

BlueCov: Integrating Test Coverage and Model Checking with JBMC

Matthias Güdemann

Munich University of Applied Sciences HM, Germany
Germany

Peter Schrammel

University of Sussex and Diffblue Ltd., UK
United Kingdom

ABSTRACT

Automated test case generation tools help businesses to write tests and increase the safety net provided by high regression test coverage when making code changes. Test generation needs to cover as much as possible of the uncovered code while avoiding generating redundant tests for code that is already covered by an existing test-suite.

In this paper we present our work on a tool for the real world application of integrating formal analysis with automatic test case generation. The test case generation is based on coverage analysis using the Java bounded model checker (JBMC). Counterexamples of the model checker can be translated into Java method calls with specific parameters.

In order to avoid the generation of redundant tests, it is necessary to measure the coverage in the *exact same* way as JBMC generates its coverage goals. Each existing coverage measurement tool uses a slightly different instrumentation and thus a different coverage criterion. This makes integration with a test case generator based on formal analysis difficult. Therefore, we developed BlueCov as a specific runtime coverage measurement tool which uses the exact same coverage criteria as JBMC does. This approach also allows for incremental test-case generation, only generating test coverage for previously untested code, e.g., to complete existing test suites.

CCS CONCEPTS

• **Software and its engineering** → **Formal software verification; Software testing and debugging.**

KEYWORDS

model-checking, test coverage, Java

1 INTRODUCTION

In recent years there has been steady progress in developing automated formal verification tools. Promising tools have emerged, not only targeting the traditional domain of embedded C programs, but also for languages such as Java that are predominantly used for larger enterprise software systems [4]. For such systems, testing is the primary way of obtaining confidence in the correct functioning of the system [12]. While it is far from feasible to formally verify such systems, automated verification tools can be used to verify small subsystems or support the automation of the testing process. An important use case is to automatically generate tests for untested parts of the system [8, 11]. For that purpose, information about uncovered code needs to be communicated to the verification tool in order to reduce the computational effort and avoid generation of redundant tests.

In this paper, we present the BlueCov tool which solves this communication problem in practice for Java programs using the

bounded model checker, JBMC [4]. Similar ideas may apply to other programming languages.

In Java, the imperative code is first compiled into a stack-based language, the *Java bytecode*, which is interpreted by the Java Virtual Machine (JVM). This compilation step abstracts away a great deal of syntactic sugar and language-level complications. Also, libraries which make up the vast majority of any real-world Java application are supplied and linked on the bytecode level. Therefore, verification tools also perform their analysis based on bytecode as input. The stack-based bytecode is translated into a control flow graph (CFG) to perform further analysis, such as abstract interpretation or symbolic execution.

Code coverage measurement is performed by bytecode-level instrumentations. This information is used to avoid generating tests for code that is already covered by an existing test-suite. The problem is now that the code structure on the bytecode level differs from the CFG used by the verification tool in details that are relevant for communicating structural code coverage information. Hence, it is not straightforward to use test coverage information obtained from coverage measurement tools, such as JaCoCo [9], in a verification tool. An exact correspondence is required, however, in order to satisfy the desired property of the analysis such as pruning paths and avoiding the generation of redundant tests.

One way of achieving this is to pick a coverage measurement tool and implement the exact same coverage criterion inside the test generator, including all the oddities of the coverage measurement tool and keep it up-to-date as the tool changes. If we wanted to support further coverage criteria, e.g. modified condition / decision coverage (MC/DC), then we would need to find an appropriate mapping from the existing coverage instrumentation or even modify its instrumentation if it does not capture all information required.

We took a more maintainable and flexible approach here by developing BlueCov, which takes the coverage goals *from the test generator* and performs the instrumentation accordingly. The coverage criterion is hence defined by the test generator, i.e., JBMC in our case.

BlueCov is much simpler than existing coverage measurement tools, also because it does not require computing coverage percentages. Coverage percentages are highly dependent on the exact coverage criterion and developers are often confused by mismatches between numbers reported from different tools. We therefore do not use BlueCov for reporting coverage percentages, but only for communication with the test generator. Originally, BlueCov was developed for and successfully used in the Diffblue Ltd Cover commercial test case generator. It has now been open-sourced¹.

This paper is an extended version of [7]. It is structured as follows: Section 2 gives some background on JBMC and coverage analysis,

¹<https://github.com/diffblue/BlueCov>

Section 3 explains the combination of JBMC and BlueCov, Section 4 illustrates the approach with an example, Section 5 gives details on the implementation and Section 6 concludes the paper.

2 BACKGROUND

2.1 JBMC— Java Bounded Model Checking

JBMC [4, 6] is a bounded model checker for Java which uses the compiled class files as input. In this sense it could also analyze other languages that are compiled to Java bytecode, e.g., Kotlin. It can also read jar files which are a form of structured zip archive of class files.

JBMC is a front-end to the CProver framework. CProver is also used by the CBMC model checker for C [10]. JBMC translates the Java bytecode into CProver’s internal GOTO program code. GOTO is a simple sequential programming language which is very close to C in its execution model. It represents the CFG of the program using GOTO instructions to model CFG edges. Java bytecode on the other hand is a stack-based low-level, assembly code like programming language.

Only necessary classes and functions are translated to GOTO. For this, JBMC performs an over-approximate reachability analysis on the call graph. Only functions which are actually used are added to the GOTO program.

JBMC uses the SMT solver built into CProver as backend, which implements a reduction of bitvector, floating-point, array and string theories to SAT. JBMC competes in the Java track at the Software Verification Competition (SV-COMP) [1, 5], has won once and was among the top three tools otherwise.

2.2 Coverage Analysis in JBMC

For test case generation JBMC currently uses the *location* coverage criterion. This criterion tries to cover each bytecode in the class file. For each CFG node corresponding to distinct bytecode instructions, JBMC inserts an assertion node `assert(false)` representing the coverage goal into the CFG. An analysis with JBMC then tries to deduce values for the input parameters of the method in such a way that the assertions are reached. These synthesized input values are then translated into executable test cases. For BlueCov other coverage criteria are supported, too, e.g. branch, path or MC/DC coverage.

The translation from JBMC results to executable Java is not trivial. Primitive values are relatively easy, but complex objects can be difficult to create. This is true in particular when one does not want to use reflection to create object instances. While reflection is very powerful, it generally does not correspond to readable tests a human tester would write. Reflection also allows for creation of object states which could not be created programmatically, e.g., because internal invariants are not respected. The Java code test case generation is proprietary to the company Diffblue Ltd Ltd. and thus not part of the presentation here.

2.3 Testcase Minimization

In many applications JBMC will have to create multiple new test cases in order to complete the test suite. Often there are different possibilities to cover the remaining coverage goals. As each generated unit test might cover several coverage goals, it makes sense to

try to minimize the number of generated unit tests to complete the coverage.

The underlying problem is the subset cover problem which is \mathcal{NP} -complete and can therefore be expensive to solve if the number of tests is large. JBMC uses a greedy heuristic to approximate the minimal number of new unit tests. First, traces are sorted by the number of covered goals. Traces are then processed in this order, collecting newly covered goals and dropping those traces that do not cover any new goal. This is similar to the approach used by [14].

This approximation works well in practice to reduce the number of generated test cases.

2.4 Java Specific Challenges

Java is an object-oriented programming language with an extensive standard library. This poses some additional challenges for model-checking of Java programs as compared to C programs. We give some challenges here and illustrate how JBMC handles these. These challenges can have a big impact on differences between reported and measured coverage because JBMC might have a different model than what is really executed.

Object Orientation Java methods often use the interfaces as input parameters. As interfaces do not have an implementation, an analysis tool has to use a class which implements the interface. This often requires loading more classes than the ones that are actually referenced in a class file. For example, using the *List* interface would create a reference to *List* but would also require loading another class that implements the interface, e.g., *ArrayList*.

Another challenge is the use of generics in the class file. Java uses type erasure which means that there is no typing information in the class files. In such case JBMC analyzes the bytecode instructions for explicit casts without checks. As the compiler does have the necessary information, it can emit instructions that do not check the result after a cast.

Java Class Library An important challenge when analyzing Java is the Java class library (JCL). It provides different data structures and algorithms. It is not feasible to simply load the JCL class files, as the resulting model would be much too big. What JBMC does is to provide a *jar* file that contains *models* of commonly used parts of the JCL. Such models do not implement the respective class functionality directly but contain a sufficient specification for JBMC.

In the running example (cf. listing 12) in this paper, the *core models* library is used to find the model for the *Math.abs* function. In this case the implementation is the same as the original implementation in the JCL as shown in listing 4.

```

1 public static float abs(float a) {
2     return (a <= 0.0F) ? 0.0F - a : a;
3 }

```

Listing 1: *Math.abs* Implementation

For other functions JBMC provides a different implementation which is functionally equivalent to the original one, but is easier for the SMT solver to handle. An example for this is the *max* (maximum) function for *float* parameters. The JBMC core models version of this function is shown in listing 28. The commented-out section contains the original code of the JCL. The JBMC implementation is functionally equivalent according to the IEEE754 standard but

is simpler for JBMC to analyze. For example, it does not use the `Float.floatToRawIntBits` method.

```

1  public static float max(float a, float b) {
2  // original code
3  // if (a != a)
4  //     return a; // a is NaN
5  // if ((a == 0.0f) &&
6  //     (b == 0.0f) &&
7  //     (Float.floatToRawIntBits(a)
8  //     == negativeZeroFloatBits)) {
9  //     // Raw conversion ok since NaN
10 //     // can't map to -0.0.
11 //     return b;
12 // }
13 // return (a >= b) ? a : b;
14
15 // JBMC core-models implementation
16 if (Float.isNaN(a) || Float.isNaN(b)) {
17     return Float.NaN;
18 } else {
19     float result = CProver.nondetFloat();
20     // choose result in such a way that it is
21     // the maximum of a and b
22
23     CProver.assume( (result == a || result == b)
24                   && result >= a && result >= b);
25     return result;
26 }
27 }

```

Listing 2: Core Model version of `Math.max`

Some JCL functions are currently not implemented or modeled. These functions simply return a nondeterministic value of the correct type when called. If this is a problem for an analysis, one can additionally model the required function which increases the precision of the analysis and of the test creation done by JBMC.

3 COMBINING JBMC AND BLUECOV

As explained above, there exist several tools which report coverage of Java programs, but none uses precisely the same criteria as JBMC does. For real-world project, it is an important goal to minimize the effort to complete the test suite. This requires measuring the test coverage in the exact way as JBMC does. Therefore, the BlueCov tool was developed to facilitate integration with JBMC. It tracks the exact Java bytecode instructions that it associates with its coverage criterion. JBMC then provides this information to BlueCov, which instruments the same bytecode instructions for runtime coverage measurement. That way, we can close the gap between the *execution coverage* that can be obtained from test execution and the *generation coverage* that JBMC reports to have achieved during test generation.

The coverage goals determined by JBMC to drive its test generation are reported as JSON output. For each coverage goal JBMC emits a Boolean flag covered or not covered. This is called the *generation coverage* of JBMC. JBMC also calculates values for the input parameters for Java methods in order to reach the coverage goals.

If there is a mismatch in the understanding of the coverage criteria between the test generator and the coverage measurement tool then the generation coverage might be different from the *execution coverage*. The execution coverage is the coverage which is reached when the code is actually executed on the JVM using the

input parameters calculated by JBMC. Figure 1 shows the overall approach.

Class Files to Properties The first step is to give the class files of the project under analysis to JBMC. It provides the option to show the coverage goals (or *properties*) considered for *generation coverage*. The output is given in JSON format and contains all the information necessary for BlueCov to perform the bytecode instrumentation.

Property Instrumentation In the second step the BlueCov tool takes the output of the JBMC coverage goals, the class files of the project and a list of class files to instrument as its input. It then proceeds by creating a database file and instruments the class files with each coverage goal.

Measure Existing Coverage In the third step the existing test suite of the project is run using the instrumented class files. During this run, BlueCov registers when an instrumented coverage goal is reached and increases a counter in the database file for each reached coverage goal.

Coverage Report After the run of the test suite, BlueCov reports the measured *execution coverage*. The report is in JSON format and contains the hitcount for each coverage goal.

Enhance Test Coverage The BlueCov coverage report is then used as input for JBMC create a minimal number of additional tests to complete the coverage of the existing test suite. JBMC reports the *generation coverage* that it achieved together with the generated test inputs to the methods.

4 EXAMPLE USING BLUECOV

We use the example in listing 12 to illustrate the steps of our approach based on BlueCov. The below function `FloatTools.sign` returns an *int* depending on the input *float* parameter. The intent is to have a sign function of the input parameter *x*. The first case in line 3 covers the “zero” value, i.e., a very small absolute value of *x* and returns 0. The second case (line 5) returns -1 in case *x* is negative and the third case (line 7) returns 1 in case *x* is positive. The return statement in line 9 was added because the compiler would otherwise emit the error *error: missing return statement*.

```

1  public class FloatTools {
2  public static int sign(float x) {
3      if(Math.abs(x) < 1e-6)
4          return 0;
5      if(x < 0)
6          return -1;
7      if(x > 0)
8          return 1;
9      return -2;
10 }
11 }

```

Listing 3: Example Program

The existing test suite is shown in listing 15, i.e., testing each of the three cases:

```

1  public class FloatToolsTest {
2  @Test
3  public void testSignZero () {
4      assertEquals(0, FloatTools.sign(-1e-10f));
5  }
6  @Test
7  public void testSignNeg () {
8      assertEquals(-1, FloatTools.sign(-10f));
9  }

```

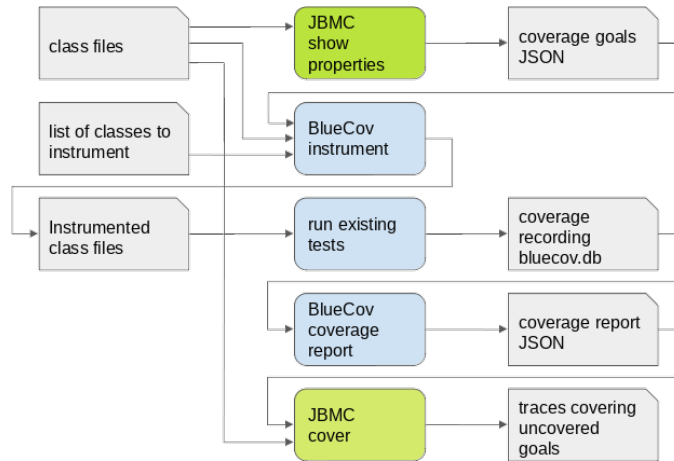


Figure 1: Overview of the Approach (JBMBC – green, BlueCov – blue)

```

10  @Test
11  public void testSignPos () {
12      assertEquals(1, FloatTools.sign(1234f));
13  }
14  }

```

Listing 4: Test Suite for Example Program

4.1 Class files to Properties

JBMBC converts the bytecode instructions of *FloatTools.sign* into its internal GOTO representation and inserts coverage goals according to the given coverage criterion. For *location* coverage, JBMBC creates 9 location coverage goals for *FloatTools.sign* corresponding to the ASSERT statements in the GOTO function below.²

```

FloatTools.sign(float) /* FloatTools.sign:(F)I */
DECL return_tmp0 : floatbv[32]
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 1
ASSERT false // goal FloatTools.sign:(F)I.coverage.1
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 1
CALL java.lang.Math.<clinit >()
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 1
CALL java.lang.Math.abs:(F)F(x)
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 1
ASSIGN return_tmp0 := java.lang.Math.abs:(F)F::return_value
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 1
DEAD java.lang.Math.abs:(F)F::return_value
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 5
ASSERT false // goal FloatTools.sign:(F)I.coverage.2
// FloatTools.java line 5 FloatTools.sign:(F)I bytecode-index 5
IF (NOT isnan(return_tmp0)) AND
    (return_tmp0 < 1.000000e-6) THEN GOTO 1
// FloatTools.java line 7 FloatTools.sign:(F)I bytecode-index 8
ASSERT false // goal FloatTools.sign:(F)I.coverage.3
// FloatTools.java line 7 FloatTools.sign:(F)I bytecode-index 8
DEAD return_tmp0
// FloatTools.java line 7 FloatTools.sign:(F)I bytecode-index 8
GOTO 2
1: DEAD return_tmp0
// FloatTools.java line 6 FloatTools.sign:(F)I bytecode-index 7
ASSERT false // goal FloatTools.sign:(F)I.coverage.4
// FloatTools.java line 6 FloatTools.sign:(F)I bytecode-index 7
ASSIGN FloatTools.sign:(F)I::return_value := 0
// FloatTools.java line 6 FloatTools.sign:(F)I bytecode-index 7
GOTO 5

```

```

// FloatTools.java line 7 FloatTools.sign:(F)I bytecode-index 11
2: ASSERT false // goal FloatTools.sign:(F)I.coverage.5
// FloatTools.java line 7 FloatTools.sign:(F)I bytecode-index 11
IF isnan(x) OR (x >= 0) THEN GOTO 3
// FloatTools.java line 8 FloatTools.sign:(F)I bytecode-index 13
ASSERT false // goal FloatTools.sign:(F)I.coverage.6
// FloatTools.java line 8 FloatTools.sign:(F)I bytecode-index 13
ASSIGN FloatTools.sign:(F)I::return_value := -1
// FloatTools.java line 8 FloatTools.sign:(F)I bytecode-index 13
GOTO 5
// FloatTools.java line 9 FloatTools.sign:(F)I bytecode-index 17
3: ASSERT false // goal FloatTools.sign:(F)I.coverage.7
// FloatTools.java line 9 FloatTools.sign:(F)I bytecode-index 17
IF isnan(x) OR (x <= 0) THEN GOTO 4
// FloatTools.java line 10 FloatTools.sign:(F)I bytecode-index 19
ASSERT false // goal FloatTools.sign:(F)I.coverage.8
// FloatTools.java line 10 FloatTools.sign:(F)I bytecode-index 19
ASSIGN FloatTools.sign:(F)I::return_value := 1
// FloatTools.java line 10 FloatTools.sign:(F)I bytecode-index 19
GOTO 5
// FloatTools.java line 11 FloatTools.sign:(F)I bytecode-index 21
4: ASSERT false // goal FloatTools.sign:(F)I.coverage.9
// FloatTools.java line 11 FloatTools.sign:(F)I bytecode-index 21
ASSIGN FloatTools.sign:(F)I::return_value := -2
5: END_FUNCTION

```

As Java bytecode is used for the analysis, there can be multiple coverage goals per source code line. The coverage goals map to the original source code as shown in listing 12.³

```

1  public class FloatTools {
2      public static int sign(float x) {
3          if(Math.abs(x) < 1e-9) // goal 1, goal 2
4              return 0; // goal 4
5          if(x < 0) // goal 3, goal 5
6              return -1; // goal 6
7          if(x > 0) // goal 7
8              return 1; // goal 8
9          return -2; // goal 9
10     }
11 }

```

Listing 5: Mapping of Goals to Source Code

JBMBC outputs information about the coverage goals in JSON format. For example, for goal `FloatTools.sign:(F)I.coverage.1` it produces an entry as shown below. This also includes information

²Simplified to improve readability; the full output can be produced with the artifact available at <https://www.dropbox.com/s/a7pe39ygj5znuh4/bluecov.zip>

³For readability, we write `goal 1` for the goal with name `FloatTools.sign:(F)I.coverage.1`.

to map the bytecode information to the line information of the source Java program.

```

1 {
2   "class": "coverage",
3   "coveredLines": "5",
4   "description": "block_2_(lines_FloatTools.java:5)",
5   "expression": "false",
6   "name": "FloatTools.sign:(F)I.coverage.1",
7   "sourceLocation": {
8     "bytecodeIndex": "1",
9     "file": "FloatTools.java",
10    "function": "FloatTools.sign:(F)I",
11    "line": "5"
12  }
13 }
```

Analogous entries exist for the other eight coverage goals. Each has a unique name and a source location which it covers. For Java, this source location contains the bytecode index of the goal, the class name, the method name and the parameter and return value types.⁴

4.2 Property Instrumentation

Using the information about the goals as generated by JBMC, BlueCov instruments the class files. It either creates an empty coverage database file or adds new entries to an existing one. The default location and name is *blueCov.db*. For each of the generated goals there is a unique identifier used to identify the goal in the database file. This UID is constructed from the fully qualified name, function parameter types and bytecode index, e.g., `FloatTools.sign:(F)I@1` for UID 0.

The bytecode instructions of the *sign* method first call the *Math.abs* function which takes a *float* parameter and returns a value of type *float*. This is shown below, first the original code then the instrumented code.

```

1 public static int sign(float); // original code
2   Code:
3     0: fload_0
4     1: invokestatic #2
5     // Method java/lang/Math.abs:(F)F
6     ...
7
8 public static int sign(float); // instrumented code
9   Code:
10    0: fload_0
11    1: getstatic #17
12    // Field company_coverage_reporter:
13    //   Lorg/cprover/coverage/CoverageLog;
14    4: ldc #18 // int 0
15    6: invokevirtual #24
16    // Method org/cprover/coverage/CoverageLog.record:(I)V
17    9: invokestatic #30
18    // Method java/lang/Math.abs:(F)F
```

In the instrumented code first the static instance of *CoverageLog* is loaded, then its *record* method is called with the UID 0. This identifier corresponds to the above goal. This call increments the hitcount for this goal in the BlueCov in-memory database.

The class itself does not contain the *clinit* static initializer as it does not contain any static fields. A static field is needed by

BlueCov and therefore it adds a static initializer to get the singleton *org.cprover.coverage.CoverageLog* object used for logging. The generated bytecode instructions of the static initializer looks as follows:

```

1 static {};
2   Code:
3     0: invokestatic #43
4     // Method org/cprover/.../
5     //   CoverageLog.getInstance:()Lorg/.../
6     //   CoverageLog;
7     3: putstatic #15
8     // Field company_coverage_reporter:
9     //   Lorg/cprover/coverage/CoverageLog;
10    6: return
```

Each of the other coverage goals is instrumented analogously. Each instrumentation uses three additional bytecode instructions, each with a different UID for the database. The UID is loaded as a parameter using the *ldc* or load constant bytecode instruction.⁵

This code instrumentation causes a constant time overhead when executing such instrumented bytecode. This could be a problem in really long-running unit tests, in particular within loops. If the exact number of hits is not required, then a further optimization is possible which records only the first time a specific location is covered and skips logging afterwards. This effectively reduces the runtime overhead to a branch execution that can always be predicted correctly after the first time.

4.3 Measure Existing Coverage

To measure the existing coverage, the *record* method of the *CoverageLog* class is called each time a coverage goal is reached. When called the first time for a specific coverage goal, the method inserts the key with value 1 into the hash map. At each further call, the associated value is incremented, updating the hitcount entries in the database. For this, BlueCov needs to be added to the Java classpath.⁵

4.4 Coverage Report

After the test suite is finished, the coverage reporter of BlueCov is called which reports the measured coverage. The reporter iterates over all entries in the database and emits the measured hitcount for each coverage goal.⁵

BlueCov emits a *coverage report* in JSON format which shows the information of Table 1. The results for our example are as expected: the condition in line 3 is executed for each unit test and therefore the associated two goals are hit 3 times each. The condition in line 5 and the associated two goals are reached 2 times and the last test in line 8 is reached one time only. Each return statement in line 4, 6 and 8 is reached exactly once. And finally the return statement in line 9 is not reached at all.

4.5 Enhance Test Coverage

At a first glance, it seems that line 9 is simply dead code which cannot be covered, but, when telling JBMC to cover the coverage

⁴The full output can be produced using the artifact available at <https://www.dropbox.com/s/a7pe39ygj5znuh4/bluecov.zip>

⁵For our example, the instrumentation, coverage measurement and coverage report can be reproduced with the artifact at <https://www.dropbox.com/s/a7pe39ygj5znuh4/bluecov.zip>

goal name	hit count
FloatTools.sign:(F)I.coverage.1	3
FloatTools.sign:(F)I.coverage.2	3
FloatTools.sign:(F)I.coverage.3	2
FloatTools.sign:(F)I.coverage.4	1
FloatTools.sign:(F)I.coverage.5	2
FloatTools.sign:(F)I.coverage.6	1
FloatTools.sign:(F)I.coverage.7	1
FloatTools.sign:(F)I.coverage.8	1
FloatTools.sign:(F)I.coverage.9	0

Table 1: Coverage Report for *FloatTools.sign*

goal 9, JVMC does in fact report that it can complete the test-suite because it found a way to execute line 9.

The input value for this is an IEEE-754 “not a number” (*NaN*) value. This is a valid value for the input parameter *x* and *NaN* values have the property that each comparison with such a value evaluates to *false*. Therefore, we can complete the test suite by adding the unit test shown in listing 5.

```

1  @Test
2  public void testSignNaN () {
3      assertEquals(-2, FloatTools.sign(Float.NaN));
4  }
```

Listing 6: Additionally Generated Unit Test

After adding this additional test and re-executing the full test-suite, BlueCov reports coverage of all location coverage goals of *FloatTools.sign*.

5 BYTECODE INSTRUMENTATION AND COVERAGE MEASUREMENT

To perform the bytecode instrumentation in BlueCov we chose the ASM⁶ library. ASM is widely used to instrument Java bytecode, e.g., in JaCoCo⁷. It uses the visitor pattern for instrumentation.

The challenge here is being able to instrument the bytecode at the right positions and to reliably get the information at runtime which instrumented bytecode is executed. In particular, it is necessary to make sure that even non-standard program termination like uncaught exceptions or direct calls to exit via JNI do not prevent the information from being stored persistently so that it can be reported back to JVMC for test generation.

5.1 Bytecode Instrumentation

Each bytecode instruction is visited and can be changed. It is interesting to note that ASM does not provide the information about the byte offset or address of an instruction. As instructions exist which can have different sizes, e.g., ‘*iload 0*’ and its specialization ‘*iload_0*’. Instead, we identify the bytecode instructions by their bytecode index. The bytecode index is the sequence number of bytecode instructions and is independent of the size in bytes. JVMC provides the bytecode index for each coverage goal.

⁶<https://asm.ow2.io/>

⁷<https://www.eclemma.org/jacoco/>

Each bytecode instruction specified for instrumentation by the coverage criterion of JVMC is instrumented by inserting the following code directly before the bytecode instruction where *uid* is the unique identifier of the instrumented location.

```
1  companyCoverageReporter.record(uid);
```

This Java code is translated into bytecode by pushing the static field *companyCoverageReporter* on the stack, then the *uid* of the instrumented location and finally calling the *record* method of the *org.cprover.coverage.CoverageLog* class.

This instrumentation is performed as shown in listing 28. Here we show the visiting method for jump instructions. Analogous methods exist for each of the other categories of bytecode instructions as well. It calls directly *instrumentByteCode* method with the current bytecode index as argument which checks whether the given bytecode index should be instrumented. If yes, the necessary bytecode instructions and arguments are inserted into the code, if no, nothing is done. On return to *visitJumpInsn*, the bytecode index counter is incremented and the next bytecode instruction is visited. In this way, all bytecode instructions are visited and each one that corresponds to a coverage goal gets instrumented with the necessary code.

```

1  public final void visitJumpInsn(final int opcode,
2                                final Label label) {
3      instrumentByteCode(bcLine);
4      bcLine += 1;
5      super.visitJumpInsn(opcode, label);
6  }
7
8  final void instrumentByteCode(final int bcLine) {
9      if (shouldBeInstrumented(bcLine)) {
10         lastMethodWasInstrumented = true;
11         // get instance from static field
12         // push value to record
13         // call `record` on CoverageLog
14         super.visitFieldInsn(Opcodes.GETSTATIC,
15                             this.className,
16                             "companyCoverageReporter",
17                             "Lorg/cprover/coverage/CoverageLog;");
18         super.visitLdcInsn(getUniqueIdentifier(bcLine));
19         super.visitMethodInsn(Opcodes.INVOKEVIRTUAL,
20                             "org/cprover/coverage/CoverageLog",
21                             "record",
22                             "(I)V",
23                             false);
24         debug("added_ID_" + getUniqueIdentifier(bcLine));
25         instrumentedLocs.add(getUniqueIdentifier(bcLine));
26     }
27 }
```

Listing 7: Bytecode Instrumentation Implementation

5.2 BlueCov Runtime Coverage Measurement

At runtime, whenever an instrumented bytecode instruction is executed, there is a call to BlueCov to register the coverage of the associated location. This increments the hit count of the corresponding location. This happens in the *CoverageLog* class of the *org.cprover.coverage* package. The *record* method is shown in listing 9.

```

1  public void record(final int key) {
2      Integer i = inMemoryMap.get(key);
3      if (i == null) {
```

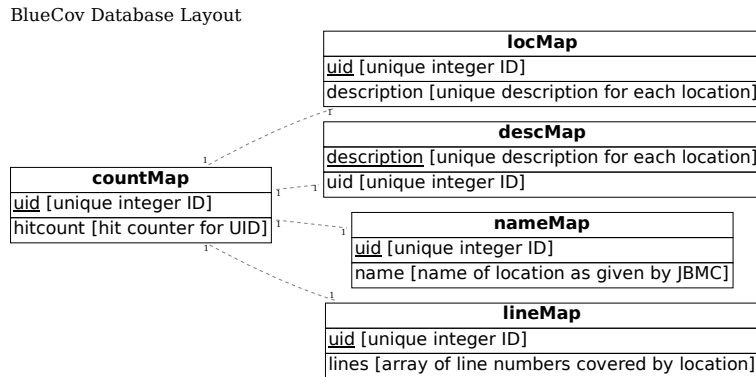


Figure 2: MapDB HTree objects

```

4     inMemoryMap.put(key, 1);
5   } else {
6     inMemoryMap.put(key, i + 1);
7   }
8 }

```

Listing 8: Recording Reaching a Coverage Goal

Storing this information persistently is realized using the *mapDB* library⁸. This library provides a combination of Java collections and on-disk database storage, i.e., there is a thread-safe hash map which is in-memory and also mirrored on disk. The ER diagram of the relevant part of the database is shown in 2.

5.2.1 Challenges. Using exact coverage measurement provides some challenges which had to be addressed. For each bytecode instruction which is instrumented for the coverage analysis, there exists one call to get the *CoverageLog* singleton object and one call to record execution of the corresponding bytecode instruction. The required time to complete this has to be minimized in order to keep execution time of the project under analysis reasonable.

To achieve this, BlueCov uses an in-memory database in its normal execution mode and only writes the in-memory content to the *mapDB* on-disk database on exit. This greatly reduces the time required to execute the projects under test.

The drawback here is that it has to be guaranteed that the in-memory database is correctly written at each possible program termination. This is achieved by installing a *shutdown hook* for the JVM runtime using the *addShutdownHook* method of the *Runtime* class. This method takes a *Thread* as parameter, its *run* method get called on normal exit but also if the JVM terminates due to a user action or system-wide event⁹. When the hook is called the code in listing 19 is executed and writes the in-memory content to the on-disk database on exit.

```

1   public void run() {
2     if (inMemory) {
3       db = makeDb();
4       countMap = db.hashMap(locCountMap)
5         .keySerializer(Serializer.INTEGER)

```

⁸<https://mapdb.org/>

⁹[https://docs.oracle.com/javase/8/docs/api/java/lang/Runtime.html#addShutdownHook\(java.lang.Thread\)](https://docs.oracle.com/javase/8/docs/api/java/lang/Runtime.html#addShutdownHook(java.lang.Thread))

```

6     .valueSerializer(Serializer.INTEGER)
7     .createOrOpen();
8     for (Integer key : inMemoryMap.keySet()) {
9       Integer orig = countMap.get(key);
10      if (orig == null) {
11        orig = 0;
12      }
13      countMap.put(
14        key, orig + inMemoryMap.get(key));
15    }
16  }
17  db.close();
18 }

```

Listing 9: Write In-Memory DB to Disk

5.3 Java Project Integration

BlueCov can easily be integrated into Maven projects. This requires adding the necessary runtime libraries to the dependencies section of the *pom.xml* file. The *bluecov.jar* file location is specified in the plugins section of the *pom.xml* file. In this way, the coverage analysis is used when Maven runs the test suite.

6 CONCLUSION

We presented an approach to focussing a Java test generator based on bounded model checking on the uncovered test goals of an existing test suite. BlueCov achieves this by performing a tailored bytecode instrumentation that captures exactly the coverage goals that correspond to a chosen coverage metric provided by the test generation tool.

This avoids mismatches in the coverage criteria implemented in existing coverage measurement tools and thus avoids the generation of redundant tests. This direct integration of formal analysis and test coverage analysis allows for reduction of cost of automated test case generation and an increase in precision. The minimization of newly generated tests increases clarity of the test suite which is important for human testers interacting with the test suite.

As open-source software, BlueCov is freely available and can be integrated into Java projects. Together with JBMC this can help in analyzing and completing test suite coverage for projects.

Related work. The CProver framework has been used to implement several test generation tools based on bounded model checking, mainly for C programs, the most versatile one being FShell [8], which offers a domain specific language for expressing arbitrary coverage goals. The tool ChainCover [13, 14] aims at reducing the initialization overhead of testing reactive systems by merging tests into test scenarios.

Semi-automated approaches to closing test coverage gaps [11] have been considered as well as automated ones, e.g. [3], but none if these works considers the problem of filling coverage gaps by exactly integrating with runtime coverage measurement in order to avoid the generation of tests that overlap with existing automatically or manually created test suites.

Future work. The proposed approach can be also taken to focus model checking on untested code, similarly to conditional model checking [2] except that the conditions are not based on information from another verification engine, but a coverage analysis tool.

REFERENCES

- [1] Beyer, D.: Automatic verification of C and Java programs: SV-COMP 2019. In: Tools and Algorithms for the Construction and Analysis of Systems. Lecture Notes in Computer Science, vol. 11429, pp. 133–155. Springer (2019). https://doi.org/10.1007/978-3-030-17502-3_9
- [2] Beyer, D., Henzinger, T.A., Keremoglu, M.E., Wendler, P.: Conditional model checking: a technique to pass information between verifiers. In: 20th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-20), SIGSOFT/FSE'12, Cary, NC, USA - November 11 - 16, 2012. p. 57. ACM (2012). <https://doi.org/10.1145/2393596.2393664>
- [3] Bloem, R., Knighofer, R., Rck, F., Tautschnig, M.: Automating test-suite augmentation. In: 2014 14th International Conference on Quality Software, Allen, TX, USA, October 2-3, 2014. pp. 67–72. IEEE (2014). <https://doi.org/10.1109/QSIC.2014.40>
- [4] Cordeiro, L., Kesseli, P., Kroening, D., Schrammel, P., Trtik, M.: JBMC: A bounded model checking tool for verifying Java bytecode. In: International Conference on Computer Aided Verification. pp. 183–190. Springer (2018)
- [5] Cordeiro, L.C., Kroening, D., Schrammel, P.: Benchmarking of Java verification tools at the Software Verification Competition (SV-COMP). ACM SIGSOFT Softw. Eng. Notes 43(4), 56 (2018). <http://arxiv.org/abs/1809.03739>
- [6] Cordeiro, L.C., Kroening, D., Schrammel, P.: JBMC: bounded model checking for Java bytecode - (competition contribution). In: Tools and Algorithms for the Construction and Analysis of Systems. Lecture Notes in Computer Science, vol. 11429, pp. 219–223. Springer (2019). https://doi.org/10.1007/978-3-030-17502-3_17
- [7] Gdemann, M., Schrammel, P.: BlueCov: Integrating Test Coverage and Model Checking with JBMC. In: Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing (SAC). ACM (2023). <https://doi.org/10.1145/3555776.3577829>
- [8] Holzer, A., Schallhart, C., Tautschnig, M., Veith, H.: Query-driven program testing. In: Verification, Model Checking, and Abstract Interpretation, 10th International Conference, VMCAI 2009, Savannah, GA, USA, January 18-20, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5403, pp. 151–166. Springer (2009). https://doi.org/10.1007/978-3-540-93900-9_15
- [9] JaCoCo: Java code coverage, <https://www.jacoco.org/jacoco/>
- [10] Kroening, D., Tautschnig, M.: CBMC – C bounded model checker. In: International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 389–391. Springer (2014)
- [11] Nellis, A., Kesseli, P., Conmy, P.R., Kroening, D., Schrammel, P., Tautschnig, M.: Assisted coverage closure. In: NASA Formal Methods - 8th International Symposium, NFM 2016, Minneapolis, MN, USA, June 7-9, 2016. Proceedings. Lecture Notes in Computer Science, vol. 9690, pp. 49–64. Springer (2016). https://doi.org/10.1007/978-3-319-40648-0_5
- [12] Schrammel, P.: How testable is business software? In: 2020 Formal Methods in Computer Aided Design, FMCAD 2020, Haifa, Israel, September 21-24, 2020. p. 1. IEEE (2020), <https://arxiv.org/abs/2011.00630>
- [13] Schrammel, P., Melham, T., Kroening, D.: Chaining test cases for reactive system testing. In: Testing Software and Systems - 25th IFIP WG 6.1 International Conference, ICTSS 2013, Istanbul, Turkey, November 13-15, 2013. Proceedings. Lecture Notes in Computer Science, vol. 8254, pp. 133–148. Springer (2013). https://doi.org/10.1007/978-3-642-41707-8_9
- [14] Schrammel, P., Melham, T., Kroening, D.: Generating test case chains for reactive systems. Int. J. Softw. Tools Technol. Transf. 18(3), 319–334 (2016). <https://doi.org/10.1007/s10009-014-0358-6>