

# Abstract Execution<sup>\*</sup>

Dominic Steinhöfel<sup>[0000–0003–4439–7129]</sup> and Reiner Hähnle<sup>[0000–0001–8000–7613]</sup>

TU Darmstadt, Dept. of Computer Science, Darmstadt, Germany  
{steinhoefel,haehnle}@cs.tu-darmstadt.de

**Abstract.** We propose a new static software analysis principle called *Abstract Execution*, generalizing Symbolic Execution: While the latter analyzes all possible execution paths of a *specific program*, Abstract Execution analyzes a *partially unspecified program* by permitting *abstract symbols* representing unknown contexts. For each abstract symbol, we faithfully represent each possible concrete execution resulting from its substitution with concrete code. There is a wide range of applications of Abstract Execution, especially for verifying *relational properties* of schematic programs. We implemented Abstract Execution in a deductive verification framework and proved correctness of eight well-known statement-level refactoring rules, including two with loops. For each refactoring we characterize the preconditions that make it semantics-preserving. Most preconditions are not mentioned in the literature.

## 1 Introduction

Reasoning about *abstract programs*, i.e. programs containing an abstract context represented by placeholder symbols, is required whenever one aims to rigorously analyze program transformation techniques. Notably in compiler validation, to argue that a specific compilation or optimization step preserves the meaning of any input program is a standard task. An established approach to this problem formalizes the abstract syntax and the semantics of the target programming language as a set of inductive definitions, then proves properties of abstract programs via structural induction over the program syntax [7]. Early work relied on pen-and-paper proofs [22,24]. Recently, interactive theorem provers are used to mechanize correctness proofs, e.g., in CompCert [21] and CakeML [31]. The main drawback is the very high effort required to mechanize a programming language and to perform interactive proofs. In this paper we take a different approach to reason about abstract programs that is automatic and based on symbolic execution. To make it work, we need to answer two questions: (i) Can one specify abstract program contexts sufficiently without giving full inductive definitions? (ii) If yes, which specification constructs are needed for abstract contexts?

We propose a new static software analysis principle called Abstract Execution (AE) that allows to *automatically* reason about *abstract* sequential programs

---

<sup>\*</sup> This work was funded by the Hessian LOEWE initiative within the Software-Factory 4.0 project. **The final authenticated version is available online at [https://doi.org/10.1007/978-3-030-30942-8\\_20](https://doi.org/10.1007/978-3-030-30942-8_20).**

with *side effects*. An essential component of AE is a specification language for abstract program contexts. This permits, in contrast to prior work [13,14,18,29], to specify irregular termination behavior (exceptions, etc.) and fine-grained assumptions on abstract programs. Here we target sequential JAVA programs, but the principles of AE are equally applicable to other sequential languages.

Abstract Execution generalizes Symbolic Execution (SE) [5,8,17]. Symbolic Execution means to execute programs with symbolic expressions as input values. When SE is embedded into a program logic [1], these symbolic values are represented by first-order Skolem constants. Skolem symbols can also be viewed as *abstract* symbols whose concrete domain value is not specified. In this sense, SE is already abstract and it amounts to execution of *abstract* programs that permit variables and fields to be initialized with unknown values. For example, *symbolic* execution of a program “i++;” amounts to execution of the *abstract* program “ $i_0++$ ”, where  $i_0$  is a fresh Skolem symbol of type **int**.

AE in addition permits not only abstract *values*, but whole *programs* to appear as undefined expressions. In the program “**if** (i>0)  $p_0$  **else**  $p_1$ ”, for example, the placeholders  $p_i$  can be substituted with *arbitrary* concrete programs, as long as the result is well-formed. From a semantic point of view, the difference between abstract values and abstract programs is that (i) an abstract program may change the value of arbitrary variables and fields, and (ii) its execution may terminate abruptly, i.e., by returning, throwing an exception, or continuing or breaking out of a loop. In other words, an abstract program can be seen as an unknown partial function between execution states, i.e. its big-step semantics  $\llbracket p \rrbracket$ . In logic, such a function can be represented by a *second-order* Skolem symbol. Therefore, AE amounts to a limited form of second-order reasoning: let  $p$  be an abstract program and  $\varphi$  a first-order formula. Then  $\vdash_{ae} [p]\varphi$  holds iff the weakest precondition of  $p$  w.r.t.  $\varphi$  is first-order derivable. For example:

$$\vdash_{ae} [\mathbf{int} \ i; \ \mathbf{boolean} \ b=b_0; \ \mathbf{if} \ (b) \ p_0 \ \mathbf{else} \ i=0;](b_0 \doteq \mathbf{FALSE} \rightarrow i \doteq 0)$$

To check the relation  $\vdash_{ae}$  *automatically*, we must provide a technical solution to the following question: how can one implement weakest precondition reasoning over *abstract* programs such that the result is a *first-order* formula? These are the main building blocks of our solution to realize limited second-order inference over programs in terms of first-order deduction: (1) *second-order* Skolemization to represent the effects of placeholders like  $p_0$  on local variables and the heap, (2) explicit modeling of all possible ways of irregular termination in separate SE branches, (3) over-approximation of returned values and thrown exceptions by first-order Skolemization, and (4) a specification language to describe the possible effect of second-order Skolem symbols as well as to define conditions when irregular termination of concrete instances for abstract programs can happen.

*Applications of Abstract Execution* AE is applicable to many problems involving reasoning about abstract programs. It can be instantiated to (at least) the following tasks: (1) Execution of abstract method calls [6], a special case of AE; (2) automatic soundness proofs of program transformation and of (3) rule-based compi-

lation [29]; (4) sound, automatic (“lazy”) symbolic execution over programs with loops and calls, (5) incremental program development and synthesis [28]; (6) a modular version of *proof carrying code* [26], where the contract of called methods needs not to be known to the certificate provider. It is impossible to discuss all these in this paper. Therefore, we focus on program transformation—task (2). We study refactoring rules as described in Fowler’s well-known books [11,12]. We model refactoring techniques as abstract programs, formalize assumptions under which a refactoring is sound, and prove behavioral equivalence of the original and refactored version for all concrete programs satisfying the abstract context.

The paper is structured as follows. In Sect. 2 we describe how to construct *abstract JAVA* programs. Sect. 3 expounds our logic for AE. Sect. 4 contains our case study about proving correctness of refactoring techniques. Sect. 5 discusses related work, Sect. 6 concludes and gives an outlook to the future. An appendix with more material is available at [key-project.org/papers/ae/](http://key-project.org/papers/ae/).

## 2 Specifying Abstract Programs

An abstract JAVA program is a program containing at least one *Abstract Placeholder Statement (APS)* symbol. The syntax to declare an APS is:

```
abstract_statement P;
```

The symbol P is an identifier for an abstract statement. Semantically, every APS with the same identifier occurring in a program or proof represents the same program. The above APS may be substituted with any *concrete* JAVA program accessing and assigning arbitrary fields and local variables, except that it is not allowed to declare local variables visible outside. Additionally, a concrete program may (1) throw any type of exception, (2) return from the method it executes, (3) break to any surrounding block label, (4) continue to and break from a surrounding loop, (5) continue to the label of a surrounding loop.

The possible behaviors of an APS are constrained by an *abstract specification*. The syntax of the specification language extends *block contracts* [1,19] of the JAVA Modelling Language (JML). JML [20] is a specification language for JAVA used to describe the behavior of JAVA classes and methods. JML specifications are embedded into JAVA code via comment lines starting with an “@” sign. *An APS is the declaration of an abstract placeholder symbol together with all specification clauses that constrain it.* We explain the involved concepts by specifying a variant of Fowler’s *Consolidate Duplicate Conditional Fragments* refactoring [11] step-by-step. The result, a fully specified program after refactoring is shown in Listing 3 (on p. 5). Table 1 summarizes all specification constructs that may be used in an APS.

Fig. 1 shows an unconstrained formalization of the refactoring. The abstract code uses an idiom to formalize abstract *expressions*: it introduces a variable representing the abstract expression (in this case, the boolean b) and precedes it with the APS `Init` that assigns to b an unknown value. The idiom works as expected if `Init` is constrained so it assigns a value exactly to b and not to any

Table 1: Specification constructs for APSs

Spec. Construct	Explanation
<code>locals(P)</code>	Refers to the Skolem (abstract) location set of local variables of an APS with symbol <code>P</code> visible from outside.
<code>declares skLocs;</code>	Specifies that an APS/method declares a list <code>skLocs</code> of Skolem location set specifiers <code>locals(·)</code> , opt. wrapped in <code>final(·)</code> modifiers, which can be used in APSs in the visible scope afterwards.
<code>assignable locs;</code>	Declares the location set <code>locs</code> to be assignable by the APS. <code>locs</code> is a list of variables, fields, and Skolem location set specifiers, optionally wrapped in a <code>hasTo(·)</code> modifier.
<code>accessible locs;</code>	Declares <code>locs</code> to be accessible by the APS.
<code>return_behavior requires φ;</code>	Specifies that the APS returns iff $\varphi$ holds.
<code>exceptional_behavior requires φ;</code>	Spec. that the APS throws an exc. iff $\varphi$ holds.
<code>break_behavior requires φ;</code>	Specifies that the APS breaks/continues during loop execution iff $\varphi$ holds.
<code>continue_behavior requires φ;</code>	Specifies that the APS breaks/continues to the (loop) label <code>lbl</code> iff $\varphi$ holds.
<code>break_behavior (lbl) requires φ;</code>	
<code>continue_behavior (lbl) requires φ;</code>	

Listing 1: Before

```

abstract_statement Init;
if (b) {
  abstract_statement P;
  abstract_statement Q1;
} else {
  abstract_statement P;
  abstract_statement Q2;
}

```

Listing 2: After

```

abstract_statement P;
abstract_statement Init;
if (b) {
  abstract_statement Q1;
} else {
  abstract_statement Q2;
}

```

Fig. 1: Unconstrained formalization of the “Consolidate Duplicate Conditional Fragments” Refactoring, “Pullout Prefix” variant [11]

other variable. JML uses the `assignable` clause to specify which locations can be assigned a value, but does not *enforce* the assignment. Hence, we extend JML with the `hasTo(·)` keyword. The specification “`//@ assignable hasTo(b);`” enforces that the specified abstract code assigns a value *exactly* to `b`.

Observe that the refactoring is unsound, whenever the APS `P` influences the value of `b`. If, for instance, `P` sets `b` to `true`, the `else` branch of the `if` statement in the refactored program is never reached. A drastic solution is to specify “`//@ ⊥ assignable \nothing;`” for `P` which excludes any assignment. This, however, restricts the refactoring rule too severely to be useful. Assume the depicted program fragments occur in the scope of a method with a variable `result` that is returned at the end. Then we might constrain `P` with “`//@ assignable result;`” which forbids assignments to `b` (because it allows assignments *exactly* to `result`), but renders `P` still useful. But this is not restrictive enough: The abstract program `Init` that initializes `b` may still *access* arbitrary locations and assign them

Listing 3: Fully specified abstract program for the refactored version

---

```

/*@ axiom mutex{throwsExc(P),
    throwsExc(Init), returns(P)};
    if (b) {
        /*@ declares locals(Q1);
            /*@ assignable \everything;
            /*@ accessible \everything;
            abstract_statement Q1;
        } else {
            /*@ declares locals(Q2);
            /*@ assignable \everything;
            /*@ accessible \everything;
            abstract_statement Q2;
        }

/*@ declares locals(P);
/*@ assignable locals(P), result;
/*@ accessible locals(P), result, args;
/*@ return_behavior requires returns(P);
/*@ exceptional_behavior requires throwsExc(P);
abstract_statement P;

/*@ assignable hasTo(b);
/*@ accessible args;
/*@ return_behavior requires false;
/*@ exceptional_behavior requires throwsExc(Init);
abstract_statement Init;

```

---

to `b`. Thus, `P` can indirectly influence the control flow by assigning a value to the variable `result` which could then affect `Init`'s choice for `b`. To address this issue we proceed as follows:

We add to all APSs in Fig. 1 (except `Init`) a “**declares**” annotation. It declares abstract “Skolem” location sets that can be instantiated with arbitrary concrete local variable declarations visible from outside. For example, to `P` we add the annotation “`/*@ declares locals(P);`”. To specify that a method containing APSs receives an unknown set “`args`” of parameters, we annotate it with “`/*@ declares args;`”. Abstract location sets in declarations can be declared final by surrounding them with “`final(·)`”. This prevents them from occurring in **assignable** clauses. Continuing the example, we add to `P` the annotation “`/*@ assignable locals(P), result;`”, to `Init` “`/*@ accessible args;`”.

Proving correctness of the refactoring still fails, however, for two reasons. The first is: we have not excluded that `Init` contains a **return** statement. Since `Init`'s only task is to initialize `b`, this should never happen. We add the annotation “`/*@ return_behavior requires false;`” to `Init`, specifying that a **return** requires the specified condition—here falsity—which excludes returning.

The second reason why the refactoring is not yet correct is that `Init` might raise an exception. This is entirely possible and a real problem: If we permit `P` to return or throw an exception, but `Init` may also throw an exception, then the refactored and the original program have different behavior. We have two options: (i) We deny `Init` to throw an exception by adding the annotation “`/*@ exceptional_behavior requires false;`” or, more generally, (ii) we enforce that *if* `Init` throws an exception, then `P` can neither throw an exception nor can it return. The latter is achieved with the help of the abstract functions `throwsExc(Init)`, `throwsExc(P)`, and `returns(P)` of *Boolean* type. They qualify the **requires** clauses that restrict exceptional and returning behavior of `Init` and `P`. A global **axiom** declares them to be mutually exclusive.

Listing 3 shows the specified refactored program. Similar annotations apply to the original version. For `Q1` and `Q2`, we permit assigning/accessing all loca-

$$\text{assignment} \frac{\Gamma \vdash \{\mathcal{U}\}\{x := e\}[\pi \ \omega]\varphi, \Delta}{\Gamma \vdash \{\mathcal{U}\}[\pi \ x=e; \ \omega]\varphi, \Delta} \quad e \text{ is side effect-free}$$

Fig. 2: SE rule for variable assignment

tions (“`\everything`”). This is default, so these declarations can be left out. The annotation “`//@ declares final(args);`” would have to be added to the surrounding method declaration. With the specification in Listing 3 we can prove *equivalence* of the original and refactored program for *any concrete instance* matching the abstract program structurally and satisfying its I/O and control flow constraints. Having proved equivalence, the refactoring can be applied in either direction. This is relevant: many of Fowler’s refactorings are bi-directional. We continue with formalizing AE over constrained APSs in a program logic.

### 3 Abstract Execution Logic

Our implementation of AE is realized on top of the symbolic execution framework of the deductive verification system KeY [1]. It is based on JAVA Dynamic Logic (JavaDL), a program logic for the JAVA language.

#### 3.1 Principles of JavaDL

JavaDL extends sorted First-Order Logic (FOL). JAVA programs appear inside logical formulas as *modalities*, of which there are two types: The box modality  $[p]\varphi$  expresses that *if* program  $p$  terminates, then the postcondition  $\varphi$  holds in any final state (partial correctness). The diamond modality  $\langle p \rangle \varphi$  additionally requires  $p$  to terminate (total correctness). To prove the validity of formulas, i.e.  $[p]\varphi$  or  $\langle p \rangle \varphi$  holds in any initial state, JavaDL has a sequent calculus comprising FOL and theory-specific rules, as well as rules realizing SE of JAVA programs. A Hoare triple  $\{\psi\} p \{\varphi\}$  is equivalent to the JavaDL formula  $\psi \rightarrow [p]\varphi$ .

The SE rules of the JavaDL calculus reduce a JAVA statement to first-order assumptions and a separate syntactic category called symbolic *updates* representing symbolic state transitions. The atomic building blocks of updates are the empty update *Skip*, representing the identity state transition, and the *elementary update*  $x := t$  for the state transition where variable  $x$  is assigned the value of term  $t$ . Two updates  $\mathcal{U}_1, \mathcal{U}_2$  can be combined into a *parallel* update  $\mathcal{U}_1 \parallel \mathcal{U}_2$ : the state changes of  $\mathcal{U}_1$  and  $\mathcal{U}_2$  are executed *simultaneously*; in case both assign to the same variable, the assignment in  $\mathcal{U}_2$  “wins”. Only in the absence of such “conflicts”, parallel composition is commutative. Updates are applied to terms  $t$  and formulas  $\varphi$ :  $\{\mathcal{U}\}t$  and  $\{\mathcal{U}\}\varphi$  represent the value of term  $t$  and truth value of formula  $\varphi$  after the state change effected by  $\mathcal{U}$ , respectively. Parallel update composition  $\{\mathcal{U}_1 \parallel \mathcal{U}_2\}\varphi$  is different from sequential composition  $\{\mathcal{U}_1\}\{\mathcal{U}_2\}\varphi$ . In the sequential case, right-hand sides of  $\mathcal{U}_2$  are interpreted in the state resulting from  $\mathcal{U}_1$ . In the parallel case, they are interpreted in the *same* pre-state. The formula  $\{\mathcal{U}_1\}\{\mathcal{U}_2\}\varphi$  is equivalent to  $\{\mathcal{U}_1 \parallel \{\mathcal{U}_1\}\mathcal{U}_2\}\varphi$ .

Fig. 2 depicts an SE rule using an update to represent the effect of an assignment of an expression without side effects to variable  $x$ . As usual, sequent calculus rules are read “bottom-up”, i.e. the rule symbolically executes the assignment by turning it into a symbolic update. SE rules operate on the first active (i.e. executable) statement of a program, here the assignment. The remaining program is contained in  $\pi\omega$ , where the prefix  $\pi$  consists of opening braces, labels, try-blocks, etc., and the postfix  $\omega$  of closing braces, blocks, and remaining statements. The program  $\pi\omega$  is a well-formed JAVA program. Sequent calculus rules have zero or more premises (sequents on top of the rule); zero premises characterize a rule that closes a proof case, more than *one* premise causes a proof to split. An example of the latter is the rule for an **if**-statement (see web appendix). The conclusion (bottom part) of a rule consists of exactly one sequent. A rule is *sound* if the correctness of the conclusion follows from the correctness of all premises. A sequent  $\Gamma \vdash \Delta$  is correct if the conjunction of the formulas in the set  $\Gamma$  implies the disjunction of those in  $\Delta$ . Details are in [1].

We need a recently introduced concept of JavaDL for reasoning about loops: *loop scopes* [30]. A loop scope  $\circlearrowleft_x p_x \circlearrowright$  is a scope for a loop body  $p$ . It results from SE of a loop **while**( $b$ ){ $p$ }. The boolean flag  $x$  encodes *completion information* about the loop: it is set to TRUE if the loop is *exited* (either normally or by irregular termination) and to FALSE if it *continues* with another iteration. Using the value of  $x$ , a postcondition can distinguish both cases.

### 3.2 Formalization of Abstract Execution

We first give a definition of the domain of *locations* used in **declares**, **accessible** and **assignable** specifications of APSs. The symbol “*allLocs*” is introduced for the “\everything” specifier in **accessible** and **assignable** specifications.

**Definition 1.** *The set  $Locs^{Concr}$  of concrete locations consists of program variables  $x$  and, for an object  $o$  and field identifier  $f$ , field locations  $(o, f)$ . The set  $Locs^{Sk}$  of Skolem location sets consists of uninterpreted functions  $loc^{Sk}$  representing arbitrary sets of concrete locations  $Locs^{Concr}$ . We define  $Locs = Locs^{Concr} \cup Locs^{Sk} \cup \{allLocs\}$ , where the symbol *allLocs* represents all concrete locations  $Locs^{Concr}$ . The set of assignable locations is defined as  $Locs^{Assgn} = Locs \cup \{loc^! \mid loc \in Locs \setminus \{allLocs\}\}$  which also includes “have-to” locations  $loc^!$ .*

In Sect. 2, we introduced the specification elements of APSs. These are formalized in the subsequent definition of the logic representation of an APS.

**Definition 2.** *Let  $id$  be an identifier symbol,  $decls \subseteq Locs^{Sk}$ ,  $assignables \subseteq Locs^{Assgn}$  and  $accessibles \subseteq Locs$ . An Abstract Placeholder Statement is a tuple*

$$(id, decls, assignables, accessibles, specs)$$

where “*specs*” represents behavioral specifications and is a tuple of the form:

$$(returnsSpec, excSpec, continuesSpec, breaksSpec, continuesSpecLbl, breakSpecLbl)$$

The elements  $returnsSpec$ ,  $excSpec$ ,  $continuesSpec$ ,  $breaksSpec$  are optional: they are the empty set or a singleton of a formula specifying when an APS returns, throws an exception, continues, and breaks, respectively. Elements  $breakSpecLbl$ ,  $continuesSpecLbl$  are partial functions from labels to formulas, specifying when an APS continues a labeled loop or breaks from a labeled block or loop.

We write “APS P” short for “the APS with identifier symbol P”. Abstract Execution reasons about the behavior of *all possible concrete programs* with which APSs “legally” may be instantiated, formally:

**Definition 3.** Let  $p_a$  be an abstract program with occurrences of APS symbols  $P_1, \dots, P_n$ . We call the substitution with concrete programs  $P_i^0$  for each  $P_i$  a legal instantiation iff (1) the result from substituting all occurrences of  $P_i$  by  $P_i^0$  in  $p_a$  is a compilable JAVA program, and (2) all  $P_i^0$  satisfy all constraints of the APSs declaring the  $P_i$ .

*Example 1.* We substitute P in Listing 3 with “`int z=result++;`”. This is legal if z is undeclared in the visible scope. The substitution instantiates the location set “`locals(P)`” with  $\{z\}$ , affecting further instantiations referring to it. The substitution of “`y=z;`” for P is illegal: First, it *assigns* a variable which is not contained in its assignable set; second, if z is not contained in the instantiation of `args`, it *accesses* an undeclared variable z; third, the program is not *compilable* if y and z are undeclared. Let `param` be in the instantiation of `args`. Substituting P with “`int z=result/param;`” and `Init` with “`b=result/param;`” is illegal since *both* could throw an exception, contradicting the axiom.

SE of an APS must over-approximate the behavior of *all* legal instantiations. To model this in the logic, we use second-order *Skolemization*. Given an APS  $(P, decls, assignables, accessibles, specs)$ , we create what we call a *Skolem update* “ $\mathcal{U}_P(assignables \approx accessibles)$ ” with *Skolem path condition* “ $C_P(accessibles)$ ” fresh for P. The term “fresh for” means that the symbols  $\mathcal{U}_P$ ,  $C_P$  are created freshly (as usual for Skolemization) when P first occurs in a proof context, but are re-used each time when P re-occurs. This ensures that each occurrence of an APS symbol represents the *same* program. We define the meaning of Skolem update and Skolem path condition by extending the notion of legal instantiation. In the definition we assume all Skolem location sets in  $Locs^{Sk}$  to be instantiated with concrete locations.

**Definition 4.** Let an APS P with  $assignables \subseteq Locs^{Assign} \setminus Locs^{Sk}$ ,  $accessibles \subseteq Locs \setminus Locs^{Sk}$  be given. An abstract update  $\mathcal{U}_P(assignables \approx accessibles)$  may be instantiated with any concrete update  $x_1 := t_1 \parallel \dots \parallel x_n := t_n$  for which the following conditions hold: (1) either all  $Locs \in accessibles$  or the  $t_i$  depend at most on locations in  $accessibles$ ; (2) either all  $Locs \in assignables$  or for each  $x_i$  one of  $x_i \in assignables$  or  $x_i^! \in assignables$ ; (3) for all  $x^! \in assignables$ , there is an  $i$  such that  $x = x_i$ . An abstract path condition  $C_P(accessibles)$  may be instantiated with any closed formula  $\varphi$  depending at most on locations in  $accessibles$ .

$$\text{simpleAERule} \frac{\Gamma \vdash \{\mathcal{U}\}\{\mathcal{U}_p(\text{allLocs} : \approx \text{allLocs})\}(C_p(\text{allLocs}) \rightarrow [\pi \ \omega]\varphi), \Delta}{\Gamma \vdash \{\mathcal{U}\}[\pi \ \text{abstract\_statement} \ P; \ \omega]\varphi, \Delta}$$

Fig. 3: Simple AE rule without abrupt termination

**Definition 5.** Abstract *JavaDL* extends *JavaDL* syntax as follows: (1) updates can be Skolem updates; (2) Skolem path conditions are also formulas; (3) programs can be abstract. Abstract sequents and sequent rules are defined as before, but range over abstract *JavaDL* formulas.

A *JavaDL* sequent calculus rule is *sound* if the validity of the conclusion follows from the validity of all premises [1]. We can leave this definition *unchanged* provided that we define validity of abstract sequents suitably:

**Definition 6.** A sequent  $S^0$  is a legal instantiation of an abstract sequent  $S$  if  $S^0$  results from substituting all Skolem updates, Skolem path conditions and APS symbols in  $S$  with legal instantiations. An abstract sequent is valid iff all its legal instantiations are valid in *JavaDL*.

One of the simplest possible AE rules (first mentioned in [29]) is shown in Fig. 3. It is only applicable for an APS whose specification and legal instantiations exclude irregular termination.

**Theorem 1.** The Abstract Execution rule `simpleAERule` (Fig. 3) is sound.

*Proof.* Let  $P^0$  be any legal instantiation of  $P$ . Since  $P^0$  cannot terminate irregularly, symbolic execution transforms the sequent  $\Gamma \vdash \{\mathcal{U}\}[\pi \ P^0 \ \omega]\varphi, \Delta$  to one of the shape  $\Gamma \vdash \{\mathcal{U}\}\{\mathcal{U}^0\}(C^0 \rightarrow [\pi \ \omega]\varphi), \Delta$ .<sup>1</sup> Assume the premise of `simpleAERule` is valid (otherwise, the rule is trivially sound). The instantiations  $\mathcal{U}^0$  of  $\mathcal{U}_p$  and  $C^0$  of  $C_p$  are legal (the *allLocs* location allows reading and writing arbitrary locations). So, by assumption,  $\Gamma \vdash \{\mathcal{U}\}\{\mathcal{U}^0\}(C^0 \rightarrow [\pi \ \omega]\varphi), \Delta$  is valid and, by soundness of SE, also the conclusion. Since  $P^0$  was chosen arbitrarily, the abstract sequent in the conclusion of `simpleAERule` is valid.  $\square$

The `simpleAERule` rule is unsatisfactory: it is too restrictive on irregular termination. It is also too abstract, because the abstract update and path condition in the premise may write and read any location. The abstract update can erase all variables and the whole heap, which prevents proving interesting properties. More useful rules can be obtained for *specific contexts* in which an APS occurs in the conclusion. Depending on the context, legal instantiations can lead to different ways of irregular termination.

Fig. 4 shows a rule for AE within a loop scope and a non-void method, but outside the scope of *loop* labels. In contrast to `simpleAERule`, `nonVoidLoopAERule`

<sup>1</sup> If the statement causes a split, like an `if` statement, we still can combine the arising sequents to a single one by state merging [27].

$$\begin{array}{c}
\text{nonVoidLoopAERule} \\
\Gamma \vdash \{\mathcal{U}\}\{\mathcal{U}_P(\text{assignables} \approx \text{accessibles})\} \\
\quad \{\text{returns} := \text{returns}_0 \parallel \text{result} := \text{result}_0 \parallel \text{exc} := \text{exc}_0 \parallel \\
\quad \text{breaks} := \text{breaks}_0 \parallel \text{continues} := \text{continues}_0 \parallel \\
\quad \text{breaksToLbl}_1 := \text{breaksToLabel}_{l_0} \parallel \dots \parallel \\
\quad \text{breaksToLbl}_n := \text{breaksToLabel}_{l_n}\} \\
\left( \begin{array}{l}
C_P(\text{accessibles}) \\
\wedge \text{mutex}(\text{returns}, \text{exc} \neq \text{null}, \text{breaksToLbl}_1, \dots, \text{breaksToLbl}_n) \\
\wedge (\text{returns} \doteq \text{TRUE} \leftrightarrow \text{returnsSpec})^? \wedge (\text{exc} \neq \text{null} \leftrightarrow \text{excSpec})^? \\
\wedge (\text{breaks} \doteq \text{TRUE} \leftrightarrow \text{breaksSpec})^? \\
\wedge (\text{continues} \doteq \text{TRUE} \leftrightarrow \text{continuesSpec})^? \\
\wedge (\text{breaksToLbl}_1 \doteq \text{TRUE} \leftrightarrow \text{breaksLbl1Spec})^? \wedge \dots \\
\wedge (\text{breaksToLbl}_n \doteq \text{TRUE} \leftrightarrow \text{breaksLblnSpec})^?
\end{array} \right) \\
\rightarrow [\pi \ l_1 : \{\dots\} \{l_n : \{ \\
\quad \circ_x \text{if}(\text{returns}) \ \mathbf{return} \ \text{result}; \ \mathbf{if}(\text{exc} \neq \text{null}) \ \mathbf{throw} \ \text{exc}; \\
\quad \quad \mathbf{if}(\text{breaks}) \ \mathbf{break}; \quad \quad \mathbf{if}(\text{continues}) \ \mathbf{continue}; \\
\quad \quad \mathbf{if}(\text{breaksToLbl}_1) \ \mathbf{break} \ l_1; \ \dots \ \mathbf{if}(\text{breaksToLbl}_n) \ \mathbf{break} \ l_n; \\
\quad \text{Rest}_1 \ \circ_x \\
\text{Rest}_2 \ \}\}\dots\} \omega] \varphi, \Delta \\
\hline
\Gamma \vdash \{\mathcal{U}\}[\pi \ l_1 : \{\dots\} \{l_n : \{ \\
\quad \circ_x \mathbf{abstract\_statement} \ P; \ \text{Rest}_1 \ \circ_x \ \text{Rest}_2 \\
\quad \}\}\dots\} \omega] \varphi, \Delta
\end{array}$$

Fig. 4: AE rule for an APS within a loop scope.

uses the *assignables/accessibles* specifications of the APS syntax. Irregular termination is modeled by **if** statements inside the loop scope in the premise. The conditionals depend on variables initialized with fresh constants in the update after the abstract update. E.g., *returns* is initialized with a constant *returns*<sub>0</sub>.

Without a specific context, SE will split at each **if** statement and follow both branches, e.g., one where P returns and one where it does not. Using the behavior specification, this can be fine-tuned: For example, in the path condition in the premise, there is an optional (marked with <sup>?</sup>) conjunct “(*returns*  $\doteq$  TRUE  $\leftrightarrow$  *returnsSpec*)<sup>?</sup>”. This lets one control the value of the guard *returns* with the formula *returnsSpec*. The behavior specifications stem from the specifications of the abstract symbol P in the conclusion as detailed in Sect. 2. The function *mutex* is interpreted such that at most one of its arguments is true at any time: here this specifies that there is not more than one reason for a program to terminate irregularly.<sup>2</sup> A proof of the following Thm. 2 is provided in the web appendix.

**Theorem 2.** *The Abstract Execution rule nonVoidLoopAERule (Fig. 4) is sound.*

*Design principles for AE rules.* The principles underlying the above rules apply to other sequential languages than JAVA as well. To create a new AE rule, we proceed as follows. Given a context in which an APS declaration is the active statement (a loop, method, labeled block, etc.), we model possible side effects of

<sup>2</sup> It is possible that, for instance, during *returning* an exception is thrown: this simply means that *exception* is the reason for termination.

that APS with separate, conditioned SE branches in the premise. For soundness it is crucial not to miss any irregular termination cases. We point out that instead of performing exhaustive structural induction, we merely distinguish *different paths of program completion*. For paths depending on values (as for a `return` or exception), Skolem constants are introduced. The conditioned premises depend on flags that establish a link to the APS’s specification; abstract updates and path conditions are added as in `nonVoidLoopAERule`. AE rules are not specific to the target application of this paper (correctness of refactoring rules), but can be used in any of the areas mentioned in the introduction.

### 3.3 Abstract Update Simplification

The JavaDL calculus comes with many simplification rules for concrete updates:  $\{\mathcal{U}_1\}\{\mathcal{U}_2\}\varphi$  simplifies to  $\{\mathcal{U}_1 \parallel \{\mathcal{U}_1\}\mathcal{U}_2\}\varphi$ , updates applied to formulas without program variables are removed, etc. In addition, spurious updates, such as those assigning variables not occurring in their scope, are removed. To reason about abstract programs, we need corresponding rules for abstract updates.

We designed a set of simplification rules for abstract updates (see Table 2): (1) Remove spurious updates: From the *assignables* part of an abstract update, delete those not occurring in the scope or that are overwritten before being read; (2) two rules handle the interplay between concrete and abstract updates; (3) two rules handle *concatenation* of abstract updates: When we cannot further simplify a formula  $\{\mathcal{U}_1\}\{\mathcal{U}_2\}\varphi$ , we connect the abstract updates by a concatenation operator, resulting in  $\{\mathcal{U}_1 \circ \mathcal{U}_2\}\varphi$ . This is not needed for concrete updates which are directly simplified as shown above. Within a concatenation, abstract updates can be commuted if their assignable/accessible specifications do not interfere.

## 4 Proving the Correctness of Refactoring Techniques

We studied five refactoring techniques from Fowler’s classic book [11] and three from the second edition [12]. We choose refactorings operating at the statement level, because they are directly expressible in JavaDL and—for the time being—exclude techniques that reorganize class hierarchies, rename constructs, or move methods. For each of the eight techniques we formalized the starting point and the result of the refactoring as a suitably specified abstract program (see Sect. 2), and then proved their *equivalence* with the AE calculus discussed in Sect. 3. Thus, we obtain soundness of, for example, *Extract Method* at the same time as of its inverse, *Inline Method*. All proofs are fully mechanized in KeY [1].

*Methodology.* For each refactoring, we create a JAVA class Refactor with two public methods: `before` contains an abstract program representing the input to the refactoring, `after` the refactored result. We start with minimal annotations in the occurring APSs including `declares` directives and standard return and assignable specifications for the abstract expressions idiom. The following JavaDL formula performs AE of an abstract program `p` on a Refactor object `o`

Table 2: Simplification rules for abstract updates

Simplification Scheme	Description
$\{\mathcal{U}_p(x, \bar{y} \approx \dots)\}\varphi(x) \rightsquigarrow \{\mathcal{U}_p(x \approx \dots)\}\varphi(x)$	Ineffective assignables are removed, abstract updates with no assignables dropped.
$\{x_1 := t_1 \parallel x_2 := t_2 \parallel \dots\} \{\mathcal{U}_p(\dots, x_1, \dots \approx \dots, \bar{x}_2, \dots)\}\varphi \rightsquigarrow \{x_1 := t_1 \parallel \dots\} \{\mathcal{U}_p(\dots, x_1, \dots \approx \dots, \bar{t}_2, \dots)\}\{\bar{x}_2 := t_2\}\varphi$	Applies variable assignments to the accessibles of the abstract update and pushes down elementary updates not assigned by the abstract update. Elementary updates that <i>have to</i> be assigned (overwritten) are dropped later.
$\{\mathcal{U}_p(\dots, \bar{y}, \dots \approx \dots)\} \{\dots \parallel \bar{x} := y \parallel \dots\}\varphi(x) \rightsquigarrow \{\mathcal{U}_p(\dots, \bar{x}, \dots \approx \dots)\} \{\dots \parallel \dots\}\varphi(x)$	Eliminates a renaming substitution “y for x”: Since the concrete update assigns to x a value chosen by $\mathcal{U}_p$ (which <i>has to</i> assign y), the abstract update can as well directly choose that value. Sound because y is not contained in $\varphi$ .
$\{\mathcal{U}_p\}\{\mathcal{U}_q\}\varphi \rightsquigarrow \{\mathcal{U}_p \circ \mathcal{U}_q\}\varphi$	Sequential to “concatenated” update application.
$\mathcal{U}_p(\text{assign}_1 \approx \text{access}_1) \circ \mathcal{U}_q(\text{assign}_2 \approx \text{access}_2) \rightsquigarrow \mathcal{U}_q(\text{assign}_2 \approx \text{access}_2) \circ \mathcal{U}_p(\text{assign}_1 \approx \text{access}_1)$	Abstract updates within concatenations can be commuted if their assignable and accessible specifiers are <i>independent</i> ; i.e., $\text{assign}_1$ has to be disjoint from $\text{assign}_2$ and $\text{access}_2$ , and similarly vice versa.

and records the result (“Flag” in the postcondition is explained later):

$$\begin{aligned} \text{AE}(p, \text{Flag}) = \{ \text{result} := \text{result}_0 \parallel \dots \} \\ \langle \mathbf{try} \{ \text{result} = \text{o.p}(\text{result}, \dots) @ \text{Refactor}; \} \\ \mathbf{catch} (\text{Throwable } t) \{ \text{result} = t; \} \rangle \text{Post}(\text{result}, \text{Flag}) \end{aligned} \quad (1)$$

Equivalence of the original and the refactored program is established by proving the formula “ $\text{AE}(\text{before}, \text{TRUE}) \leftrightarrow \text{AE}(\text{after}, \text{TRUE})$ ”. This is loaded into KeY and an automatic proof is started. In all but one refactoring technique (first in Table 3), the proof cannot be finished and open goals remain. The reason is—quite simply—that Fowler’s refactoring techniques are not sound in general [10]. As he points out, they rely on robust test suites and a “try-compile-and-test” loop trusted to unveil potential faults introduced by a refactoring.

In our setting we have the opportunity to restrict the programs that can be soundly refactored via suitable annotations of the APSs used to describe refactoring source and target. Fowler in most cases does not mention these restrictions. Fortunately, inspection of unclosable proof goals provides clear hints on the nature of the required annotations.

A typical example is when all open proof goals expect an APS  $P$  to throw an exception (assumption  $\text{exc} \neq \text{null}$  occurs in each unprovable goal). This is addressed by adding a constraint on  $P$  that forbids to throw exceptions. Another common situation concerns too liberal **accessible/assignable** specifications. These lead to open goals that contain a sequent of the form  $\{\mathcal{U}\}\varphi \vdash \{\mathcal{U}'\}\varphi$  that

Listing 4: Before

---

```

done = false; i = 0;
while (!done && i < threshold) {
  if (condition) {
    abstract_statement Body;
    done = true;
  }
  i++;
}
return result;

```

---

Listing 5: After

---

```

i = 0;
while (i < threshold) {
  if (condition) {
    abstract_statement Body;
    break;
  }
  i++;
}
return result;

```

---

Fig. 5: The *Remove Control Flag* Refactoring Technique

becomes valid when the abstract updates  $\mathcal{U}$  and  $\mathcal{U}'$  are identical. Any differences give hints on possible annotations that permit to close the proof.

Once a proof is complete, the formalization of a refactoring should be checked for validity, i.e. whether the intention of the refactoring technique has been faithfully captured. Specifications should not be more restrictive than necessary and permit substituting non-trivial programs for APSs. For example, it is easy, but useless, to find a proof with “**assignable \nothing;**” specifications. For each behavioral restriction, there should be a convincing justification. We discovered non-trivial and justifiable restrictions for almost all the investigated refactoring techniques (summarized in Table 3). Source code samples are in the appendix.

*Complexity of checking legal substitutions.* A closed equivalence proof about abstract programs asserts that those programs behave equivalently for all legal substitutions of concrete programs for APS symbols (Def. 3). Consequently, for each concrete program one must check that all the constraints specified in APSs are satisfied. These include syntactic restrictions (e.g., when inlining a method, there are no recursive calls in the body) as well as behavioral ones. The latter are not necessarily automatically checkable or even decidable. For example, to decide whether a program throws an exception, is equivalent to reachability. Even so, a formalized and proven refactoring technique makes its requirements *explicit* that before were mentioned only informally (if at all). Not in general, but in practice quite often, constraints can be proven in KeY.

*Multiple specifications.* It can make sense to create *multiple* formalizations of the same refactoring technique. The restrictions that ensure soundness can differ depending on (1) the program *context* (inside/outside a loop or labeled block), and (2) the termination mode (normal, exceptional, break, etc.). The “consolidate duplicate conditional fragments” technique in Table 3 exemplifies this.

*Programs with loops.* We studied two refactoring techniques involving loops: *Remove Control Flag* [11] and *Split Loop* [12]. To handle loops we make use of the Flag in formula (1) to separate runs leaving the loop from those leading to further executions of the loop body. The value of Flag is controlled by the loop scope parameter (the invariant rule in Fig. 7 in the appendix contains details).

Table 3: Studied refactoring techniques and discovered behavioral constraints.

Refactoring Technique	Discovered Restrictions	Justification & Remarks
Consolidate duplicate conditional fragments (Extract Postfix) [11]	none	—
Consolidate duplicate conditional fragments (Extract Prefix) [11]	<ul style="list-style-type: none"> <li>• <b>if</b> guard may only throw exc. if prefix terminates normally;</li> <li>• pulled out statement may assign heap and parameters iff guard does not access it.</li> </ul>	Irregular termination in both guard and prefix affect final result. Influence on accessibles of guard can change control flow (whether <b>if</b> or <b>else</b> is taken).
Consolidate duplicate conditional fragments (Postfix of <b>try-catch</b> to <b>finally</b> ) [11]	<ul style="list-style-type: none"> <li>• Program in <b>try</b> block may not return;</li> <li>• program in <b>catch</b> block may not return or throw exception.</li> </ul>	A <b>finally</b> block is always executed, even after a return. This changes the returned result before and after the refactoring.
Consolidate duplicate conditional fragments (Postfix of <b>try-catch</b> , no <b>finally</b> ) [11]	none	Fowler talks about moving the postfix “to the final block”, leaving unspecified whether this refers to the <b>finally</b> block or to the statements after <b>try</b> . Only in the latter case no restrictions apply.
Decompose Conditional [11]	Special case of “Extract Method”, see below.	
Extract Method [11]	Extracted fragment <ul style="list-style-type: none"> <li>• must not return;</li> <li>• may only assign heap and a local variable.</li> </ul>	A return after extraction does not affect the top-level method.
Replace Exception with Test [11]	There has to be a “roll-back” program in the <b>catch/else</b> block, program in <b>try/if</b> may only assign variables reset by that program.	The program in <b>try/if</b> may change part of the state before throwing an exception, therefore the result after the exception/test can differ.
Move Statements to Callers [12]	Neither the moved statement nor the remaining program may return.	If the remaining program in the called method returns early, the moved one is not executed after, but before refactoring.
Slide Statements [12]	All programs participating in the sliding, i.e. the swapped parts and the one in the middle <ul style="list-style-type: none"> <li>• must not return or throw an exception;</li> <li>• must be <i>independent</i>.</li> </ul>	<i>Independent</i> means the participating programs write to locations that the others do not access. Abrupt termination would change the result.
Split Loop [12]	Explained in text.	Example with loop.
Remove Control Flag [11]	Explained in text.	Example with loop.

The *Remove Control Flag* refactoring in Fig. 5 is interesting, because the number of runs of the loop before and after the refactoring differs by one (the guard needs to be executed one extra time before). This complicates the proof since we obtain  $Post(\_result, FALSE)$  for one case and  $Post(\_result, TRUE)$  for the other. We solve this by harmonizing the iteration structure via an *unrolling technique* [16] and an intermediate refactoring. Alternatively, one can code the unrolling inside a modified loop invariant rule (Fig. 9 in the appendix).

The *Split Loop* refactoring splits a loop with two independent parts into two successive loops. We had to supply several annotations to the APSs. For instance, the first part of the loop body must not break or continue (since otherwise, the second part is skipped, which is not the case after the refactoring), while the second part must not return, throw an exception, or break, since then we would have to relate runs continuing loop iteration with others exiting the loop.

*Performance* All proofs are performed automatically in KeY without user interaction. For the refactorings with loops, currently small proof scripts ( $\approx 40$  lines) are needed for loop coupling. The proofs have 2,900–40,000 nodes (median: 7,100) and take 6–300 s (median: 29 s) to complete. All problem files with detailed statistics, together with a KeY version implementing AE, are available at [key-project.org/papers/ae/](http://key-project.org/papers/ae/).

## 5 Related Work

The idea of Abstract Execution was first mentioned in our earlier work [29], where it is used to formalize the correctness of compilation rules of a JAVA-to-LLVM IR compiler. There, APSs could not be annotated and irregular termination was excluded; also, every APS can assign and access any location. In the present paper we lift these restrictions, provide an implementation and a case study. Abstract execution of APSs can be seen as a generalization of *abstract operation contracts* [6,15] to abstract *block* contracts. In the former, contracts are abstract, but programs concrete; we generalize this to abstract *programs*. This amounts to encoding limited second-order inference (no induction, no higher-order quantification) over programs into first-order (dynamic) logic.

The principal use cases for AE reside in the area of *relational verification* [4], which includes, but is not limited to: general-purpose relational proofs about programs [2,16], correctness proofs for refactorings [13], regression verification [14], proven-correct compilation [21,31] and compiler optimizations [23,18], program synthesis [28], information flow properties (e.g., by self-composition [3,9]).

There are several approaches to prove relational properties of *concrete* programs (e.g., LLRÉVE [16]). Barthe et al. [2] propose the construction of *product programs* from two variable-disjoint programs as a general-purpose technique for verifying relational program properties. After execution of the product program, the result is checked for correctness, i.e. equality. This works even for structurally different programs. Instead, we execute both programs in isolation in an equivalence proof. This has the drawback of requiring a certain structural

similarity of the programs and explicit loop coupling, but it is more resilient in the presence of irregular or non-termination. Product programs and AE are not mutually exclusive: One can create a *product* of *abstract* programs.

Garrido & Meseguer [13] prove correctness of three JAVA refactoring techniques based on an executable Maude semantics of JAVA. They focus on refactorings not targeted by us (e.g., *Pull Up Field*). It is unclear whether this approach works for statement-based refactorings including loops. Alive [23] permits proving automatically the correctness of “peephole optimizations” for LLVM. While this approach reasons about classes of programs, it is parametric only in register names and imposes other serious restrictions (e.g., no loops). Eilertsen et al. [10] generate semantic correctness assertions that ensure preservation of program semantics after refactoring. They work on concrete programs and perform run-time, not static checking.

Godlin & Strichman [14] perform “Regression Verification” by transforming loops into recursive functions and replacing recursive calls with uninterpreted function symbols. The latter are similar to APSs, however, side effects or irregular termination cannot be modeled, because functions are pure. Mechtaev et al. [25] propose a mechanism for proving existential second-order properties based on symbolic functions. Their goal is to find existential witnesses for those functions by synthesis from a user-specified grammar. In contrast, we aim at *universal* properties, and APSs represent statements (with side effects), not functions. The PEC system [18] for proving the correctness of compiler optimizations uses *meta variables* ranging over expressions, variables and statements. The latter are “single-entry-single-exit”, whereas APSs can have multiple exit points, including irregular termination. In addition, we permit annotations that constrain possible behavior. The property to be proven in [18] is a certain bi-simulation relation which is somewhat inflexible and requires lockstep execution.

## 6 Conclusion and Future Work

We proposed *Abstract Execution* of programs that contain APS symbols, a new software analysis principle extending symbolic execution. AE permits to *automatically* reason about partially unspecified programs. APSs allow irregular termination and include specification of assignable and accessible locations as well as of termination behavior. This generalizes other approaches going into similar directions. We implemented our method and applied it to eight JAVA refactoring techniques, of which two require reasoning about loops. Our formalization of the refactoring techniques makes implicit requirements explicit. It helps to better understand and safely apply refactorings. We plan to investigate how to support structurally different programs (e.g., comparing iterative and recursive versions of the same algorithm), concurrent programs, and we intend to look at other application areas. To prove the correctness of compiler optimizations automatically using AE is a natural follow-up to our work on refactoring techniques.

## References

1. Ahrendt, W., Beckert, B., Bubel, R., Hähnle, R., Schmitt, P., Ulbrich, M. (eds.): Deductive Software Verification—The KeY Book: From Theory to Practice, LNCS, vol. 10001. Springer (2016)
2. Barthe, G., Crespo, J.M., Kunz, C.: Relational Verification Using Product Programs. In: Butler, M.J., Schulte, W. (eds.) FM: Formal Methods, 17th Intl. Symp. on Formal Methods, Limerick, Ireland. LNCS, vol. 6664, pp. 200–214. Springer (2011)
3. Barthe, G., D’Argenio, P.R., Rezk, T.: Secure Information Flow by Self-Composition. In: 17th IEEE Computer Security Foundations Workshop, CSFW-17, Pacific Grove, CA, USA. pp. 100–114. IEEE Computer Society (2004)
4. Beckert, B., Ulbrich, M.: Trends in Relational Program Verification. In: Principled Software Development - Essays Dedicated to Arnd Poetzsch-Heffter on the Occasion of his 60th Birthday. pp. 41–58 (2018)
5. Boyer, R.S., Elspas, B., Levitt, K.N.: SELECT—A Formal System for Testing and Debugging Programs by Symbolic Execution. ACM SIGPLAN Notices **10**(6), 234–245 (Jun 1975)
6. Bubel, R., Hähnle, R., Pelevina, M.: Fully Abstract Operation Contracts. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation, 6th International Symposium, ISoLA 2014, Corfu, Greece. LNCS, vol. 8803, pp. 120–134. Springer (Oct 2014)
7. Burstall, R.M.: Proving Properties of Programs by Structural Induction. The Computer Journal **12**(1), 41–48 (1969)
8. Burstall, R.M.: Program Proving as Hand Simulation with a Little Induction. In: Information Processing ’74, pp. 308–312. Elsevier/North-Holland (1974)
9. Darvas, A., Hähnle, R., Sands, D.: A Theorem Proving Approach to Analysis of Secure Information Flow. In: Hutter, D., Ullmann, M. (eds.) Proc. 2nd International Conference on Security in Pervasive Computing. LNCS, vol. 3450, pp. 193–209. Springer (2005)
10. Eilertsen, A.M., Bagge, A.H., Stolz, V.: Safer Refactorings. In: Proc. 7th Intern. Symp. on Leveraging Applications of Formal Methods, ISoLA. pp. 517–531 (2016)
11. Fowler, M.: Refactoring: Improving the Design of Existing Code. Object Technology Series, Addison-Wesley (Jun 1999)
12. Fowler, M.: Refactoring: Improving the Design of Existing Code. Addison-Wesley Signature Series, Addison-Wesley Professional (Nov 2018), 2nd edition
13. Garrido, A., Meseguer, J.: Formal Specification and Verification of Java Refactorings. In: Proc. 6th IEEE Intern. Workshop on Source Code Analysis and Manipulation. pp. 165–174. SCAM ’06, IEEE Computer Society (2006)
14. Godlin, B., Strichman, O.: Regression Verification: Proving the Equivalence of Similar Programs. Softw. Test., Verif. Reliab. **23**(3), 241–258 (2013)
15. Hähnle, R., Schaefer, I., Bubel, R.: Reuse in Software Verification by Abstract Method Calls. In: Bonacina, M.P. (ed.) Proc. 24th Conference on Automated Deduction (CADE), Lake Placid, USA. LNCS, vol. 7898, pp. 300–314. Springer (2013)
16. Kiefer, M., Klebanov, V., Ulbrich, M.: Relational Program Reasoning Using Compiler IR - Combining Static Verification and Dynamic Analysis. J. Autom. Reasoning **60**(3), 337–363 (2018)
17. King, J.C.: Symbolic Execution and Program Testing. Communications of the ACM **19**(7), 385–394 (Jul 1976)
18. Kundu, S., Tatlock, Z., Lerner, S.: Proving Optimizations Correct Using Parameterized Program Equivalence. In: Proc. PLDI 2009. pp. 327–337 (2009)

19. Lanzinger, F.: A Divide-and-Conquer Strategy with Block and Loop Contracts for Deductive Program Verification. Bachelor Thesis, Institute of Theoretical Informatics, Karlsruhe Institute of Technology (April 2018)
20. Leavens, G.T., Poll, E., Clifton, C., Cheon, Y., Ruby, C., Cok, D., Müller, P., Kiniry, J., Chalin, P., Zimmerman, D.M., Dietl, W.: JML Reference Manual (May 2013), <http://www.eecs.ucf.edu/~leavens/JML//OldReleases/jmlrefman.pdf>, draft revision 2344
21. Leroy, X.: Formal Verification of a Realistic Compiler. *Communications of the ACM* **52**(7), 107–115 (2009)
22. London, R.L.: Correctness of a Compiler for a Lisp Subset. In: *Proc. of ACM Conf. on Proving Assertions About Programs*. pp. 121–127. ACM (1972)
23. Lopes, N.P., Menendez, D., Nagarakatte, S., Regehr, J.: Practical Verification of Peephole Optimizations with Alive. *Commun. ACM* **61**(2), 84–91 (2018)
24. McCarthy, J., Painter, J.: Correctness of a Compiler for Arithmetic Expressions. *Mathematical Aspects of Computer Science* **1** (1967)
25. Mehtaev, S., Griggio, A., Cimatti, A., Roychoudhury, A.: Symbolic Execution with Existential Second-Order Constraints. In: *Proc. 2018 Joint Meeting on European Software Engineering Conf. and Symp. on the Foundations of Software Engineering*. pp. 389–399 (2018)
26. Necula, G.C.: Proof-carrying code. In: *Proc. 24th ACM Symposium on Principles of Programming Languages, Paris, France*. pp. 106–119. ACM Press (Jan 1997)
27. Scheurer, D., Hähnle, R., Bubel, R.: A General Lattice Model for Merging Symbolic Execution Branches. In: Ogata, K., Lawford, M., Liu, S. (eds.) *Proc. 18th Intern. Conf. on Formal Engineering Methods (ICFEM)*. LNCS, vol. 10009, pp. 57–73. Springer (Nov 2016)
28. Srivastava, S., Gulwani, S., Foster, J.S.: From Program Verification to Program Synthesis. In: *Proc. 37th POPL*. pp. 313–326 (2010)
29. Steinhöfel, D., Hähnle, R.: Modular, Correct Compilation with Automatic Soundness Proofs. In: *Proc. 8th Intern. Symp. on Leveraging Applications of Formal Methods (ISoLA)*. LNCS, vol. 11244, pp. 424–447. Springer (2018)
30. Steinhöfel, D., Wasser, N.: A New Invariant Rule for the Analysis of Loops with Non-standard Control Flows. In: Polikarpova, N., Schneider, S. (eds.) *Proc. of 13th Intern. Conf. on Integrated Formal Methods*. pp. 279–294. Springer (2017)
31. Tan, Y.K., Myreen, M.O., Kumar, R., Fox, A., Owens, S., Norrish, M.: A New Verified Compiler Backend for CakeML. In: *Proc. 21st Intern. Conf. on Functional Programming*. pp. 60–73. ACM (2016)