



7-17-2019

Tests as Maintainable Assets Via Auto-generated Spies: A case study involving the Scala collections library's Iterator trait

Konstantin Läufer
Loyola University Chicago

John O'Sullivan
Loyola University Chicago, josullivan1@luc.edu

George K. Thiruvathukal
Loyola University Chicago, gkt@cs.luc.edu

Author Manuscript

This is a pre-publication author manuscript of the final, published article.

Recommended Citation

Konstantin Läufer, John O'Sullivan, and George K. Thiruvathukal. 2019. Tests as Maintainable Assets Via Auto-generated Spies. In Proceedings of Tenth ACM SIGPLAN Scala Symposium, London, United Kingdom, July 17, 2019 (Scala '19), 6 pages.

This Conference Proceeding is brought to you for free and open access by the Faculty Publications at Loyola eCommons. It has been accepted for inclusion in Computer Science: Faculty Publications and Other Works by an authorized administrator of Loyola eCommons. For more information, please contact ecommons@luc.edu.

[Creative Commons License](#)

This work is licensed under a [Creative Commons Attribution-Share Alike 4.0 License](#).

© 2019 Association for Computing Machinery

Tests as Maintainable Assets Via Auto-generated Spies

A case study involving the Scala collections library's `Iterator` trait

Konstantin Läufer
Loyola University Chicago
laufer@cs.luc.edu

John O'Sullivan
Loyola University Chicago
josullivan@cs.luc.edu

George K. Thiruvathukal
Loyola University Chicago and
Argonne National Laboratory
gkt@cs.luc.edu, gkt@anl.gov

Abstract

In testing stateful abstractions, it is often necessary to record interactions, such as method invocations, and express assertions over these interactions. Following the *Test Spy* design pattern, we can reify such interactions programmatically through additional mutable state. Alternatively, a mocking framework, such as Mockito, can *automatically generate* test spies that allow us to record the interactions and express our expectations in a declarative domain-specific language. According to our study of the test code for Scala's `Iterator` trait, the latter approach can lead to a significant reduction of test code complexity in terms of metrics such as code size (in some cases over 70% smaller), cyclomatic complexity, and amount of additional mutable state required. In this tools paper, we argue that the resulting test code is not only more maintainable, readable, and intentional, but also a better stylistic match for the Scala community than manually implemented, explicitly stateful test spies.

ACM Reference Format:

Konstantin Läufer, John O'Sullivan, and George K. Thiruvathukal. 2019. Tests as Maintainable Assets Via Auto-generated Spies: A case study involving the Scala collections library's `Iterator` trait. In *Tenth ACM SIGPLAN Scala Symposium (Scala '19)*, July 17, 2019, London, United Kingdom. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3337932.3338814>

1 Introduction

In our university-level programming languages course [11], the `scanLeft` method (similar to `foldLeft` with intermediate results) in the Scala collections library is an important part of our overall pedagogy and the subject of many examples, including running averages and other forms of sliding

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Scala '19, July 17, 2019, London, United Kingdom

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6824-7/19/07...\$15.00

<https://doi.org/10.1145/3337932.3338814>

analysis on unbounded streams, e.g., dynamic word clouds, stock market analysis, etc. In general, such prefix scans are useful and efficient building blocks for interactive, event-based, and other stream processing systems.

In this—and other—courses, we also emphasize the notion of *tests as assets*. The software development community increasingly views automated tests as longer-term, maintainable assets along with the production code itself [9, 13], and a body of work on design patterns for automated testing has emerged [14]. In particular, a *Mock Object* replaces an object the system-under-test (SUT) depends on, is usually preconfigured to provide certain behaviors the SUT expects, and dynamically verifies the expected interactions coming from the SUT. By contrast, a *Test Spy* also takes the place of a dependency of the SUT but behaves like the original dependency while recording the SUT's *indirect outputs*, i.e., interactions with the dependency in terms of method invocation frequency and arguments, for later verification. Both of these are subpatterns of *Test Double*.

In the context of our course, we noticed and reported a seven-year-old bug in the `scanLeft` method of the `Iterator` trait, which provides some lazy stateful behaviors that are challenging to test. Indeed, the original test for `Iterator.scanLeft` does not fully test the correctness of this method under certain conditions. This led us to study the Scala collections library's source code, where we noticed that the test suite includes several instances of manually implemented, explicitly stateful test spies. While the corrected `scanLeft` implementation was successful in terms of clarity, conciseness, and idiomatic style, we found it difficult to understand the code and the actual cases being tested.

In this paper, we look for opportunities to use *automatically generated* test spies as a systematic way to improve the tests for `scanLeft` and other `Iterator` methods and bring them in line with the notion of *tests as assets*. While this serves as the underlying case study for this paper, the technique we describe is of general value as a *programming pearl*. More broadly, it brings *existing* tools and techniques to the Scala community in the hope that they will be useful. We have organized the rest of the paper as follows:

- a detailed explanation of the case study based on the `scala.collection.Iterator` trait
- an overview of auto-generated test spies and how they can replace manually implemented test dependencies

```

1 object CumAvgFunctional extends App {
2   val lines = scala.io.Source.stdin.getLines
3   val values = lines.map(_.toDouble)
4   val results = values.scanLeft((0, 0.0)) {
5     case ((count, sum), value) =>
6       (count + 1, sum + value)
7   }
8   results.drop(1).foreach { case (count, sum) =>
9     println(count + ": " + (sum / count))
10  }
11 }

```

Figure 1. The imperative version of the cumulative running average filter is a simple, monolithic while loop. The functional version shown here is a pipeline of modular, separately testable stages, arguably making it more readable and maintainable. Both run in linear time and constant space.

- a side-by-side comparison of the official `scanLeft` test vs. an equivalent test with an auto-generated spy
- a comparison of complexity metrics, including code size and cyclomatic complexity (a quantitative measure of the number of linearly independent paths through a program's source code) [12], before and after refactoring several tests from manually-implemented to auto-generated test spies, and
- a summary of human and technical challenges that must be addressed for the proposed approach to be used more widely by Scala developers.

The novelty of our work is an initial exploration to identify, understand, and improve test assets, focused on a “by hand” analysis of the test code complexity. Although our efforts have been focused on the Scala library itself, owing to the critical role Scala's core library plays in all Scala development, the methods we describe are of general value and also play a role in our own development efforts. We are not aware of other work focused specifically on understanding test code complexity and using the resulting metrics to improve test assets. We believe our work has important implications in education and professional Scala development by training the next generation of Scala developers how to write effective, comprehensible, and maintainable tests.

2 Case study: `Iterator.scanLeft`

The `scala.collection.Iterator` trait did not have a `scanLeft` method until it was requested in Scala Issue 4054 [4] and implemented in January 2011 as shown in Figure 3.

This version, however, causes the example from Figure 1 to behave incorrectly as seen in Figure 2b: Instead of printing the first updated average right after reading the first value, it prints this only after reading the second value; it then prints each subsequent update delayed by one input value, and the final update only *after* EOF. The reason is that the iterator returned by `scanLeft` does not return the current item until

after the (premature) call to `self.next()` on line 600 in Figure 3 returns. The included test does not catch this bug because it focuses on the correctness of the resulting items, irrespective of the interactions with the original iterator.

We reported this bug as Scala Issue 10709 [6] in February 2018, after attempting to use `scanLeft` in the context of our spring 2018 Scala-based programming languages course [11]. The Scala team promptly fixed this issue as of Scala 2.12.5, by reimplementing the method using a flat four-state machine, replaced as of Scala 2.13.x [7] with an arguably more elegant and straightforward implementation based on the State pattern [3].

The corresponding test, shown in Figure 4a, uses a test spy in the form of a custom iterator with additional state to test for incremental correctness along with the “right amount of laziness.”

Discussion There are several possible factors contributing to the fact that this bug remained unreported for so long:

- `Iterator.scanLeft` might be rarely used, or rarely used *incrementally* as described above.
- Developers might have resorted to a workaround but not taken the time to report the actual bug.
- State-dependent behaviors are challenging to comprehend, document, implement, and test.

Furthermore, there is considerable *essential complexity* [2] to a stateful behavior such as `scanLeft`: it returns a decorator [3] around the receiver, after which one is no longer allowed to interact directly with the receiver, while subsequent interactions with the decorator have side effects on the original receiver, such as reading lines from input. This complexity carries over to the spy-based test in terms of the dynamic interactions shown in Figure 5.

We conjecture that tests that rely on manually implemented spies typically suffer from two additional shortcomings: (1) *accidental complexity* that affects maintainability, and (2) unhelpful failure messages. We will now explore how automatically generated spies can help address both issues.

3 Auto-generated test spies can help

As an alternative to manually implementing test doubles (see Section 1) to represent the SUT's dependencies, mocking frameworks support the automatic generation of mock objects, usually based on their type. Some mocking frameworks, such as Mockito [5], additionally support the automatic generation of *test spies*. We chose Mockito because of its maturity, support for idiomatic Scala syntax [8], and unique support for spying on final and anonymous classes.

Concretely, as shown in Figure 4b, we can use Mockito to wrap a test spy around a simple iterator instance, invoke `scanLeft`, and then interact with the iterator resulting from this invocation; these interactions still correspond to Figure 5. During these interactions, we can test not only the (overall) correctness of the values of the resulting iterator, but also the

```

$ ./target/universal/stage/bin/cum-avg-imp
> 6
1: 6.0
> 7
2: 6.5
> 2
3: 5.0
> ^D
(a) correct behavior

```

```

$ ./target/universal/stage/bin/cum-avg-fun
> 6
> 7
1: 6.0
> 2
2: 6.5
> ^D
3: 5.0
(b) incorrect behavior

```

Figure 2. Correct (left) and incorrect (right) sample runs of the cumulative running average filters. In the incorrect case, the updated average goes up to the previous instead of the current input value. (We prefix input lines with >.)

```

595 def scanLeft[B](z: B)(op: (B, A) => B): Iterator
    [B] = new Iterator[B] {
596   var hasNext = true
597   var elem = z
598   def next() = if (hasNext) {
599     val res = elem
600     if (self.hasNext) elem = op(elem, self.next())
601     else hasNext = false
602     res
603   } else Iterator.empty.next()
604 }

```

```

9 object Test {
10  def main(args: Array[String]) {
11    val it = Iterator.from(1).map(n => n * n).
      scanLeft(0)(_+_ )
12
13    assert(it.next == 0)
14    assert(it.next == 1)
15    assert(it.next == 5)
16    // etc.
17  }
18 }

```

Figure 3. Original version of `scanLeft` and associated correctness test. In this implementation, the resulting iterator does not return the current item until *after* the (premature) call to `self.next()` on line 600 returns (shown in boldface above). The test does not catch this bug because it focuses on the overall correctness of the sequence of items returned.

Method under test	Test method	Mut _b	LOC _b	LOC _a	LOC _Δ	CC _b	CC _a
Iterator.sliding	IteratorTest.groupedIteratorShouldNotAskForUnneededElement	2	10	4	-60%	3	1
Iterator.++	IteratorTest.noExcessiveHasNextInJoinIterator	2	15	8	-47%	6	2
Iterator.toStream	IteratorTest.toStreamIsSufficientlyLazy	1	9	7	-22%	3	3
Iterator.scanLeft	IteratorTest.`scan is lazy enough`	3	26	7	-73%	5	3
HashMap.getOrElseUpdate	HashMapTest.getOrElseUpdate_evalOnce	1	5	4	-20%	2	2

Table 1. Comparison of test code complexity metrics (number of mutable objects, LOC, and Cyclomatic Complexity counted manually) before/after refactoring from manually implemented to auto-generated test spies. (Mut_a = 0 for all test methods.) Even when there is no or little reduction in the metrics, there is usually a gain in readability and clarity of intent. This table includes most affected tests in `IteratorTest` and `HashMapTest`; there may be other applicable tests elsewhere in the codebase.

(incremental) correctness of the effects of these interactions on the original iterator. In this case, the correct amount of laziness means to have invoked `next()` on the original iterator *i* times, where *i* is the position of the current value in the resulting iterator. Specifically, when we invoke `next()` on the resulting iterator for the first time, we expect to see the initial *z* value of `scanLeft`, and there should not yet have been any invocations of `next()` on the original iterator.

This test is three times as short as the official version. We argue that it is not only more comprehensible, maintainable, and effective at conveying the intent, but it also produces a

more useful error message pointing directly to the offending eager invocation of `next()`.

Discussion By eliminating the need for custom iterator implementations with mutable state, auto-generated test spies allow the programmer to focus on the functional correctness of the SUT. This promotes the view of test code as readable, comprehensible, teachable, and maintainable assets, which has the potential to work its way back into the API documentation. If more developers feel encouraged to write tests, this might foster a test-driven mindset in the community.

```

300 @Test def `scan is lazy enough`() : Unit = {
301   val results = ListBuffer.empty[Int]
302   val it = new AbstractIterator[Int] {
303     var cur = 1 ; val max = 3
304     override def hasNext = {
305       results += -cur ; cur < max }
306     override def next() = {
307       val res = cur ; results += -res
308       cur += 1 ; res }
309   }
310   val xy = it.scanLeft(10)((sum, x) => {
311     results += -(sum + x) ; sum + x
312   })
313   val scan = ListBuffer.empty[Int]
314   for (i <- xy) { scan += i ; results += i }
315   assertEquals(List(10,11,13), scan)
316   assertEquals(List
317     (10,-1,-1,-11,11,-2,-2,-13,13,-3), results)
318 }

```

```

java.lang.AssertionError: expected:<List(10, -1,
-1, -11, 11, -2, -2, -13, 13, -3)> but was:<
ListBuffer(-1, -1, -11, 10, -2, -2, -13, ...)>
at org.junit.Assert.fail(Assert.java:88)
at AssertUtil.assertEquals(...)
at IteratorTest.scan_is_lazy_enough(IteratorTest
.scala:316) <-- line 316 above

```

(a) test with manually implemented spy

```

9 @Test def `scan is lazy enough with spy`() = {
10   val it = spy(Iterator(1, 2, 3))
11   val expected = Array(0, 1, 3, 6)
12   val result = it.scanLeft(0)(_ + _)
13   for (i <- expected.indices) {
14     result.next() shouldBe expected(i)
15     it.next() wasCalled i.times
16   }
17 }

```

```

org.mockito.exceptions.verification.
  NeverWantedButInvoked:
elements.next();
Never wanted here:
-> at scala.collection.IndexedSeqLike$Elements.
  next(IndexedSeqLike.scala:59)
But invoked here:
-> at scala.collection.Iterator$$anon$14.next(
  Iterator.scala:600) <-- line 600 in Figure 3

```

```

at scala.collection.IndexedSeqLike$Elements.next
(IndexedSeqLike.scala:59)
at OurIteratorTest.
  $anonfun$scan_is_lazy_enough_with_spy$2(
  OurIteratorTest.scala:15) <-- line 15 above

```

(b) test with auto-generated spy

Figure 4. Incremental correctness and laziness tests for `Iterator.scanLeft`. (a) Manually implemented, explicitly stateful spy defined (interwoven with the SUT itself) on lines 301–309 and 311, exercised on line 314, and verified on line 316. (b) Auto-generated, declarative spy using Mockito Scala defined on line 10, exercised on lines 13–14, and verified on line 15. We argue that the failure message for (b) is more helpful by directly indicating the site of the unwanted invocation of `next()`.




Figure 5. Collaboration diagram for a spy-based test of a stateful SUT's method that returns a decorator of the SUT. The test verifies through the spy that the interactions with the MUT's result have the desired effect on the original SUT.

Table 1 shows other test methods that can benefit from this approach; those typically exhibit *code smells* such as `var` or mutable state other than the stateful SUT, e.g., a buffer. Because Test Spy is a common subpattern of Test Double [14], we expect the technique of spying on a stateful SUT to be beneficial in similar scenarios to Figure 5, even when the result of the method-under-test (MUT) depends on the SUT in more general ways than decorating/wrapping the SUT. For instance, we have used generated test spies for unit testing a view component in a Scala-based Android app [10].

4 Conclusions and Future Work

We have shown that Mockito's auto-generated test spies eliminate much accidental complexity from certain state-based tests. We plan to conduct a similar investigation of Scala `view`, `Stream`, and the `LazyList` class added in 2.13.x.

More broadly, we hope to use repository mining to identify other Scala projects that can benefit from making tests more immutable and declarative, understand the effect of these refactorings on test code complexity/quality [1] and process metrics [15], and investigate the possibility of tool support.

Finally, it turns out that a test coverage tool would not have indicated the problem with the original test for `Iterator.scanLeft` shown in Figure 3. Common metrics, such as statement and branch coverage, remain unchanged (after compensating for a finite vs. indefinite iterator as SUT). Further study might reveal whether suitable existing test coverage metrics are effective for these complex stateful behaviors. Nevertheless, taming test code complexity using test spies can be an effective strategy for improving the comprehension and maintainability of test cases.

References

- [1] Len Bass, Paul Clements, and Rick Kazman. 2012. *Software Architecture in Practice* (3rd ed.). Addison-Wesley Professional.
- [2] Frederick P. Brooks, Jr. 1987. No Silver Bullet—Essence and Accidents of Software Engineering. *Computer* 20, 4 (April 1987), 10–19. <https://doi.org/10.1109/MC.1987.1663532>
- [3] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. 1995. *Design Patterns: Elements of Reusable Object-oriented Software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [4] GitHub. 2010. missing scanLeft in Iterator #4054. (Dec. 2010). <https://github.com/scala/bug/issues/4054>
- [5] GitHub. 2017. Mockito 2.x. <http://site.mockito.org/>. (2017).
- [6] GitHub. 2018. Iterator.scanLeft has incorrect element-by-element behavior (delayed by one item) #10709. (Feb. 2018). <https://github.com/scala/bug/issues/10709>
- [7] GitHub. 2018. scala/collection-strawman: Implementation of the new Scala 2.13 Collections. (March 2018). <https://github.com/scala/collection-strawman>
- [8] GitHub. 2019. Mockito Scala 1.x. <https://github.com/mockito/mockito-scala>. (2019).
- [9] Danielle Gonzalez, Joanna C. S. Santos, Andrew Popovich, Mehdi Mirakhorli, and Mei Nagappan. 2017. A Large-scale Study on the Usage of Testing Patterns That Address Maintainability Attributes: Patterns for Ease of Modification, Diagnoses, and Comprehension. In *Proceedings of the 14th International Conference on Mining Software Repositories (MSR '17)*. IEEE Press, Piscataway, NJ, USA, 391–401. <https://doi.org/10.1109/MSR.2017.8>
- [10] Konstantin Läufer. 2014. Android Stopwatch example for COMP 413: Intermediate Object-Oriented Programming (Scala-based fall 2014 section). <https://github.com/loyolachicagocode/stopwatch-android-scala>. (2014).
- [11] Konstantin Läufer. 2018. Lecture Notes: Theory (and Practice) of Programming Languages. <http://lucproglanguagecourse.github.io>. (2018). Loyola University Chicago.
- [12] Thomas J. McCabe. 1976. A Complexity Measure. In *Proceedings of the 2Nd International Conference on Software Engineering (ICSE '76)*. IEEE Computer Society Press, Los Alamitos, CA, USA, 407–. <http://dl.acm.org.flagship.luc.edu/citation.cfm?id=800253.807712>
- [13] John D. McGregor. 2002. Building Reusable Test Assets for a Product Line. In *Software Reuse: Methods, Techniques, and Tools*, Cristina Gacek (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 345–346.
- [14] Gerard Meszaros. 2006. *XUnit Test Patterns: Refactoring Test Code*. Prentice Hall PTR, Upper Saddle River, NJ, USA.
- [15] George K. Thiruvathukal, Shilpika, Nicholas J. Hayward, and Konstantin Läufer. 2018. Metrics Dashboard: A Hosted Platform for Software Quality Metrics. *CoRR* abs/1804.02053 (2018). arXiv:1804.02053 <http://arxiv.org/abs/1804.02053>