

Reasoning About Exceptional Behavior At the Level of Java Bytecode^{*}

Marco Paganoni¹ and Carlo A. Furia¹[0000–0003–1040–3201]

Software Institute, USI Università della Svizzera italiana, Lugano, Switzerland
marco.paganoni@usi.ch bugcounting.net

Abstract. A program’s exceptional behavior can substantially complicate its control flow, and hence accurately reasoning about the program’s correctness. On the other hand, formally verifying realistic programs is likely to involve exceptions—a ubiquitous feature in modern programming languages.

In this paper, we present a novel approach to verify the exceptional behavior of Java programs, which extends our previous work on BYTEBACK. BYTEBACK works on a program’s bytecode, while providing means to specify the intended behavior at the source-code level; this approach sets BYTEBACK apart from most state-of-the-art verifiers that target source code. To explicitly model a program’s exceptional behavior in a way that is amenable to formal reasoning, we introduce Vimp: a high-level bytecode representation that extends the Soot framework’s Grimp with verification-oriented features, thus serving as an intermediate layer between bytecode and the Boogie intermediate verification language. Working on bytecode through this intermediate layer brings flexibility and adaptability to new language versions and variants: as our experiments demonstrate, BYTEBACK can verify programs involving exceptional behavior in all versions of Java, as well as in Scala and Kotlin (two other popular JVM languages).

1 Introduction

Nearly every modern programming language supports exceptions as a mechanism to signal and handle unusual runtime conditions (so-called *exceptional behavior*) separately from the main control flow (the program’s *normal* behavior). Exceptions are usually preferable to lower-level ad hoc solutions (such as error codes and defensive programming), because deploying them does not pollute the source code’s structured control flow. However, by introducing extra, often implicit execution paths, exceptions may also complicate reasoning about all possible program behavior—and thus, ultimately, about program correctness.

In this paper, we introduce a novel approach to perform deductive verification of Java programs involving exceptional behavior. Exceptions were baked into the Java programming language since its inception, where they remain widely used [13, 15, 17]; nevertheless, as the example of Sec. 2 demonstrates, they can be somewhat of a challenge to reason about formally. To model together normal and exceptional control flow

^{*} Work partially supported by SNF grant 200021-207919 (LastMile).

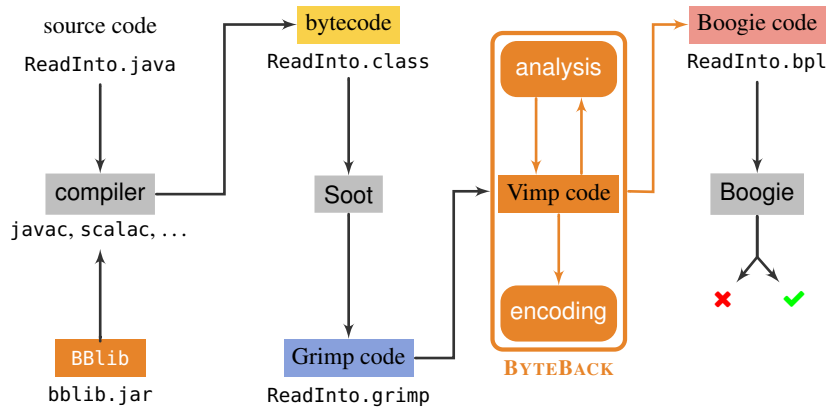


Fig. 1: An overview of BYTEBACK’s verification workflow.

paths, and to seamlessly support any exception-related language features, our verification approach crucially targets a program’s *bytecode* intermediate representation—instead of the source code analyzed by state-of-the-art verifiers such as KeY [1] and OpenJML [6]. We introduced the idea of performing formal verification at the level of bytecode in previous work [20]. In this paper, we build on those results and implement support for exceptions in the BYTEBACK deductive verifier.

The key idea of our BYTEBACK approach (pictured in Fig. 1) is using JVM bytecode solely as a convenient intermediate representation; users of BYTEBACK still annotate program source code in a very similar way as if they were working with a source-level verifier. To this end, we extend the specification library introduced with BYTEBACK (called BLib) with features to specify exceptional behavior (for example, conditions under which a method terminates normally or exceptionally) using custom Java expressions; thus, such specifications remain available in bytecode after compiling an annotated program using a standard Java compiler. BYTEBACK analyzes the bytecode and encodes the program’s semantics, its specification, and other information necessary for verification into Boogie [3]—a widely-used intermediate language for verification; then, verifying the Boogie translation is equivalent to verifying the original Java program against its specification.

As we demonstrate with experiments in Sec. 4, performing verification on bytecode offers several advantages: *i) Robustness* to source-language changes: while Java evolves rapidly, frequently introducing new features (also for exceptional behavior), bytecode is generally stable; thus, our verification technique continues to work with the latest Java versions. *ii) Multi-language* support: BYTEBACK and its BLib specification library are designed so that they can be applied, in principle, to specify programs in any language that is bytecode-compatible; while the bulk of our examples are in Java, we will demonstrate verifying exceptional behavior in Scala and Kotlin—two modern languages for the JVM. *iii) Flexibility* of modeling: since exceptional behavior becomes explicit in bytecode, the BYTEBACK approach extensively and seamlessly deals with any intricate exceptional behavior (such as implicit or suppressed exceptions).

Contributions and positioning. In summary, the paper makes the following contributions: *i)* Specification features to specify exceptional behavior of JVM languages. *ii)* A

<pre> 1 @Require(r = null ∨ ¬r.closed) 2 @Raise(NullPointerException, r = null) 3 @Return(a ≠ null ∧ a.length = 0 ∧ r ≠ null) 4 @Ensure(r = null ∨ r.closed) 5 static void into(final Resource r, final int[] a) { 6 try (r) { 7 int i = 0; 8 while (true) { invariant(0 ≤ i ≤ a.length); 9 invariant(r = null ∨ ¬r.closed); 10 a[i] = r.read(); ++i; 11 } 12 } catch (IndexOutOfBoundsException 13 NoSuchElementException e) { return; } 14 } </pre>	<pre> class Resource implements AutoCloseable { boolean closed; boolean hasNext; @Raise(IllegalStateException, closed) @Raise(NoSuchElementException, ¬hasNext) @Return(¬closed ∧ hasNext) int read() { /* ... */ } // ... } </pre>
---	--

(a) Method `into` copies `r`'s content into array `a`. It is annotated with normal and exceptional pre- and postconditions using a simplified BBLib syntax.

(b) An outline of class `Resource`'s interface with Boolean attributes `closed` and `hasNext`, and method `read`.

Fig. 2: Annotated Java method `into` and class `Resource`, which demonstrate some pitfalls of specifying and reasoning about exceptional behavior.

verification technique that encodes bytecode exceptional behavior into Boogie. *iii*) An implementation of the specification features and the verification technique that extend the BBLib library and BYTEBACK verifier. *iv*) Vimpr: a high-level bytecode format suitable to reason about functional correctness, built on top of Soot's Grimp format [11,26]. *v*) An experimental evaluation with 37 programs involving exceptional behavior in Java, Scala, and Kotlin. *vi*) For reproducibility, BYTEBACK and all experimental artifacts are available in a replication package [19]. While we build upon BBLib and BYTEBACK, introduced in previous work of ours [20], this paper's contributions substantially extend them with support for exceptional behavior, as well as other Java related features (see Sec. 3). For simplicity, henceforth “BBLib” and “BYTEBACK” denote their current versions, equipped with the novel contributions described in the rest of the paper.

2 Motivating Example

Exceptions can significantly complicate the control flow of even seemingly simple code; consequently, correctly reasoning about exceptional behavior can be challenging even for a language like Java—whose exception-handling features have not changed significantly since the language's origins.

To demonstrate, consider Fig. 2a's method `into`, which inputs a reference `r` to a `Resource` object and an integer array `a`, and copies values from `r` into `a`—until either `a` is filled up or there are no more values in `r` to read. Fig. 2b shows the key features of class `Resource`—which implements Java's `AutoCloseable` interface, and hence can be used similarly to most standard I/O classes. Method `into`'s implementation uses a **try**-with-resources block to ensure that `r` is closed whenever the method terminates—normally or exceptionally. The **while** loop terminates as soon as any of the following conditions holds: *i*) `r` is `null`; *ii*) `k` reaches `a.length` (array `a` is full); *iii*) `r.read()` throws a `NoSuchElementException` (`r` has no more elements). Method `into` returns with an (propagated) exception only in case *i*); case *ii*'s `IndexOutOfBoundsException` and case *iii*'s `NoSuchElement` exceptions are caught by the **catch** block that *suppresses* them with a **return**—so that `into` can return normally.

Fig. 2a’s annotations, which use a simplified syntax for BBLib—BYTEBACK’s specification library—specify part of `into`’s expected behavior. Precondition `@Require` expresses the constraint that object `r` must be open (or `null`) for `into` to work as intended. Annotation `@Raise` declares that `into` terminates with a `NullPointerException` exception if `r` is `null`; conversely, `@Return` says that `into` returns *normally* if `a` is a non-`null` empty array and `r` is also not `null`. Finally, postcondition `@Ensure` says that, if it’s not `null`, `r` will be closed when `into` terminates—regardless of whether it does so normally or exceptionally. The combination of language features and numerous forking control-flow paths complicate reasoning about—and even specifying—`into`’s behavior. While exception handling has been part of Java since version 1, try-with-resources and multi-catch (both used in Fig. 2a) were introduced in Java 7; thus, even a state-of-the-art verifier like KeY [1] lacks support for them. OpenJML [6] can reason about all exception features up to Java 7; however, the try-with-resources using an existing `final` variable became available only in Java 9.⁽¹⁾ Furthermore, OpenJML implicitly checks that verified code does not throw any implicit exceptions (such as `NullPointerException` or `IndexOutOfBoundsException` exceptions)—thus disallowing code such as Fig. 2a’s, where implicitly throwing exceptions is part of the expected behavior. These observations are best thought of as design choices—rather than limitations—of these powerful Java verification tools: after all, features such as multi-catch are syntactic sugar that makes programs more concise but does not affect expressiveness; and propagating uncaught implicit exceptions can be considered an anti-pattern⁽²⁾ [14, 27]. However, they also speak volumes to the difficulty of fully supporting all cases of exceptional behavior in a feature-laden language like Java.

As we detail in the rest of the paper, BYTEBACK’s approach offers advantages in such scenarios. Crucially, the implicit control flow of exception-handling code becomes explicit when compiled to bytecode, which eases analyzing it consistently and disentangling the various specification elements. For instance, the `while` loop’s several exceptional exit points become apparent in `into`’s bytecode translation, and BYTEBACK can check that the declared invariant holds in all of them, and that postcondition `@Ensure` holds in all matching method return points. Furthermore, bytecode is more stable than Java—thus, a verifier like BYTEBACK is more robust to source language evolution. Thanks to these capabilities, BYTEBACK can verify the behavior of Fig. 2’s example.

3 Specifying and Verifying Exceptional Behavior

This section describes the new features of BYTEBACK to specify and verify exceptional behavior. Fig. 1 shows BYTEBACK’s workflow, which we revisited to support these new features. Users of BYTEBACK—just like with every deductive verifier—have to annotate the source code to be verified with a specification and other annotations. To this end, BYTEBACK offers BBLib: an annotation library that is usable with any language that is bytecode compatible. Sec. 3.1 describes the new BBLib features to specify exceptional behavior. Then, users compile the annotated source code with the language’s bytecode compiler. BYTEBACK relies on the Soot static analysis framework to analyze bytecode; precisely, Soot offers Grimp: a higher-level alternative bytecode representation. BYTEBACK processes Grimp and translates it to Vimpr: a verification-oriented

extension of Grimp that we introduced in the latest BYTEBACK version and is described in Sec. 3.2. As we discuss in Sec. 3.3, BYTEBACK transforms Grimp to Vimp in steps, each taking care of a different aspect (expressions, types, control flow, and so on). Once the transformation is complete, BYTEBACK encodes the Vimp program into the Boogie intermediate verification language [3]; thanks to Vimp’s custom design, the Boogie encoding is mostly straightforward. Finally, the Boogie tool verifies the generated Boogie program, and reports success or any verification failures (which can be manually traced back to source-code specifications that could not be verified).

3.1 Specifying Exceptional Behavior

Users of BYTEBACK add behavioral specifications to a program’s source using BBLib: BYTEBACK’s standalone Java library, offering annotation tags and static methods suitable to specify functional behavior. Since BBLib uses only basic language constructs, it is compatible with most JVM languages (as we demonstrate in Sec. 4); and all the information added as BBLib annotations is preserved at the bytecode level. This section first summarizes the core characteristics of BBLib (to make the paper self contained), and then describes the features we introduced to specify exceptional behavior.

Specification Expressions. Expressions used in BYTEBACK specifications must be *aggregable*, that is pure (they can be evaluated without side effects) and branchless (they can be evaluated without branching instructions). These are common requirements to ensure that specification expressions are well-formed [24]—hence, equivalently expressible as purely *logic expressions*. Correspondingly, BBLib forbids impure expressions, and offers aggregable replacements for the many Java operators that introduce branches in the bytecode, such as the standard Boolean operators (&&, ||, ...) and comparison operators (==, <, ...). Tab. 1 shows several of BBLib’s aggregable operators, including some that have no immediately equivalent Java expression (such as the quantifiers). Using only BBLib’s operators in a specification ensures that it remains in a form that BYTEBACK can process after the source program has been compiled to bytecode [20]. Thus, BBLib operators map to Vimp logic operators (Sec. 3.2), which directly translate to Boogie operators with matching semantics (Sec. 3.4).

	IN JAVA/LOGIC	IN BBLib
comparison	$x < y, x \leq y, x == y$ $x != y, x \geq y, x > y$	<code>lt(x, y), lte(x, y), eq(x, y)</code> <code>neq(x, y), gte(x, y), gt(x, y)</code>
conditionals	$c ? t : e$	<code>conditional(c, t, e)</code>
propositional	$!a, a \ \&\& \ b, a \ \ b, a \ \implies \ b$	<code>not(a), a & b, a b, implies(a, b)</code>
quantifiers	$\forall x: T \bullet P(x)$ $\exists x: T \bullet P(x)$	<code>T x = Binding.T(); forall(x, P(x))</code> <code>T x = Binding.T(); exists(x, P(x))</code>

Table 1: BBLib’s aggregable operators, used instead of Java’s impure or branching operators in specification expressions.

Method Specifications. To specify the input/output behavior of methods, BBLib offers annotations `@Require` and `@Ensure` to express a method’s pre- and postconditions. For example, `@Require(p) @Ensure(q) t m(args)` specifies that p and q are method

m’s pre- and postcondition; both p and q denote names of *predicate* methods, which are methods marked with annotation `@Predicate`. As part of the verification process, BYTEBACK checks that every such predicate method p is well-formed: *i*) p returns a **boolean**; *ii*) p ’s signature is the same as m’s or, if p is used in m’s postcondition and m’s return type is not **void**, it also includes an additional argument that denotes m’s returned value; *iii*) p ’s body returns a single aggregable expression; *iv*) if p ’s body calls other methods, they must also be aggregable. For example, the postcondition q of a method `int fun(int x)` that always returns a value greater than the input x is expressible in BBLib as `@Predicate boolean q(int x, int result) {return gt(result, x);}`. Postcondition predicates may also refer to a method’s pre-state by means of `old(a)` expressions, which evaluates to a ’s value in the method’s pre-state.

For simplicity of presentation, we henceforth abuse the notation and use an identifier to denote a predicate method’s *name*, the *declaration* of the predicate method, and the *expression* that the predicate method returns. Consider, for instance, Fig. 2a’s precondition; in BBLib syntax, it can be declared as `@Require("null_or_open")`, where "null_or_open" is the name of a method `null_or_open`, whose body returns expression `eq(r, null) | not(r.closed)`, which is `into`’s actual precondition.

Exceptional Postconditions. Predicate methods specified with `@Ensure` are evaluated on a method’s post-state regardless of whether the method terminated normally or with an exception; for example, predicate `x_eq_y` in Fig. 3 says that attribute x always equals argument y when method m terminates. To specify exceptional behavior, a method’s postcondition predicate may include an additional argument e of type `Throwable`, which denotes the thrown exception if the method terminated with an exception or satisfies BBLib’s predicate `isVoid(e)` if the method terminated normally. Predicate `x_pos` in Fig. 3 is an example of exceptional behavior, as it says that m throws an exception of type `PosXExc` when attribute x is greater than zero upon termination; conversely, predicate `y_neg` specifies that m terminates normally when argument y is negative or zero.

SHORTHAND	EQUIVALENT POSTCONDITION
<code>@Raise(exception = E.class, when = p)</code>	<code>@Ensure(implies(old(p), e instanceof E))</code>
<code>@Return(when = p)</code>	<code>@Ensure(implies(old(p), isVoid(e)))</code>
<code>@Return</code>	<code>@Return(when = true)</code>

Table 2: BBLib’s `@Raise` and `@Return` and annotation shorthands.

Shorthands. For convenience, BBLib offers annotation shorthands `@Raise` and `@Return` to specify when a method terminates exceptionally or normally. Tab. 2 shows the semantics of these shorthands by translating them into equivalent postconditions. The `when` argument refers to a method’s pre-state through the `old` expression, since it is common to relate a method’s exceptional behavior to its inputs. Thus, Fig. 3’s postcondition `y_neg` is equivalent to `@Return(lte(y, 0))`, since m does not change y ’s value; conversely, `@Raise(PosXExc.class, gt(x, 0))` is *not* equivalent to `x_pos` because m sets x to y ’s value.

Intermediate Specification. BBLib also supports the usual intra-method specification elements: assertions, assumptions, and loop invariants. Given an aggregable ex-

```

29 class C {
30
31     int x = 0;
32
33     @Ensure("x_eq_y") // x = y when m terminates normally or exceptionally
34     @Ensure("x_pos") // if x > 0 then m throws an exception
35     @Ensure("y_neg") // if y ≤ 0 then m terminates normally
36     void m(int y) { x = y; if (x > 0) throw new PosXExc(); }
37
38     @Predicate public boolean y_neg(int y, Throwable e)
39     { return implies(lte(y, 0), isVoid(e)); }
40
41     @Predicate public boolean x_pos(int y, Throwable e)
42     { return implies(gt(x, 0), e instanceof PosXExc); }
43
44     @Predicate public boolean x_eq_y(int y)
45     { return eq(x, y); }
46
47 }

```

Fig. 3: Examples of exceptional postconditions in BBlib.

pression e , **assertion**(e) specifies that e must hold whenever execution reaches it; **assumption**(e) restricts verification from this point on to only executions where e holds; and **invariant**(e) declares that e is an invariant of the loop within whose body it is declared. As we have seen in Fig. 2a’s running example, loop invariants hold, in particular, at all exit points of a loop—including exceptional ones.

3.2 The Vimp Intermediate Representation

In our previous work [20], BYTEBACK works directly on Grimp—a high-level bytecode representation provided by the Soot static analysis framework [11, 26]. Compared to raw bytecode, Grimp conveniently retains information such as types and expressions, which eases BYTEBACK’s encoding of the program under verification into Boogie [3]. However, Grimp remains a form of bytecode, and hence it represents well executable instructions, but lacks support for encoding logic expressions and specification constructs. These limitations become especially inconvenient when reasoning about exceptional behavior, which often involves logic conditions that depend on the types and values of exceptional objects. Rather than reconstructing this information during the translation from Grimp to Boogie, we found it more effective to extend Grimp into Vimp, which fully supports logic and specification expressions.

Our bespoke Vimp bytecode representation can encode all the information relevant for verification. This brings several advantages: *i*) it decouples the input program’s static analysis from the generation of Boogie code, achieving more flexibility at either ends of the toolchain; *ii*) it makes the generation of Boogie code straightforward (mostly one-to-one); *iii*) BYTEBACK’s transformation from Grimp to Vimp becomes naturally *modular*: it composes several simpler transformations, each taking care of a different aspect and incorporating a different kind of information. The rest of this section presents

Vimp’s key features, and how they are used by `BYTEBACK`’s Grimp-to-Vimp transformation \mathcal{V} .¹ As detailed in Sec. 3.4, \mathcal{V} composes the following feature-specific transformations: \mathcal{V}_{exc} makes the exceptional control flow explicit; \mathcal{V}_{agg} aggregates Grimp expressions into compound Vimp expressions; $\mathcal{V}_{\text{inst}}$ translates Grimp instructions (by applying transformation \mathcal{V}_{stm} to statements, transformation \mathcal{V}_{exp} to expressions within statements, and $\mathcal{V}_{\text{types}}$ to expression types); $\mathcal{V}_{\text{loop}}$ handles loop invariants.

Expression Aggregation. Transformation \mathcal{V}_{agg} *aggregates* specification expressions (see Sec. 3.1), so that each corresponds to a single Vimp pure and branchless expression. In a nutshell, $\mathcal{V}_{\text{agg}}(s)$ takes a piece s of aggregable Grimp code, converts it into static-single assignment form, and then recursively replaces each variable’s single usage with its unique definition. For example, consider Fig. 2a’s loop invariant: it corresponds to `a := lte(0, i); b := lte(i, a.length); c := a & b` in Grimp, and becomes `c := lte(0, i) & lte(i, a.length)` in Vimp.

Expected Types. Transformation $\mathcal{V}_{\text{type}}$ reconstructs the *expected type* of expressions when translating them to Vimp. An expression e ’s expected type depends on the context where e is used; in general, it differs from e ’s type in Grimp, since Soot’s type inference may not distinguish between Boolean and integer expressions—which both use the `int` bytecode type.

Boolean Expressions. Another consequence of bytecode’s lack of a proper **boolean** representation is that Grimp uses integer operators also as Boolean operators (for example the unary minus `-` for “not”). In contrast, Vimp supports the usual Boolean operators `¬`, `∧`, `∨`, `⇒`, and constants `true` and `false`. Transformation \mathcal{V}_{exp} uses them to translate Vimp expressions e whose expected type $\mathcal{V}_{\text{type}}(e)$ is boolean; this includes specification expressions (which use `BLib`’s replacement operators), but also regular Boolean expressions in the executable code. For example, $\mathcal{V}_{\text{exp}}(-a) = \neg \mathcal{V}_{\text{exp}}(a)$, $\mathcal{V}_{\text{exp}}(k) = \text{true}$ for every constant $k \geq 1$, and $\mathcal{V}_{\text{exp}}(h) = \text{false}$ for every constant $h < 1$.

Transformation $\mathcal{V}_{\text{exp}}(e)$ also identifies quantified expressions—expressed using a combination of `BLib`’s `Contract` and `Binding`—after aggregating them, and renders them using Vimp’s quantifier syntax:

$$\begin{aligned} \mathcal{V}_{\text{exp}}(\text{Contract.forall}(\text{Binding.T}(), e)) &= \forall \mathcal{V}_{\text{type}}(\text{Binding.T}()) \ v :: \mathcal{V}_{\text{exp}}(e) \\ \mathcal{V}_{\text{exp}}(\text{Contract.exists}(\text{Binding.T}(), e)) &= \exists \mathcal{V}_{\text{type}}(\text{Binding.T}()) \ v :: \mathcal{V}_{\text{exp}}(e) \end{aligned}$$

Assertion Instructions. Vimp includes instructions **assert**, **assume**, and **invariant**, which transformation \mathcal{V}_{stm} introduces for each corresponding instance of `BLib` assertions, assumptions, and loop invariants. Transformation $\mathcal{V}_{\text{loop}}$ relies on Soot’s loop analysis capabilities to identify loops in Vimp’s unstructured control flow; then, it expresses their invariants by means of assertions and assumptions. As shown in Fig. 4, $\mathcal{V}_{\text{loop}}$ checks that the invariant holds upon loop entry (label *head*), at the end of each iteration (*head* again), and at every exit point (label *exit*).

¹ In the following, we occasionally take some liberties with Grimp and Vimp code, using a readable syntax that mixes bytecode instruction and Java statement syntax; for example, `m()` represents an invocation of method `m` that corresponds to a suitable variant of bytecode’s `invoke`.

<pre> k = 0; while (k < 10) { invariant(lte(k, 10) & lte(k, X)); k++; if (k ≥ X) break; } return k; </pre>	<pre> k := 0; head: if k ≥ 10 goto exit; invariant k ≤ 10 ∧ k ≤ X; k := k + 1; if k ≥ X goto exit; back: goto head; exit: return k; </pre>	<pre> k := 0; head: assert k ≤ 10 ∧ k ≤ X; if k ≥ 10 goto exit; assume k ≤ 10 ∧ k ≤ X; k := k + 1; if k ≥ X goto exit; back: goto head; exit: assert k ≤ 10 ∧ k ≤ X; return k; </pre>
---	--	---

Fig. 4: A loop in Java (left), its unstructured representation in Vimip (middle), and the transformation $\mathcal{V}_{\text{loop}}$ of its invariant into assertions and assumptions (right).

<pre> try { x = o.size(); if (x = 0) throw new E(); } catch (E e) { x = 1; } </pre>	<pre> ℓ₁: x := o.size(); ℓ₂: if x != 0 goto ℓ₅; ℓ₃: e := new E(); ℓ₄: throw e; ℓ₅: goto ℓ₆; hE: e := @caught; x := 1; ℓ₆: ... </pre>	<pre> ℓ₁: x := o.size(); if @thrown = void goto skip₂; V_{exc}(throw @thrown); skip₂: ℓ₂: if x ≠ 0 goto ℓ₅; ℓ₃: e := new E(); ℓ₄: @thrown := e; if ¬(@thrown instanceof E) goto skip₅; goto hE; skip₅: ℓ₅: goto ℓ₆; hE: e := @thrown; @thrown := void; x := 1; ℓ₆: ... </pre>
---	---	---

Fig. 5: A try-catch block in Java (left), its unstructured representation as a trap in Grimp (middle, empty lines are for readability), and its transformation \mathcal{V}_{exc} in Vimip with explicit exceptional control flow (right).

3.3 Modeling Exceptional Control Flow

Bytecode stores a block’s exceptional behavior in a data structure called the *exception table*.^[3] Soot represents each table entry as a *trap*, which renders a try-catch block in Grimp bytecode. Precisely, a trap t is defined by: *i*) a block of instructions B_t that may throw exceptions; *ii*) the type E_t of the handled exceptions; *iii*) a label h_t to the handler instructions (which terminates with a jump back to the end of B_t). When executing B_t throws an exception whose type conforms to E_t , control jumps to h_t . At the beginning of the handler code, Grimp introduces $e := @caught$, which stores into a local variable e of the handler a reference $@caught$ to the thrown exception object. Fig. 5 shows an example of try-catch block in Java (left) and the corresponding trap in Grimp (middle): ℓ_1, \dots, ℓ_5 is the instruction block, E is the exception type, and hE is handler’s entry label. The rest of this section describes BYTEBACK’s transformation \mathcal{V}_{exc} , which transforms the implicit exceptional control flow of Grimp traps into explicit control flow in Vimip.

Explicit Exceptional Control-Flow. Grimp’s variable $@caught$ is called $@thrown$ in Vimip. While $@caught$ is read-only in Grimp—where it only refers to the currently handled exception— $@thrown$ can be assigned to in Vimip. This is how BYTEBACK makes exceptional control flow explicit: assigning to $@thrown$ an exception object e signals that e has been thrown; and setting $@thrown$ to `void` marks the current execution as normal. Thus, BYTEBACK’s Vimip encoding sets $@thrown := \text{void}$ at the beginning of a program’s execution, and then manipulates the special variable to reflect the bytecode semantics of exceptions as we outline in the following. With this approach, the Vimip encoding of a **try** block simply results from encoding each of the block’s instructions explicitly according to their potentially exceptional behavior.

Throw Instructions. Transformation \mathcal{V}_{exc} desugars **throw** instructions into explicit assignments to **@thrown** and jumps to the suitable handler. A **throw e** instruction within the blocks of n traps t_1, \dots, t_n —handling exceptions of types E_1, \dots, E_n with handlers at labels h_1, \dots, h_n —is transformed into:

$$\mathcal{V}_{\text{exc}}(\text{throw } e) = \left(\begin{array}{l} \text{@thrown := e;} \\ \text{if } \neg(\text{@thrown instanceof } E_1) \text{ goto } \underline{\text{skip}}_1; \\ \text{goto } h_1; \\ \underline{\text{skip}}_1: \text{if } \neg(\text{@thrown instanceof } E_2) \text{ goto } \underline{\text{skip}}_2; \\ \text{goto } h_2; \\ \underline{\text{skip}}_2: \text{if } \neg(\text{@thrown instanceof } E_3) \text{ goto } \underline{\text{skip}}_3; \\ \vdots \\ \underline{\text{skip}}_{n-1}: \text{if } \neg(\text{@thrown instanceof } E_n) \text{ goto } \underline{\text{skip}}_n; \\ \text{goto } h_n; \\ \underline{\text{skip}}_n: \text{return; // propagate exception to caller} \end{array} \right)$$

The assignment to **@thrown** stores a reference to the thrown exception object e ; then, a series of checks determine if e has type that conforms to any of the handled exception types; if it does, execution jumps to the corresponding handler.

Transformation \mathcal{V}_{exc} also replaces the assignment $e := \text{@caught}$ that Grimp puts at the beginning of every handler with $e := \text{@thrown}; \text{@thrown} := \text{void}$, signaling that the current exception is handled, and thus the program will resume normal execution.

Exceptions in Method Calls. A called method may throw an exception, which the caller should propagate or handle. Accordingly, transformation \mathcal{V}_{exc} adds after every method call instructions to check whether the caller set variable **@thrown** and, if it did, to handle the exception within the caller as if it had been directly thrown by it.

$$\mathcal{V}_{\text{exc}}(m(a_1, \dots, a_m)) = \left(\begin{array}{l} m(a_1, \dots, a_m); \\ \text{if } (\text{@thrown} = \text{void}) \text{ goto } \underline{\text{skip}}; \\ \mathcal{V}_{\text{exc}}(\text{throw @thrown}) \\ \underline{\text{skip}}: \text{/* code after call */} \end{array} \right)$$

Potentially Excepting Instructions. Some bytecode instructions may implicitly throw exceptions when they cannot execute normally. In Fig. 2a’s running example, $r.\text{read}()$ throws a `NullPointerException` exception if r is `null`; and the assignment to $a[i]$ throws an `IndexOutOfBoundsException` exception if i is not between 0 and $a.\text{length} - 1$. Transformation \mathcal{V}_{exc} recognizes such potentially excepting instructions and adds explicit checks that capture their implicit exceptional behavior. Let op be an instruction that throws an exception of type E_{Op} when condition F_{Op} holds; \mathcal{V}_{exc} transforms op as follows.

$$\mathcal{V}_{\text{exc}}(op) = \left(\begin{array}{l} \text{if } \neg F_{\text{Op}} \text{ goto } \underline{\text{normal}}; \\ \mathcal{V}_{\text{exc}} \left(\begin{array}{l} e := \text{new } E_{\text{Op}}(); \\ \text{throw } e; \end{array} \right) \\ \underline{\text{normal}}: op \end{array} \right)$$

By chaining multiple checks, transformation \mathcal{V}_{exc} handles instructions that may throw multiple implicit exceptions. For example, here is how it encodes the potentially

excepting semantics of array lookup $a[i]$, which fails if a is **null** or i is out of bounds.

$$\mathcal{V}_{\text{exc}}(\text{res} := a[i]) = \left(\begin{array}{l} \text{if } \neg(a = \text{null}) \text{ goto } \underline{\text{normal}}_1; \\ \mathcal{V}_{\text{exc}} \left(\begin{array}{l} e_1 := \text{new NullPointerException}(); \\ \text{throw } e_1; \end{array} \right) \\ \underline{\text{normal}}_1: \text{if } \neg(0 \leq i \wedge i < a.\text{length}) \text{ goto } \underline{\text{normal}}_2; \\ \mathcal{V}_{\text{exc}} \left(\begin{array}{l} e_2 := \text{new IndexOutOfBoundsException}(); \\ \text{throw } e_2; \end{array} \right) \\ \underline{\text{normal}}_2: \text{res} := a[i]; \end{array} \right)$$

3.4 Transformation Order and Boogie Code Generation

BYTEBACK applies the transformations \mathcal{V} from Grimp to Vimp in a precise order that incrementally encodes the full program semantics respecting dependencies.



Like raw bytecode, the source Grimp representation on which BYTEBACK operates is a form of three-address code, where each instruction performs exactly one operation (a call, a dereferencing, or a unary or binary arithmetic operation). *i*) BYTEBACK first applies \mathcal{V}_{exc} to make the exceptional control flow explicit. *ii*) Then, it aggregates expressions using \mathcal{V}_{agg} . *iii*) Transformation $\mathcal{V}_{\text{inst}}$ is applied next to every instruction; in turn, $\mathcal{V}_{\text{inst}}$ relies on transformations \mathcal{V}_{exp} , \mathcal{V}_{stm} , and $\mathcal{V}_{\text{type}}$ (presented in Sec. 3.2) to process the expressions and types manipulated by the instructions (the instructions themselves do not change, and hence $\mathcal{V}_{\text{inst}}$'s definition is straightforward). *iv*) Finally, it applies $\mathcal{V}_{\text{loop}}$ to encode loop invariants as intermediate assertions; since this transformation is applied after \mathcal{V}_{exc} , the loop invariants can be checked at all loop exit points—normal and exceptional.

The very last step \mathcal{B} of BYTEBACK's pipeline takes the fully transformed Vimp program and encodes it as a Boogie program. Thanks to Vimp's design, and to the transformation \mathcal{V} applied to Grimp, the Vimp-to-Boogie translation is straightforward. In addition, BYTEBACK also generates a detailed Boogie axiomatization of all logic functions used to model various parts of JVM execution—which we described in greater detail in previous work [20]. One important addition is an axiomatization of subtype relations among exception types, used by Boogie's `instanceof` function to mirror the semantics of the homonymous Java operator. Consider the whole tree T of exception types used in the program:² each node is a type, and its children are its direct subtypes. For every node C in the tree, BYTEBACK produces one axiom asserting that every child X of C is a subtype of C ($X \preceq C$), and one axiom for every pair X, Y of C 's children asserting that any descendant types x of X and y of Y are *not* related by subtyping (in other words, the subtrees rooted in X and Y are disjoint): $\forall x, y: \text{Type} \bullet x \preceq X \wedge y \preceq Y \implies x \not\preceq y \wedge y \not\preceq x$.

3.5 Implementation Details

BBLib as Specification Language. We are aware that BBLib's syntax and conventions may be inconvenient at times; they were designed to deal with the fundamental con-

² Since the root of all exception types in Java is `class Throwable`—a concrete class—an exception type cannot be a subtype of multiple exception classes, and hence T is strictly a tree.

EXCEPTING INSTRUCTIONS	EXCEPTION	CONDITION
dereferencing <code>o_</code> , <code>o[i]</code>	<code>NullPointerException</code>	<code>o = null</code>
array access <code>a[i]</code>	<code>IndexOutOfBoundsException</code>	$\neg (\theta \leq i \wedge i < a.length)$

Table 3: Potentially excepting instructions currently supported by BYTEBACK.

straints that any specification must be expressible in the source code and still be fully available for analysis in bytecode after compilation. This rules out the more practical approaches (e.g., comments) adopted by source-level verifiers. More user-friendly notations could be introduced on top of BBlib—but doing so is outside the present paper’s scope.

Attaching Annotations. As customary in deductive verification, BYTEBACK models calls using the modular semantics, whereby every called method needs a meaningful specification of its effects within the caller. To support more realistic programs that call to Java’s standard library methods, BBlib supports the `@Attach` annotation: a class `S` annotated with `@Attach(I.class)` declares that any specification of any method in `S` serves as specification of any method with the same signature in `I`. We used this mechanism to model the fundamental behavior of widely used methods in Java’s, Scala’s, and Kotlin’s standard libraries. As a concrete example, we specified that the constructors of common exception classes do not themselves raise exceptions.

Implicit Exceptions. Sec. 3.3 describes how BYTEBACK models potentially excepting instructions. The mechanism is extensible, and the current implementation supports the ubiquitous `NullPointerException` and `IndexOutOfBoundsException` exceptions, as shown in Tab. 3. Users can selectively enable or disable these checks for implicitly thrown exceptions either for each individual method, or globally for the whole program.

Dependencies. BYTEBACK is implemented as a command-line tool that takes as input a classpath and a set E of class files within that path. The analysis collects all classes on which the entry classes in E recursively depend—where “ A depends on B ” means that A inherits from or is a client of B . After collecting all dependencies, BYTEBACK feeds them through its verification toolchain (Fig. 1) that translates them to Boogie. In practice, BYTEBACK is configured with a list of *system packages*—such as `java.lang`—that are treated differently: their implementations are ignored (i.e., not translated to Boogie for verification), but their interfaces and any specifications are retained to reason about their clients. This makes the verification process more lightweight,

Features and Limitations. The main limitations of BYTEBACK’s previous version [20] were a lack of support for exception handling and `invokedynamic`. As discussed in the rest of the paper, BYTEBACK now fully supports reasoning about exceptional behavior. We also added a, still limited, support for `invokedynamic`: any instance of `invokedynamic` is conservatively treated as a call whose effects are unspecified; furthermore, we introduced ad hoc support to reason about concatenation and comparison of string literal—which are implemented using `invokedynamic` since Java 9.³ A full support of `invokedynamic` still belongs to future work.

³ <https://docs.oracle.com/javase/9/docs/api/java/lang/invoke/StringConcatFactory.html>

Other remaining limitations of BYTEBACK’s current implementation are a limited support of string objects, and no modeling of numerical errors such as overflow (i.e., numeric types are encoded with infinite precision). Adding support for all of these features is possible by extending BYTEBACK’s current approach.

4 Experiments

We demonstrated BYTEBACK’s capabilities by running its implementation on a collection of annotated programs involving exceptional behavior in Java, Scala, and Kotlin.

4.1 Programs

Tab. 4 lists the 37 programs that we prepared for these experiments; all of them involve *some* exceptional behavior in different contexts.⁴ More than half of the programs (20/37) are in Java: 17 only use language features that have been available since Java 8, and another 3 rely on more recent features available since Java 17. To demonstrate how targeting bytecode makes BYTEBACK capable of verifying, at least in part, other JVM languages, we also included 9 programs written in Scala (version 2 of the language), and 8 programs written in Kotlin (version 1.8.0). Each program/experiment consists of one or more classes with their dependencies, which we annotated with BBlib to specify exceptional and normal behavior, as well as other assertions needed for verification (such as loop invariants). The examples total 7 810 lines of code and annotations, with hundreds of annotations and 1 070 methods (including BBlib specification methods) involved in the verification process. According to their features, the experiments can be classified into two groups: feature experiments and algorithmic experiments.

Feature Experiments. Java 8 programs 1–7, Java 17 program 18, Scala programs 21–25, and Kotlin programs 30–33 are *feature* experiments: each of them exercises a small set of exception-related language features; correspondingly, their specifications check that BYTEBACK’s verification process correctly captures the source language’s semantics of those features. For example, experiments 4, 24, and 32 feature different combinations of try-catch blocks and throw statements that can be written in Java, Scala, and Kotlin, and test whether BYTEBACK correctly reconstructs all possible exceptional and normal execution paths that can arise. A more specialized example is experiment 5, which verifies the behavior of loops with both normal and exceptional exit points.

Algorithmic Experiments. Java 8 programs 8–17, Java 17 programs 19–20, Scala programs 26–29, and Kotlin programs 34–37 are *algorithmic* experiments: they implement classic algorithm that also use exceptions to signal when their inputs are invalid. The main difference between feature and algorithmic experiments is specification: algorithmic experiments usually have more complex pre- and postconditions than feature experiments, which they complement with specifications of exceptional behavior on the

⁴ We focus on these exception-related programs, but the latest version of BYTEBACK also verifies correctly the 35 other programs we introduced in previous work to demonstrate its fundamental verification capabilities.

corner cases. For example, experiments 8, 26, and 34 implement array reversal algorithms; their postconditions specify that the input array is correctly reversed; and other parts of their specification say that they result in an exception if the input array is null. Experiment 20 is an extension of Fig. 2’s running example, where the algorithm is a simple stream-to-array copy implemented in a way that may give rise to various kinds of exceptional behavior.

Experiments 16 and 17 are the most complex programs in our experiments: they include a subset of the complete implementations of Java’s `ArrayList` and `LinkedList` standard library classes,^[4] part of which we annotated with basic postconditions and a specification of their exceptional behavior (as described in their official documentation). In particular, `ArrayList`’s exceptional specification focuses on possible failures of the class constructor (for example, when given a negative number as initial capacity); `LinkedList`’s specification focuses on possible failures of some of the read methods (for example, when trying to get elements from an empty list). Thanks to BBLib’s features (including the `@Attach` mechanism described in Sec. 3.5), we could add annotations without modifying the implementation of these classes. Note, however, that we verified relatively simple specifications, focusing on exceptional behavior; a dedicated support for complex data structure functional specifications [9,22] exceeds BBLib’s current capabilities and belongs to future work.

Implicit Exceptions. As explained in Sec. 3.5, users of `BYTEBACK` can enable or disable checking of implicitly thrown exceptions. Experiments 2, 20, 22, and 31 check implicit `null`-pointer exceptions; experiments 1, 20, 21, and 30 check implicit out-of-bounds exceptions; all other experiments do not use any implicitly thrown exceptions, and hence we disabled the corresponding checks.

4.2 Results

All experiments ran on a Fedora 36 GNU/Linux machine with an Intel Core i9-12950HX CPU (4.9GHz), running Boogie 2.15.8.0, Z3 4.11.2.0, and Soot 4.3.0. To account for measurement noise, we repeated the execution of each experiment five times and report the average wall-clock running time of each experiment, split into `BYTEBACK` bytecode-to-Boogie encoding and Boogie verification of the generated Boogie program. We ran Boogie with default options except for experiment 19, which uses the `/infer:j` option (needed to derive the loop invariant of the enhanced `for` loop, whose index variable is implicit in the source code).

All of the experiments verified successfully. To sanity-check that the axiomatization or any other parts of the encoding introduced by `BYTEBACK` are consistent, we also ran Boogie’s so-called smoke test on the experiments;⁵ these tests inject `assert false` in reachable parts of a Boogie program, and check that none of them pass verification.

As you can see in Tab. 4, `BYTEBACK`’s running time is usually below 1.5 seconds; and so is Boogie’s verification time. Unsurprisingly, programs 16 and 17 are outliers, since they are made of larger classes with many dependencies; these slow down both `BYTEBACK`’s encoding process and Boogie’s verification, which have to deal with many annotations and procedures to analyze and verify.

⁵ Smoke tests provide no absolute guarantee of consistency, but are often practically effective.

#	EXPERIMENT	LANG	ENCODING	VERIFICATION	SOURCE	BOOGIE	MET	ANNOTATIONS		
			TIME [s]		SIZE [LOC]		<i>P</i>	<i>S</i>	<i>E</i>	
1	Implicit Index Out of Bounds	J 8	1.2	1.0	87	366	16	5	2	8
2	Implicit Null Dereference	J 8	1.0	1.0	84	429	26	4	0	10
3	Multi-Catch	J 8	1.2	0.9	67	311	10	1	1	4
4	Throw-Catch	J 8	1.1	1.1	164	504	46	10	0	17
5	Throw-Catch in Loop	J 8	1.1	1.0	97	398	11	1	0	9
6	Try-Finally	J 8	1.0	1.0	125	386	15	2	4	6
7	Try-With-Resources	J 8	1.3	1.2	199	1 635	26	3	4	13
8	Array Reverse	J 8	1.1	1.1	72	221	9	5	3	2
9	Binary Search	J 8	1.2	0.8	52	169	6	4	4	1
10	GCD	J 8	1.1	0.8	50	200	6	3	1	1
11	Linear Search	J 8	1.2	0.8	62	193	9	6	6	2
12	Selection Sort (double)	J 8	1.2	1.9	110	234	16	9	9	3
13	Selection Sort (int)	J 8	1.1	3.2	110	234	16	9	9	3
14	Square of Sorted Array	J 8	1.1	0.8	61	187	7	4	1	1
15	Sum	J 8	1.1	0.8	45	175	5	2	1	1
16	ArrayList	J 8	5.7	5.7	2 653	7 160	294	14	0	24
17	LinkedList	J 8	2.1	2.9	2 472	3 041	366	8	2	17
18	Try-With-Resources on Local	J 17	1.0	0.9	44	220	6	1	1	1
19	Summary	J 17	0.9	0.8	48	171	5	2	2	1
20	Read Resource	J 17	1.1	0.8	117	447	14	12	6	7
21	Implicit Index Out of Bounds	S 2	1.2	0.8	44	231	7	2	1	4
22	Implicit Null Dereference	S 2	1.0	0.9	43	275	6	1	0	6
23	Multi-Catch	S 2	1.2	0.9	45	297	7	1	1	2
24	Throw-Catch	S 2	1.1	1.1	121	460	22	7	1	12
25	Try-Finally	S 2	1.3	1.0	117	455	15	2	4	3
26	Array Reverse	S 2	1.1	1.2	62	221	8	4	2	2
27	Counter	S 2	1.2	0.8	48	183	8	3	3	4
28	GCD	S 2	1.1	0.8	51	203	5	2	1	1
29	Linear Search	S 2	1.0	0.8	46	155	6	4	3	1
30	Implicit Index Out of Bounds	K 1.8	1.2	0.8	45	279	7	2	4	4
31	Implicit Null Dereference	K 1.8	1.2	0.9	41	309	6	1	0	6
32	Throw-Catch	K 1.8	1.2	1.1	121	442	22	7	0	12
33	Try-Finally	K 1.8	1.3	1.0	108	409	15	2	4	3
34	Array Reverse	K 1.8	1.2	1.0	60	226	8	4	2	2
35	Counter	K 1.8	1.2	0.9	46	177	8	3	3	4
36	GCD	K 1.8	1.1	0.8	50	202	5	2	1	1
37	Linear Search	K 1.8	1.3	0.8	43	193	6	4	3	1
total			47.8	44.0	7 810	21 398	1 070	156	89	199
average			1.3	1.2	211	578	29	4	2	5

Table 4: Verification experiments with exceptional behavior used to demonstrate BYTEBACK’s capabilities. Each row reports: a numeric identifier # and a short description of the EXPERIMENT; the source programming LANGUAGE (Java 8, Java 17, Scala 2, Kotlin 1.8); the wall-clock time (in seconds) taken for ENCODING bytecode into Boogie, and for the VERIFICATION of the Boogie program; the size (in non-empty lines of code) of the SOURCE program with its annotations, and of the generated BOOGIE program; the number of METHODS that make up the program and its BLib specification; and the number of ANNOTATIONS introduced for verification, among: specification predicates *P* (**@Predicate**), pre- and postconditions *S* (**@Require**, **@Ensure**), and exception annotations *E* (**@Raise**, **@Return**).

Only about 11% of the time listed under Tab. 4’s column ENCODING is taken by BYTEBACK’s actual encoding; the rest is spent to perform class resolution (Sec. 3.5) and to initialize Soot’s analysis—which dominate BYTEBACK’s overall running time.

5 Related Work

The state-of-the-art deductive verifiers for Java include OpenJML [6], KeY [1], and Krakatoa [7]; they all process the source language directly, and use variants of JML specification language—which offers support for specifying exceptional behavior.

Exceptional Behavior Specifications. Unlike BBlib, where postconditions can refer to both exceptional and normal behavior, JML clearly separates between the two, using **ensures** and **signals** clauses (as demonstrated in Fig. 6). These JML features are supported by OpenJML, KeY, and Krakatoa according to their intended semantics.

<pre>//@ ensures this.a == a; //@ signals (Throwable) this.a == a; public void m(int a) { this.a = a; if (e) throw new RuntimeException(); }</pre>	<pre>@Ensure(this.a = a) public void m(int a) { this.a = a; if (e) throw new RuntimeException(); }</pre>
--	--

Fig. 6: Equivalent exceptional specifications in JML (left) and BBlib (right).

Implicit Exceptional Behavior. Implicitly thrown exceptions, such as those occurring when accessing an array with an out-of-bounds index, may be handled in different ways by a verifier: *i*) ignore such exceptions; *ii*) implicitly check that such exceptions never occur; *iii*) allow users to specify these exceptions like explicit ones. OpenJML and Krakatoa [12] follow strategy *ii*), which is sound but loses some precision since it won’t verify some programs (such as Sec. 2’s example); KeY offers options to select any of these strategies, which gives the most flexibility; BYTEBACK offers options *i*) and *iii*), so that users can decide how thorough the analysis of exceptional behavior should be.

Java Exception Features. OpenJML, KeY, and Krakatoa [7] all support try-catch-finally blocks, which have been part of Java since its very first version. The first significant extension to exceptional feature occurred with Java 7, which introduced multi-catch and try-with-resources blocks.⁽⁵⁾ KeY and Krakatoa support earlier versions of Java, and hence they cannot handle either feature. OpenJML supports many features of Java up to version 8, and hence can verify programs using multi-catch or try-with-resources—with the exception of try-with-resources using an existing **final** variable, a feature introduced only in Java 9. As usual, our point here is not to criticize these state-of-the-art verification tools, but to point out how handling the proliferation of Java language features becomes considerably easier when targeting bytecode following BYTEBACK’s approach.

Intermediate Representation Verifiers. A different class of verifiers—including JayHorn [10, 25], SeaHorn [8], and SMACK [23]—target intermediate representations (JVM bytecode for JayHorn, and LLVM bitcode for SeaHorn and SMACK). Besides this similarity, these tools’ capabilities are quite different from BYTEBACK’s: they implement analyses based on model-checking (with verification conditions expressible as constrained Horn clauses, or other specialized logics), which provide a high degree

of automation (e.g., they do not require loop invariants) to verify simpler, lower-level properties (e.g., reachability). Implicitly thrown exceptions are within the purview of tools like JayHorn, which injects checks before each instruction that may dereference a null pointer, access an index out of bounds, or perform an invalid cast. In terms of usage, this is more similar to a specialized static analysis tool that checks the absence of certain runtime errors [2, 4, 21] than to fully flexible, but onerous to use, deductive verifiers like BYTEBACK.

BML [5] is a specification language for bytecode; since it is based on JML, it is primarily used as a way of expressing a high-level Java behavioral specification at the bytecode level. This is useful for approaches to proof-carrying code [18] and proof transformations [16], where one verifies a program’s source-code and then certifies its bytecode compilation by directly transforming the proof steps.

6 Conclusions

Reasoning about exceptional behavior at the level of Java bytecode facilitates handling exception-related features in any version of Java, as well as in other JVM languages like Scala and Kotlin. More generally, the BYTEBACK approach that we extended in this paper can complement the core work in source-level deductive verification and make it readily available to the latest languages and features.

References

1. Ahrendt, W., Beckert, B., Bruns, D., Bubel, R., Gladisch, C., Grebing, S., Hähnle, R., Hentschel, M., Herda, M., Klebanov, V., Mostowski, W., Scheben, C., Schmitt, P.H., Ulrich, M.: The KeY platform for verification and analysis of Java programs. In: Gianakopoulou, D., Kroening, D. (eds.) *Verified Software: Theories, Tools and Experiments - 6th International Conference, VSTTE 2014, Vienna, Austria, July 17-18, 2014, Revised Selected Papers*. Lecture Notes in Computer Science, vol. 8471, pp. 55–71. Springer (2014). https://doi.org/10.1007/978-3-319-12154-3_4, https://doi.org/10.1007/978-3-319-12154-3_4
2. Banerjee, S., Clapp, L., Sridharan, M.: NullAway: practical type-based null safety for Java. In: Dumas, M., Pfahl, D., Apel, S., Russo, A. (eds.) *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2019, Tallinn, Estonia, August 26-30, 2019*. pp. 740–750. ACM (2019). <https://doi.org/10.1145/3338906.3338919>, <https://doi.org/10.1145/3338906.3338919>
3. Barnett, M., Chang, B.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.P. (eds.) *Formal Methods for Components and Objects, 4th International Symposium, FMCO 2005, Amsterdam, The Netherlands, November 1-4, 2005, Revised Lectures*. Lecture Notes in Computer Science, vol. 4111, pp. 364–387. Springer (2005). https://doi.org/10.1007/11804192_17, https://doi.org/10.1007/11804192_17
4. Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O’Hearn, P.W., Papakonstantinou, I., Purbrick, J., Rodriguez, D.: *Moving Fast with Software Verification*.

- In: Havelund, K., Holzmann, G.J., Joshi, R. (eds.) *NASA Formal Methods - 7th International Symposium, NFM 2015, Pasadena, CA, USA, April 27-29, 2015, Proceedings. Lecture Notes in Computer Science*, vol. 9058, pp. 3–11. Springer (2015). https://doi.org/10.1007/978-3-319-17524-9_1, https://doi.org/10.1007/978-3-319-17524-9_1
5. Chrzaszcz, J., Huisman, M., Schubert, A.: BML and related tools. In: de Boer, F.S., Bonsangue, M.M., Madelaine, E. (eds.) *Formal Methods for Components and Objects, 7th International Symposium, FMCO 2008, Sophia Antipolis, France, October 21-23, 2008, Revised Lectures. Lecture Notes in Computer Science*, vol. 5751, pp. 278–297. Springer (2008). https://doi.org/10.1007/978-3-642-04167-9_14, https://doi.org/10.1007/978-3-642-04167-9_14
 6. Cok, D.R.: OpenJML: Software verification for Java 7 using JML, OpenJDK, and Eclipse. In: Dubois, C., Giannakopoulou, D., Méry, D. (eds.) *Proceedings 1st Workshop on Formal Integrated Development Environment, F-IDE 2014, Grenoble, France, April 6, 2014. EPTCS*, vol. 149, pp. 79–92 (2014). <https://doi.org/10.4204/EPTCS.149.8>, <https://doi.org/10.4204/EPTCS.149.8>
 7. Filliâtre, J., Marché, C.: The Why/Krakatoa/Caduceus platform for deductive program verification. In: Damm, W., Hermanns, H. (eds.) *Computer Aided Verification, 19th International Conference, CAV 2007, Berlin, Germany, July 3-7, 2007, Proceedings. Lecture Notes in Computer Science*, vol. 4590, pp. 173–177. Springer (2007). https://doi.org/10.1007/978-3-540-73368-3_21, https://doi.org/10.1007/978-3-540-73368-3_21
 8. Gurfinkel, A., Kahsai, T., Komuravelli, A., Navas, J.A.: The SeaHorn verification framework. In: Kroening, D., Pasareanu, C.S. (eds.) *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 9206, pp. 343–361. Springer (2015). https://doi.org/10.1007/978-3-319-21690-4_20, https://doi.org/10.1007/978-3-319-21690-4_20
 9. Hiep, H.A., Maathuis, O., Bian, J., de Boer, F.S., de Gouw, S.: Verifying OpenJDK’s LinkedList using KeY (extended paper). *Int. J. Softw. Tools Technol. Transf.* **24**(5), 783–802 (2022). <https://doi.org/10.1007/s10009-022-00679-7>, <https://doi.org/10.1007/s10009-022-00679-7>
 10. Kahsai, T., Rümmer, P., Sanchez, H., Schäf, M.: JayHorn: A framework for verifying Java programs. In: Chaudhuri, S., Farzan, A. (eds.) *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 9779, pp. 352–358. Springer (2016). https://doi.org/10.1007/978-3-319-41528-4_19, https://doi.org/10.1007/978-3-319-41528-4_19
 11. Lam, P., Bodden, E., Lhoták, O., Hendren, L.: The Soot framework for Java program analysis: a retrospective. In: *Cetus Users and Compiler Infrastructure Workshop (CETUS 2011) (Oct 2011)*, <https://www.bodden.de/pubs/lb1h11soot.pdf>
 12. Marché, C., Paulin-Mohring, C., Urbain, X.: The KRAKATOA tool for certification of JAVA/JAVACARD programs annotated in JML. *J. Log. Algebraic Methods Program.* **58**(1-2), 89–106 (2004). <https://doi.org/10.1016/j.jlap.2003.07.006>, <https://doi.org/10.1016/j.jlap.2003.07.006>
 13. Marcilio, D., Furia, C.A.: How Java programmers test exceptional behavior. In: *Proceedings of the 18th Mining Software Repositories Conference (MSR)*. pp. 207–218. IEEE (May 2021)
 14. Marcilio, D., Furia, C.A.: What is thrown? Lightweight precise automatic extraction of exception preconditions in Java methods. In: *Proceedings of the 38th IEEE International Conference on Software Maintenance and Evolution (ICSME)*. pp. 340–351. IEEE Computer Society (October 2022)

15. Melo, H., Coelho, R., Treude, C.: Unveiling exception handling guidelines adopted by Java developers. In: SANER. IEEE (2019)
16. Müller, P., Nordio, M.: Proof-transforming compilation of programs with abrupt termination. In: Proceedings of SAVCBS. pp. 39–46. ACM (2007), <https://doi.org/10.1145/12923216.1292321>
17. Nakshatri, S., Hegde, M., Thandra, S.: Analysis of exception handling patterns in Java projects: An empirical study. IEEE/ACM MSR (2016)
18. Necula, G.C.: Proof-carrying code. In: Lee, P., Henglein, F., Jones, N.D. (eds.) POPL. pp. 106–119. ACM Press (1997). <https://doi.org/10.1145/263699.263712>
19. Paganoni, M., Furia, C.A.: ByteBack iFM 2023 Replication Package (Sep 2023). <https://doi.org/10.5281/zenodo.8335240>, <https://doi.org/10.5281/zenodo.8335240>
20. Paganoni, M., Furia, C.A.: Verifying functional correctness properties at the level of Java bytecode. In: Chechik, M., Katoen, J.P., Leucker, M. (eds.) Proceedings of the 25th International Symposium on Formal Methods (FM). Lecture Notes in Computer Science, vol. 14000, pp. 343–363. Springer (March 2023)
21. Papi, M.M., Ali, M., Jr., T.L.C., Perkins, J.H., Ernst, M.D.: Practical pluggable types for Java. In: Ryder, B.G., Zeller, A. (eds.) Proceedings of the ACM/SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2008, Seattle, WA, USA, July 20–24, 2008. pp. 201–212. ACM (2008). <https://doi.org/10.1145/1390630.1390656>, <https://doi.org/10.1145/1390630.1390656>
22. Polikarpova, N., Tschannen, J., Furia, C.A.: A fully verified container library. *Formal Aspects of Computing* **30**(5), 495–523 (September 2018)
23. Rakamaric, Z., Emmi, M.: SMACK: decoupling source language details from verifier implementations. In: Biere, A., Bloem, R. (eds.) Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18–22, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8559, pp. 106–113. Springer (2014). https://doi.org/10.1007/978-3-319-08867-9_7, https://doi.org/10.1007/978-3-319-08867-9_7
24. Rudich, A., Darvas, Á., Müller, P.: Checking well-formedness of pure-method specifications. In: Cuéllar, J., Maibaum, T.S.E., Sere, K. (eds.) FM 2008: Formal Methods, 15th International Symposium on Formal Methods, Turku, Finland, May 26–30, 2008, Proceedings. Lecture Notes in Computer Science, vol. 5014, pp. 68–83. Springer (2008). https://doi.org/10.1007/978-3-540-68237-0_7, https://doi.org/10.1007/978-3-540-68237-0_7
25. Rümmer, P.: JayHorn: a Java model checker. In: Murray, T., Ernst, G. (eds.) Proceedings of the 21st Workshop on Formal Techniques for Java-like Programs, FTfJP@ECOOP 2019, London, United Kingdom, July 15, 2019. p. 1:1. ACM (2019). <https://doi.org/10.1145/3340672.3341113>, <https://doi.org/10.1145/3340672.3341113>
26. Vallée-Rai, R., Co, P., Gagnon, E., Hendren, L.J., Lam, P., Sundaresan, V.: Soot – a Java bytecode optimization framework. In: MacKay, S.A., Johnson, J.H. (eds.) Proceedings of the 1999 conference of the Centre for Advanced Studies on Collaborative Research, November 8–11, 1999, Mississauga, Ontario, Canada. p. 13. IBM (1999), <https://dl.acm.org/citation.cfm?id=782008>
27. Weimer, W., Necula, G.C.: Finding and preventing run-time error handling mistakes. In: Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2004, October 24–28, 2004, Vancouver, BC, Canada. pp. 419–431. ACM (2004). <https://doi.org/10.1145/1028976.1029011>, <https://doi.org/10.1145/1028976.1029011>

URL References

1. <https://jcp.org/en/jsr/detail?id=334>
2. <https://docs.oracle.com/javase/tutorial/essential/exceptions/runtime.html>
3. <https://docs.oracle.com/javase/specs/jvms/se7/html/jvms-2.html#jvms-2.10>
4. <https://docs.oracle.com/javase/8/docs/api/java/util/package-summary.html>
5. <https://www.oracle.com/java/technologies/javase/jdk7-relnotes.html>