomba F, Gall HC (2020) Pizza versus pinsa: On the perception and measurability of unit test code quality. In: 2020 IEEE international conference on software maintenance and evolution (ICSME), pp 336–347

- Li N, Offutt J (2016) Test oracle strategies for model-based testing. *IEEE Trans Softw Eng* 43(4):372–395. <https://doi.org/10.1109/TSE.2016.2597136>
- Lu ZX, Vercammen S, Demeyer S (2020) Semi-automatic test case expansion for mutation testing. In: 2020 IEEE workshop on validation, analysis and evolution of software tests (VST). IEEE, pp 1–7
- Lukasczyk S, Kroiß F, Fraser G (2020) Automated unit test generation for python. In: Aleti A, Panichella A (eds) Search-based software engineering. Springer International Publishing, Cham, pp 9–24
- Mairhofer S, Feldt R, Torkar R (2011) Search-based software testing and test data generation for a dynamic programming language. In: Proceedings of the 13th annual conference on genetic and evolutionary computation, GECCO '11. Association for Computing Machinery, New York, pp 1859–1866. <https://doi.org/10.1145/2001576.2001826>
- Milani Fard A, Mirzaaghaei M, Mesbah A (2014) Leveraging existing tests in automated test generation for web applications. In: Proceedings of the 29th ACM/IEEE international conference on Automated software engineering, pp 67–78
- Mirshokraie S, Mesbah A, Pattabiraman K (2015) Jseft: Automated javascript unit test generation. In: 2015 IEEE 8th international conference on software testing, verification and validation (ICST), pp 1–10
- Nijkamp N, Brandt C, Zaidman A (2021) Naming amplified tests based on improved coverage. In: 2021 IEEE 21st international working conference on source code analysis and manipulation (SCAM). IEEE, pp 237–241
- Oosterbroek W, Brandt C, Zaidman A (2021) Removing redundant statements in amplified test cases. In: 2021 IEEE 21st international working conference on source code analysis and manipulation (SCAM), IEEE, pp 242–246
- Pacheco C, Lahiri SK, Ernst MD, Ball T (2007) Feedback-directed random test generation. In: 29th international conference on software engineering (ICSE'07), pp 75–84
- Panichella A, Kifetew FM, Tonella P (2018) Automated test case generation as a many-objective optimisation problem with dynamic selection of the targets. *IEEE Trans Softw Eng* 44(2):122–158. <https://doi.org/10.1109/TSE.2017.2663435>
- Papadakis M, Kintis M, Zhang J, Jia Y, Traon YL, Harman M (2019) Mutation testing advances: An analysis and survey. *Adv Comput* 112:275–378. <https://doi.org/10.1016/bs.adcom.2018.03.015>. <http://www.sciencedirect.com/science/article/pii/S0065245818300305>
- Patrick M, Jia Y (2017) Kd-art: Should we intensify or diversify tests to kill mutants? *Inf Softw Technol* 81:36–51
- Peck MM, Bouraqadi N, Fabresse L, Denker M, Teruel C (2015) Ghost: A uniform and general-purpose proxy implementation. *Sci Comput Program* 98:339–359
- Petke J, Haraldsson SO, Harman M, Langdon WB, White DR, Woodward JR (2018) Genetic improvement of software: A comprehensive survey. *IEEE Trans Evol Comput* 22(3):415–432. <https://doi.org/10.1109/TEVC.2017.2693219>
- Petrović G, Ivanković M (2018) State of mutation testing at google. In: Proceedings of the 40th international conference on software engineering: software engineering in practice, ICSE-SEIP '18. Association for Computing Machinery, New York, pp 163–171. <https://doi.org/10.1145/3183519.3183521>
- Petrovic G, Ivankovic M, Fraser G, Just R (2021) Practical mutation testing at scale: A view from google. *IEEE Trans Softw Eng*, 1–1. <https://doi.org/10.1109/TSE.2021.3107634>
- Pezze M, Rubinov K, Wuttke J (2013) Generating effective integration test cases from unit ones. In: 2013 IEEE Sixth international conference on software testing, verification and validation, IEEE, pp 11–20
- Pluquet F, Marot A, Wuyts R (2009) Fast type reconstruction for dynamically typed programming languages. In: Noble J (ed) Proceedings of the 5th symposium on dynamic languages, DLS 2009, october 26, 2010. ACM, Orlando, pp 69–78. <https://doi.org/10.1145/1640134.1640145>
- Rößler J, Fraser G, Zeller A, Orso A (2012) Isolating failure causes through test case generation. In: Proceedings of the 2012 international symposium on software testing and analysis, pp 309–319
- Rojas JM, Fraser G, Arcuri A (2016) Seeding strategies in search-based unit test generation. *Software Testing, Verification and Reliability* 26(5):366–401
- Sălcianu A, Rinard M (2005) Purity and side effect analysis for java programs. In: International workshop on verification, model checking, and abstract interpretation. Springer, pp 199–215
- Thummalapenta S, Marri MR, Xie T, Tillmann N, De Halleux J (2011) Retrofitting unit tests for parameterized unit testing. In: International conference on fundamental approaches to software engineering. Springer, pp 294–309
- Tillmann N, Schulte W (2006) Unit tests reloaded: Parameterized unit testing with symbolic execution. *IEEE Softw*, 23(4). <https://doi.org/10.1109/MS.2006.117>
- Tonella P (2004) Evolutionary testing of classes. In: Proceedings of the 2004 ACM SIGSOFT international symposium on Software testing and analysis - ISSTA '04. <https://doi.org/10.1145/1007512.1007528>

- Vera-Pérez OL, Danglot B, Monperrus M, Baudry B (2019) Suggestions on test suite improvements with automatic infection and propagation analysis. arXiv:1909.04770
- Vera-Pérez OL, Monperrus M, Baudry B (2018) Descartes: A pitest engine to detect pseudo-tested methods: Tool demonstration. In: 2018 33rd IEEE/ACM international conference on automated software engineering (ASE), pp 908–911
- Wibowo JTP, Hendradjaya B, Widayani Y (2015) Unit test code generator for lua programming language. In: 2015 international conference on data and software engineering (ICoDSE), pp 241–245
- Wilkinson H (2019) Vm support for live typing: Automatic type annotation for dynamically typed languages. In: Proceedings of the Conference Companion of the 3rd International Conference on Art, Science, and Engineering of Programming, Programming '19. Association for Computing Machinery, New York. <https://doi.org/10.1145/3328433.3328443>
- Wilkinson H, Chillo N, Brunstein G (2009) Mutation testing. European Smalltalk User Group (ESUG 09). Brest, France. http://www.esug.org/data/ESUG2009/Friday/Mutation_Testing.pdf
- Wohlin C, Runeson P, Höst M, Ohlsson MC, Regnell B, Wesslén A (2000) Experimentation in software engineering. Kluwer Academic Publishers, Boston
- Xie T (2006) Augmenting automatically generated unit-test suites with regression oracle checking. Lect Notes Comput Sci, 380–403. https://doi.org/10.1007/11785477_23
- Xu Z, Cohen MB, Rothermel G (2010) Factors affecting the use of genetic algorithms in test suite augmentation. In: Proceedings of the 12th annual conference on Genetic and evolutionary computation, pp 1365–1372
- Xu Z, Rothermel G (2009) Directed test suite augmentation. In: 2009 16th asia-pacific software engineering conference. IEEE, pp 406–413
- Xuan J, Xie X, Monperrus M (2015) Crash reproduction via test case mutation: let existing test cases help. In: Proceedings of the 2015 10th joint meeting on foundations of software engineering, pp 910–913
- Yoo S, Harman M (2012) Test data regeneration: generating new test data from existing test data. Software Testing, Verification and Reliability 22(3):171–201
- Yoshida H, Tokumoto S, Prasad MR, Ghosh I, Uehara T (2016) Fsx: fine-grained incremental unit test generation for c/c++ programs. In: Proceedings of the 25th international symposium on software testing and analysis, pp 106–117
- Zaidman A, Rompaey BV, van Deursen A, Demeyer S (2011) Studying the co-evolution of production and test code in open source and industrial developer test processes through repository mining. International Journal on Empirical Software Engineering 16(3):325–364. <https://doi.org/10.1007/s10664-010-9143-7>