
From Trees to DAGs: A General Lattice Model for Symbolic Execution

**Von Bäumen zu gerichteten azyklischen Graphen: Ein allgemeines
Verbandsmodell für symbolische Ausführung**

Master-Thesis von Dominic Scheurer aus Darmstadt

Tag der Einreichung:

1. Gutachten: Prof. Dr. Reiner Hähnle
 2. Gutachten: Nathan Wasser, Dr. Richard Bubel
-



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Fachbereich Informatik
Software Engineering (SE)

From Trees to DAGs:

A General Lattice Model for Symbolic Execution

Von Bäumen zu gerichteten azyklischen Graphen: Ein allgemeines Verbandsmodell für symbolische Ausführung

Vorgelegte Master-Thesis von Dominic Scheurer aus Darmstadt

1. Gutachten: Prof. Dr. Reiner Hähnle

2. Gutachten: Nathan Wasser, Dr. Richard Bubel

Tag der Einreichung:

Erklärung zur Master-These

Hiermit versichere ich, die vorliegende Master-These ohne Hilfe Dritter nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die aus Quellen entnommen wurden, sind als solche kenntlich gemacht. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Darmstadt, den 9. April 2015

(D. Scheurer)

Abstract

Symbolic Execution is a precise static program analysis technique for software testing and verification. In the course of the analysis, programs are transformed into symbolic execution trees containing up to exponentially many branches in the number of branch points. We address this so-called “path explosion problem” in the context of program verification by proposing a general lattice-based framework for join operations that allows for the merging of branches during symbolic execution. Several concrete join techniques are presented as instances of this framework and are implemented for the deductive verification system KeY. We show that our operations indeed reduce the number of states and branches significantly for certain examples, and apply a join technique to information flow analysis in a short case study to demonstrate that state joining can increase the precision of analyses in principle.

Zusammenfassung (German Abstract)

Symbolische Ausführung ist eine präzise Technik zur statischen Analyse von Programmen im Bereich des Testens und der Verifikation von Software. Im Zuge der Analyse werden Programme in symbolische Ausführungsbäume transformiert, welche bis zu exponentiell viele Zweige (in Abhängigkeit von der Zahl der Verzweigungspunkte) enthalten. Wir gehen dieses sogenannte “Pfadexplosionsproblem” im Kontext der Programmverifikation an, indem wir ein allgemeines verbandsbasiertes Rahmenwerk für Verbindungsoperationen angeben, welches die Zusammenführung von Zuständen während der symbolischen Ausführung erlaubt. Verschiedene konkrete Verbindungstechniken werden als Instanzen dieses Rahmenwerks vorgestellt und implementiert für das deduktive Verifikationssystem KeY. Wir zeigen, dass unsere Operationen für gewisse Beispiele tatsächlich die Zahl der Zustände und Abzweigungen signifikant reduzieren, und wenden eine Verbindungstechnik auf den Bereich der Informationsflussanalyse an, um zu demonstrieren, dass die Zusammenführung von Zuständen prinzipiell dazu in der Lage ist, die Präzision von Analyseverfahren in diesem Bereich zu erhöhen.

Contents

1	Introduction	4
1.1	Motivation	4
1.2	Outline	5
2	Preliminaries	6
2.1	Basic Definitions	6
2.2	Dynamic Logic	6
2.2.1	Java DL Syntax	7
2.2.2	Java DL Semantics	9
2.3	Symbolic Execution	10
3	A Lattice Model for Symbolic Execution	13
3.1	Concretization and Weakening	13
3.2	The General Lattice Model	18
3.3	Constructing Join Nodes	21
3.3.1	Full Anonymization	21
3.3.2	If-Then-Else	24
3.3.3	If-Then-Else by additional Path Condition Constraints	25
3.3.4	Abstract Weakening	27
3.3.5	Heap Treatment	29
4	Implementation	31
4.1	JoinRule and JoinRuleUtils	31
4.2	The Abstraction Framework	33
4.3	The Partner Goals Closing Rule	33
4.4	Macro for Execution until Join Points	34
4.5	Sample Join Rule Implementations	34
5	Evaluation and Case Study	37
5.1	Performance Evaluation	37
5.2	Case Study: Information Flow Analysis	39
6	Related Work and Conclusion	45
6.1	Related Work	45
6.2	Conclusion and Future Work	46
	Large Figures	48
	Lists of Acronyms, Figures, Tables, Algorithms, Listings	56
	Bibliography	60
	Further Reading	62

1 Introduction

1.1 Motivation

Symbolic Execution [Bur74; Kin76] is a method to systematically explore all execution paths in a program for all possible input values. In contrast to concrete execution, symbolic execution treats input values as symbols. Whenever the execution depends on the unknown concrete value of a program variable, it splits into subbranches that are thereupon followed independently (\rightarrow Figure 1.1). The result is a *symbolic execution tree*, consisting of *symbolic execution states*, which resembles the unrolled control flow graph. Since its inception in the 1970s, symbolic execution has been employed in two fundamentally different scenarios: (i) The *state exploration* for the purpose of, for instance, test case generation or debugging [BEL75; Kin76; God12; JMN13; CS13], and (ii) the *formal verification* of programs against functional properties [Bur74; DE82; BHS07]. The strength of symbolic execution is its precision. However, there are some drawbacks: First, “classic” symbolic execution is not capable of, for instance, fixpoint iteration for unbounded loops, and relies on repeated loop unwinding. Extensions allow for the manual specification of loop invariants to facilitate the termination of such executions. Thus, symbolic execution techniques usually lack full automation. Second, the splits of the tree at branch points where the execution depends on concrete values cause an up to exponential increase of the tree size (“*path explosion problem*”) (see, for example, [CS13]). Existing approaches in literature addressing the path explosion problem in a debugging / testing context often use subsumption techniques to stop execution of redundant paths [APV06; BCE08; Jaf+12; JMN13; CJM14] or employ guided search strategies for finding good test cases faster [BS08; Xie+09]. In a verification context, these techniques are not applicable since the complete symbolic execution tree has to be considered to prove the desired properties.

Abstract Interpretation [CC77; Cou01] is a static analysis method in which concrete values are abstracted to suitable values of a chosen abstract domain. The analysis follows the control flow of a program; in particular, and in contrast to symbolic execution, branches are merged at places where the control flow converges. Thereby, abstract values are joined according to the join operation of the abstract domain. Systems based on abstract interpretation can achieve full automation [Cou+05], in particular because of their capability of performing a fixpoint iteration for (unbounded) loops. The abstraction, on the other hand, induces a natural loss of precision. A wrong choice of the abstract domain can render it impossible to prove certain kinds of properties.

An obvious idea for tackling the path explosion problem in a verification context is to take up ideas of abstract interpretation by joining branches of symbolic execution trees at suitable points. In the course of this, the underlying data structure for symbolic execution is transformed from a tree to a *Directed Acyclic Graph (DAG)*. Several approaches [HSS09; Kuz+12; Sen+14] realize this by using if-then-else constructs to merge symbolic execution states. All these approaches preserve the precision of symbolic

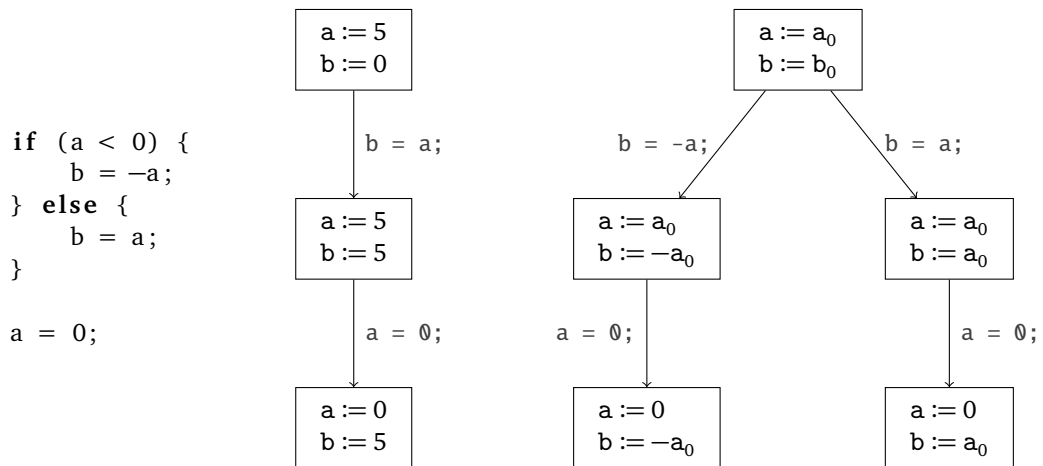


Figure 1.1: Concrete Execution vs. Symbolic Execution

execution and propose different kinds of optimizations. However, they are inflexible regarding their restriction to a fixed kind of join technique, considering that standard abstract interpretation systems allow for an arbitrary selection of abstract domains.

KeY [BHS07] is a deductive verification system for first-order Java Dynamic Logic. During the verification process, Java programs are executed symbolically by the means of special calculus rules, resulting in so-called *updates* representing the effects of the analyzed program. Afterward, properties about the program can be proven by first-order reasoning (augmented by the treatment of updates) with a sequent calculus. The proof procedure works in a semi-automatic fashion: Powerful automatic strategies in KeY significantly reduce the amount of interaction, which might though be required for instance in the case of difficult existential quantifiers or for the specification of loop invariants. The system, in its role as a symbolic execution engine, is the foundation of further applications like symbolic debuggers [Häh+10] and visualizers [HHB14].

In this thesis, we propose a novel framework for embedding join operations into symbolic execution in the context of software verification. Our goal is to contribute to solving the path explosion problem and, by making symbolic execution more flexible, to open it up to new kinds of program analyses built thereupon.

1.2 Outline

Chapter 2 contains preliminaries for the remainder of the thesis. In Chapter 3, we define the notion of *concretizations* of symbolic execution states and a partial order relation called *weakening* between symbolic execution states. In the tradition of [CC77], we base our framework upon *lattice structures* induced by join operations. Besides basic lattice properties, the join operations are required to satisfy two additional correctness properties. We define a join rule for operations conforming with the lattice framework and formally prove a corresponding soundness theorem. Furthermore, we specify concrete join operations based on, e.g., if-then-else constructs and lattice-based abstraction in the sense of abstract interpretation, and show that these operations conform with our framework. We implemented the presented operations in the KeY system; the implementation is outlined in Chapter 4. In Chapter 5, we report the results of a first evaluation of our implementation. Our experiments show that joining branches in KeY proof trees reduces the number of nodes and branches in the tree for several example programs, hence our approach indeed constitutes a step towards solving the path explosion problem. A small application to information flow analysis suggests that branch joining might be able to improve the precision of information flow analyses based on symbolic execution, or could serve as a basis for new analyses. Chapter 6 provides a comparison to related work as well as an outlook on possible future improvements and extensions of our system.

2 Preliminaries

This chapter introduces basic mathematical notions employed in the thesis, selected important concepts concerning the syntax and semantics of Java Dynamic Logic (Java DL), which is the logic that our formalisms are based on, and some fundamentals on Symbolic Execution (SE).

2.1 Basic Definitions

Definition 2.1 (Power Set and Star Operation). For a set A , the power set (i.e., set of all subsets) 2^A is defined as $2^A := \{S : S \subseteq A\}$. We denote by A^n the set of tuples in $\underbrace{A \times A \times \cdots \times A}_{n \text{ times}}$, where \times is the Cartesian set product, by A^* the set $\bigcup_{n \geq 0} A^n$, and by A^+ the set $\bigcup_{n \geq 1} A^n$. \diamond

Definition 2.2 (Projection of Tuples). For sets A_1, \dots, A_n , let $\bar{a} = (a_1, \dots, a_n) \in A_1 \times \cdots \times A_n$ be a tuple of length n . Then we denote by $\text{proj}_i(\bar{a})$ the i -th projection of \bar{a} , i.e. $\text{proj}_i(\bar{a}) := a_i$. For sets of tuples $A \subseteq A_1 \times \cdots \times A_n$, we define $\text{proj}_i(A) := \{\text{proj}_i(\bar{a}) : \bar{a} \in A\}$. \diamond

Lattices and Semilattices

The concept of a *lattice* can be defined in two seemingly different ways: (i) as a partially ordered set (“poset”) with special properties, namely the *existence* of unique least upper and greatest lower bounds, and (2) as a structure / algebra with operations \sqcup (join) and \sqcap (meet) for *computing* least upper and greatest lower bounds. As shown in [Grä78, Theorem 1], these definitions are actually equivalent. From a poset lattice, one can construct an equivalent algebra lattice by defining $a \sqcap b := \inf\{a, b\}$ and $a \sqcup b := \sup\{a, b\}$; from an algebra lattice, one can construct a poset lattice by defining $a \leq b$ iff $a \sqcup b = b$ (or equivalently, $a \leq b$ iff $a \sqcap b = a$). The restriction of an algebra lattice to only one of the operations \sqcap and \sqcup yields the notion of a *semilattice*, that is a join-semilattice or a meet-semilattice, depending on the included operation. We subsequently provide a definition of semilattices as structures following [Grä78].

Definition 2.3 (Semilattice). A *semilattice* (A, \circ) consists of a non-empty set A and one binary operation \circ , such that the properties (L1), (L2) and (L3) are satisfied for $a, b, c \in A$:

- (L1) *Idempotency*: $a \circ a = a$.
- (L2) *Commutativity*: $a \circ b = b \circ a$
- (L3) *Associativity*: $(a \circ b) \circ c = a \circ (b \circ c)$ \diamond

Control Flow Graphs

A Control Flow Graph (CFG) [All70] is a directed graph capturing the control flow relationships in a program. The CFG serves as a basis for many kinds of program analyses like abstract interpretation [CC77]. It consists of a unique entry node START and other nodes called *basic blocks* comprising a linear sequence of program instructions. The edges of the graph correspond to the control flow of the underlying program. Edges arising from conditional branchings in the control flow are labeled with the corresponding branch condition.

Example 2.4. Figure 2.1 shows the CFG for the simple Java program of Listing 2.1. The boxes are basic blocks, the last block $y = 0$; is a “program terminating block”. \diamond

2.2 Dynamic Logic

Java Dynamic Logic (Java DL) extends (typed) first-order logic by including Java programs as well as syntactic elements capturing state changes in the language. For programs, there exist two *modalities*: $\langle p \rangle \varphi$ expresses that the program p terminates and afterward the formula φ holds, whereas $[p] \varphi$ expresses the weaker condition that *if* p terminates, the formula φ holds afterward. State changes can be declared by so-called (syntactic) *updates* which roughly correspond to substitution functions. Subsequently, we introduce syntactic and semantic concepts of Java DL which are used in this thesis. For all notions in this section that are mentioned but not explicitly defined, we refer to [Ben11] (the main source for the fundamental definitions) and [BHS07, Chapter 3].


```

1 x = 0;
2 while (z < y) {
3     z = z + y;
4     x++;
5 }
6 y = 0;

```

Listing 2.1: Example Java program




Figure 2.1: CFG for Listing 2.1.

2.2.1 Java DL Syntax

Subsequently we provide the definitions for the syntactic categories of signatures, terms, formulae, and updates of Java DL. All definitions originate from [Ben11], but may be slightly simplified and shortened. For the complete definitions, please consult [Ben11].

A Java DL signature is basically a signature of typed first-order logic with subtyping, equipped with a Java program Prg . As in first-order logic, the language of Java DL is parametric in a signature which defines the primitives from which terms and formulae may be built.

Definition 2.5 (Java DL Signature). A *signature* is a tuple $\Sigma = (\mathcal{T}, \preceq, PV, LgV, Func, Pred, \alpha, Prg)$ consisting of (i) a finite set of *types* \mathcal{T} such that *Any*, *Boolean*, *Int*, *Null*, *LocSet*, *Field*, *Heap*, *Object* $\in \mathcal{T}$, also containing all reference types of Prg , (ii) a partial order $\preceq \subseteq \mathcal{T} \times \mathcal{T}$ on \mathcal{T} , called the *subtype relation*, as depicted in Figure 2.2, (iii) an infinite set LgV of *logical variables*, (iv) an infinite set PV of *program variables* such that all local variables a of type T in Prg also appear as $a \in PV$ with type A , where $A = T$ if T is a reference type, $A = Boolean$ if $T = boolean$, and $A = Int$ if $T \in \{byte, short, int\}$, (v) an infinite set $Func$ of *function symbols*, (vi) an infinite set $Pred$ of *predicate symbols*, (vii) a static *typing function* α such that $\alpha(v) \in \mathcal{T}$ for $v \in PV \cup LgV$, $\alpha(f) \in \mathcal{T}^* \times \mathcal{T}$ for $f \in Func$, and $\alpha(p) \in \mathcal{T}^+$ for $p \in Pred$, and (viii) a program Prg in the intersection between Java and Java Card, i.e. a set of Java classes and interfaces.

We require that the following symbols are present in every signature for each type $A \in \mathcal{T}$: $heap \in PV$, $cast_A \in Func$, $select_A \in Func$, $store \in Func$, $create \in Func$ and $created \in Func$, with $\alpha(cast_A) = (Any, A)$, $\alpha(select_A) = ((Heap, Object, Field), A)$, $\alpha(store) = ((Heap, Object, Field, Any), Heap)$, $\alpha(create) = ((Heap, Object), Heap)$, $\alpha(created) = Field$. \diamond

The distinction between rigid and non-rigid predicate and function symbols prevailing in [BHS07] is dropped. In this framework, rigid symbols had the same interpretation in all states, whereas non-rigid symbols could change their meaning between state transitions. Since general non-rigid functions and predicates became obsolete in the more recent framework of [Ben11], program variables are now contained in the dedicated set PV as the henceforth only non-rigid function symbols.

One of the main results of [Ben11] is the integration of an explicit model of the Java heap as the value of a special program variable $heap$ into Java DL. A heap structure is logically represented as a term consisting of nested $store$ expressions; for instance, the update $heap := store(heap, o, f, 2)$ changes the initial heap by setting the field $o.f$ to 2. The result is again a heap expression that can be the input of further store expressions.

In the following, if not otherwise specified, we assume an underlying signature $\Sigma = (\mathcal{T}, \preceq, PV, LgV, Func, Pred, \alpha, Prg)$ for the program of interest Prg as given.




Figure 2.2: Structure of Java DL type hierarchies. Source: [Ben11]

From now on, we expect all Java / Java Card programs in Java DL to be “legal” fragments. Instead of giving a complete formal definition of legal syntax and semantics for Java programs, which would go beyond the scope of this thesis, we refer to the *Java Language Specification* [Gos+05]. For the complete definition, see [Ben11, Definition 5.2].

Subsequently, we define the syntax of Java DL terms, formulae and updates by a BNF specification.

Definition 2.6 (Java DL Syntax). The sets Terms_Σ^A of *terms* of type A , Form_Σ of *formulae* and Upd_Σ of *updates* are defined by the following grammar:

$$\begin{aligned}
\text{Terms}_\Sigma^A &::= x \mid \mathbf{a} \mid f \left(\text{Terms}_\Sigma^{B'_1}, \dots, \text{Terms}_\Sigma^{B'_n} \right) \mid \\
&\quad \text{if } (\text{Form}_\Sigma) \text{ then } (\text{Terms}_\Sigma^A) \text{ else } (\text{Terms}_\Sigma^A) \mid \{ \text{Upd}_\Sigma \} \text{Terms}_\Sigma^A \\
\text{Form}_\Sigma &::= \text{true} \mid \text{false} \mid p \left(\text{Terms}_\Sigma^{B'_1}, \dots, \text{Terms}_\Sigma^{B'_n} \right) \mid \neg \text{Form}_\Sigma \mid \text{Form}_\Sigma \wedge \text{Form}_\Sigma \mid \\
&\quad \text{Form}_\Sigma \vee \text{Form}_\Sigma \mid \text{Form}_\Sigma \rightarrow \text{Form}_\Sigma \mid \text{Form}_\Sigma \leftrightarrow \text{Form}_\Sigma \mid \\
&\quad \forall Ax; \text{Form}_\Sigma \mid \exists Ax; \text{Form}_\Sigma \mid [p] \text{Form}_\Sigma \mid \langle p \rangle \text{Form}_\Sigma \mid \{ \text{Upd}_\Sigma \} \text{Form}_\Sigma \\
\text{Upd}_\Sigma &::= \text{skip} \mid \mathbf{a} := \text{Terms}_\Sigma^A \mid \text{Upd}_\Sigma \parallel \text{Upd}_\Sigma \mid \{ \text{Upd}_\Sigma \} \text{Upd}_\Sigma
\end{aligned}$$

for any variable $x \in \text{LgV}$ s.th. $\alpha(x) = A$, any program variable $\mathbf{a} \in \text{PV}$ s.th. $\alpha(\mathbf{a}) = A$, any function symbol $f \in \text{Func}$ s.th. $\alpha(f) = ((B_1, \dots, B_n), A)$ and predicate symbol $p \in \text{Pred}$ s.th. $\alpha(p) = (B_1, \dots, B_n)$, where $B'_1 \preceq B_1, \dots, B'_n \preceq B_n$, any legal program fragment p in the context of Prg , and any type $A' \in \mathcal{T}$ with $A' \preceq A$. The set Terms_Σ of (arbitrarily typed) terms is defined as $\text{Terms}_\Sigma := \bigcup_{A \in \mathcal{T}} \text{Terms}_\Sigma^A$. As usual, we call a Java DL term, formula and update *closed* if it contains no free (unbound) logic variables. \diamond

The terms of Java DL are similar to terms of first-order logic, except for program variables (“non-rigid constant symbols”) and the if-then-else constructs, which are additions. Informally, a term $\text{if}(\varphi) \text{ then } (t_1) \text{ else } (t_2)$ evaluates to t_1 if φ holds and to t_2 otherwise. A further peculiarity of Java DL is the concept of *updates* that are employed to syntactically represent the effect of terminating program executions. The KeY system (symbolically) executes the statements of given Java programs and records the effect in an update, until the end of the program is reached (\rightarrow Section 2.3). Intuitively, an *elementary* update $\mathbf{a} := t$ assigns the value of the term t to the program variable \mathbf{a} ; a *parallel* update $U_1 \parallel U_2$ executes the updates U_1 and U_2 in parallel. Curly braces transform updates into update applications. Example 2.11 demonstrates the evaluation of updates.

Subsequently, we introduce a generalized notion of substitutions that allows, besides the usual substitution of free variables, also the substitution of constant symbols. Following [Rüm03], this concept can be characterized as *nullary f -substitution*. We need this extended definition to allow for the substitution of Skolem constants by quantifiable variables: In our join methods in Chapter 3, we sometimes introduce fresh Skolem constants in the spirit of a universally quantified variable; for faithfully using those as preconditions in sequents, we need to quantify over them.

Definition 2.7 (Substitution). Let $\varphi \in \text{Form}_\Sigma$ be a Java DL formula. By $\varphi[t'/t]$, where $t, t' \in \text{Terms}_\Sigma^A$ for any type A , and t is *either a logic variable or a constant*, we denote the formula resulting from a substitution of t by t' in φ . If $t \in \text{LgV}$ is a logic variable, we only substitute unbound instances of t , i.e. those that are not in the scope of an existential or universal quantifier for that variable. \diamond

For convenience, we also introduce the following tuple notation.

Notation 2.8 (Tuple Notation). We abbreviate tuples of variables or constants (x_1, x_2, \dots, x_n) by \bar{x} and call n the *length* of the tuple. By $\forall/\exists\bar{x}\varphi$ we understand $\forall/\exists x_1; \dots \forall/\exists x_n; \varphi$; the notation $\varphi[\bar{t}/\bar{t}']$ denotes $\varphi[t_1/t'_1][\dots/\dots][t_n/t'_n]$. \diamond

2.2.2 Java DL Semantics

Java DL syntax elements are interpreted by Kripke structures which, simply speaking, allow for constructing transition systems with first-order models as vertices. A given Kripke structure assigns the same meaning to all function or predicate symbols; however, it may differ in the interpretation of the program variables. Whenever the value of a program variable is changed, the Kripke structure proceeds to another *state*. Together with variable assignments, Kripke structures interpret arbitrary Java DL formulae (of suitable signatures) of our language. The following definitions again originate from [Ben11]; some of those only occur in a shortened, simplified form subsequently.

Definition 2.9 (Java DL Kripke Structure). A *Java DL Kripke Structure* $K_{\Sigma_p} = (D, \delta, I, S, \rho)$ consists of (i) a set D of semantical values, called the *domain*, (ii) a dynamic typing function $\delta : D \rightarrow \mathcal{T}$, which gives rise to the *subdomains* $D^A = \{x \in D : \delta(x) \preceq A\}$ for all types $A \in \mathcal{T}$, (iii) an *interpretation function* I mapping every function symbol $f \in \text{Func}$ with $\alpha(f) = (A_1, \dots, A_n, A)$ to a function $I(f) : D^{A_1}, \dots, D^{A_n} \rightarrow D^A$ and every predicate symbol $p \in \text{Pred}$ with $\alpha(p) = (A_1, \dots, A_n)$ to a relation $I(p) \subseteq D^{A_1} \times \dots \times D^{A_n}$, (iv) a set S of *states*, which are functions $\sigma \in S$ mapping every program variable $a \in \text{PV}$ with $\alpha(a) = A$ to a value $\sigma(a) \in D^A$, and (v) a function ρ that associates with every program fragment p a transition relation $\rho(p) \subseteq S^2$ s.th. $(\sigma_1, \sigma_2) \in \rho(p)$ iff p , when started in σ_1 , terminates normally (i.e., not by throwing an exception) in σ_2 [Gos+05]. We consider Java programs to be deterministic, so for all program fragments p and all $\sigma_1 \in S$, there is at most one σ_2 s.th. $(\sigma_1, \sigma_2) \in \rho(p)$. We require that every Kripke structure satisfies the following:

- S is the set of *all* functions mapping program variables to properly typed values (it is therefore completely determined by D and δ)
- $D^{\text{Boolean}} = \{tt, ff\}$, $D^{\text{Int}} = \mathbb{Z}$, $D^{\text{Null}} = \{\text{null}\}$, $D^{\text{Heap}} = D^{\text{Object}} \times D^{\text{Field}} \rightarrow D^{\text{Any}}$
- $I(\text{true}) = tt$, $I(\text{false}) = ff$
- $I(\text{cast}_A)(x) = \begin{cases} x & \text{if } x \in D^A \\ \text{null} & \text{if } x \notin D^A \text{ and } A \preceq \text{Object} \\ \emptyset & \text{if } x \notin D^A \text{ and } A = \text{LocSet} \\ ff & \text{if } x \notin D^A \text{ and } A = \text{Boolean} \end{cases}$
- $I(\text{select}_A)(h, o, f) = I(\text{cast}_A)(h(o, f))$ for all $h \in D^{\text{Heap}}$, $o \in D^{\text{Object}}$, $f \in D^{\text{Field}}$
- $I(\text{store})(h, o, f, x)(o', f') = \begin{cases} x & \text{if } o = o', f = f' \text{ and } f \neq I(\text{created}) \\ h(o', f') & \text{otherwise} \end{cases}$
for all $h \in D^{\text{Heap}}$, $o, o' \in D^{\text{Object}}$, $f, f' \in D^{\text{Field}}$, $d \in D^{\text{Any}}$
- $I(\text{create})(h, o)(o', f) = \begin{cases} tt & \text{if } o = o', o \neq \text{null} \text{ and } f = I(\text{created}) \\ h(o', f) & \text{otherwise} \end{cases}$
for all $h \in D^{\text{Heap}}$, $o, o' \in D^{\text{Object}}$, $f \in D^{\text{Field}}$. \diamond

For the same reasons why we did not formalize the syntactical correctness of Java programs, we also omit a definition of the semantics of Java programs. Instead, a “black box” function ρ is utilized to capture the behavior of legal program fragments p . The symbolic execution rules of the Java DL calculus provide a formalization of the Java semantics. Listing all those rules would go beyond the scope of this thesis.

Kripke structures allow for the definition of the semantics of Java DL terms, formulae and updates based on a valuation function val .

Definition 2.10 (Java DL Semantics). Given a Kripke structure $K_{\Sigma} = (D, \delta, I, S, \rho)$, a state $\sigma \in S$ and a *variable assignment* $\beta : \text{LgV} \rightarrow D$ (where for $x \in \text{LgV}$ with $\alpha(x) = A$ we have $\beta(x) \in D^A$), we evaluate every term $t \in \text{Terms}_{\Sigma}^A$ to a value $val_{(K_{\Sigma}, \sigma, \beta)}(t) \in D^A$, every formula $\varphi \in \text{Form}_{\Sigma}$ to a truth value $val_{(K_{\Sigma}, \sigma, \beta)}(\varphi) \in \{tt, ff\}$, and every update $u \in \text{Upd}_{\Sigma}$ to a state transformer $val_{(K_{\Sigma}, \sigma, \beta)}(u) : S \rightarrow S$ as defined in [Ben11, Figure 5.2].

We write $(K_{\Sigma}, \sigma, \beta) \models \varphi$ for $val_{(K_{\Sigma}, \sigma, \beta)}(\varphi) = tt$. A formula $\varphi \in \text{Form}_{\Sigma}$ is called *logically valid*, in symbols $\models \varphi$, iff $(K_{\Sigma}, \sigma, \beta) \models \varphi$ for all Kripke structures K_{Σ} , all states $\sigma \in S$, and all variable assignments β . Furthermore, we write $(K_{\Sigma}, \sigma) \models \varphi$ if $(K_{\Sigma}, \sigma, \beta) \models \varphi$ holds for all variable assignments

β ; in particular, we write $(K_\Sigma, \sigma) \models \varphi$ for *closed* formulae φ . For closed terms t without program variables we write $val_{K_\Sigma}(t)$. \diamond

The following example illustrating the semantics of Java DL updates originates from [RRR13].

Example 2.11 (Update Semantics). Consider the formula $\{i := j + 1\} i \geq j$. Evaluating $\{i := j + 1\} i \geq j$ in a state σ is identical to evaluating the subformula $i \geq j$ in a state σ' which coincides with σ except for the value of i that is evaluated to the value of $val_{(K_\Sigma, \sigma, \beta)}(j + 1)$. Evaluation of the parallel update $i := j \parallel j := i$ in a state σ leads to the successor state σ' that is identical to σ except that the values of i and j are swapped. The parallel update $i := 3 \parallel i := 4$ has a *conflict* as i is assigned different values. In such a case the last occurring assignment $i := 4$ overrides all previous ones of the same location variable. Evaluation of the formula $\{i := j\} \{j := i\} \varphi$ in a state σ results in evaluating φ in a state where i has the value of j in σ , and j remains unchanged. The update *skip*, the empty update, does not change the interpreting state. \diamond

Definition 2.12 (Update Normal Form). An update is in *normal form* if it has the shape $U_1 \parallel \dots \parallel U_n$, $n \geq 0$, where each U_i is an elementary update and there is no conflict between U_i and U_j for any $i \neq j$. \diamond

2.3 Symbolic Execution

Symbolic Execution (SE), in contrast to concrete execution, treats program variables, in particular program inputs, as symbols – as long as they are not assigned concrete values. Whenever the execution depends on the concrete, but unknown, value of a variable (in an if statement, for instance), execution splits into subbranches. Thus, the result of the symbolic execution of a program is a Symbolic Execution Tree (SET) in which each node represents a symbolic execution state. An SET resembles an “unrolled” CFG and may in principle, for instance in the presence of loops, be infinite. SE states track changes made to program locations in course of the execution (the *symbolic state*), as well as the constraints on (symbolic) values that lead to the execution of the current path (the *path condition*). Those notions are defined subsequently, following [Kin76; Häh+10; JH14; HHB14].

Definition 2.13 (Symbolic Execution State). A *Symbolic Execution State* is a triple (U, C, φ) consisting of (1) the *symbolic state* $U \in \text{Upd}_\Sigma$, an update in normal form with only closed terms as right sides, tracking changes made to program variables, (2) a set of *closed* Java DL formulae $C \in 2^{\text{Form}_\Sigma}$ encoding the current *path condition*, and (3) a Java DL formula $\varphi \in \text{Form}_\Sigma$ usually containing a modality, which we call the *program counter*, representing the Java code that remains to be executed after that state. \diamond

We denote the (underspecified) set of all symbolic execution states for a given program Prg by $SEStates_{Prg}$. Symbolic execution of a program yields an SET consisting of SE states in $SEStates_{Prg}$. Complete symbolic execution trees for a program p with desired post condition φ are finite acyclic trees whose root is labeled with the node $(U_0, C_0, [p] \varphi)$ or $(U_0, C_0, \langle p \rangle \varphi)$, and whose leaves only contain the empty program counter (i.e., the formula *true*). Every child node is generated from its parent according to the semantics of the symbolically executed program statement.

Example 2.14. Figure 2.3 shows the partial Symbolic Execution Tree for the Java program in Listing 2.1. Program counters are abstracted as line number pointing to the next statement to execute, where the special number -1 refers to the end of the program. The tree is infinite, since it is unknown whether the initial value of z is smaller than that of y , and, e.g., whether z is initially greater or equal than 0 (in this case, $z + y < y$ would be false). An alternative to the unwinding steps used in the tree are loop invariants, which would make the tree finite. \diamond

Symbolic Execution in KeY

The KeY theorem prover is based upon a *sequent calculus*. A sequent is a pair of sets of formulae $\Gamma, \Delta \subseteq 2^{\text{Form}_\Sigma}$, the antecedent and the succedent, of the form $\Gamma \Longrightarrow \Delta$. Its semantics is defined by $\bigwedge_{\varphi \in \Gamma} \varphi \rightarrow \bigvee_{\psi \in \Delta} \psi$. A *sequent calculus rule* has one conclusion and zero or more premises. It is applied to a sequent s by matching its conclusion against s . The instantiated premises are then added as children of s , thus generating a proof tree [RRR13]. The rules in the KeY calculus not concerning symbolic execution correspond to usual rules of a sequent calculus for first-order logic like the calculus LK by Gentzen [Gen64]. In addition to those, the calculus of KeY contains a large set of rules

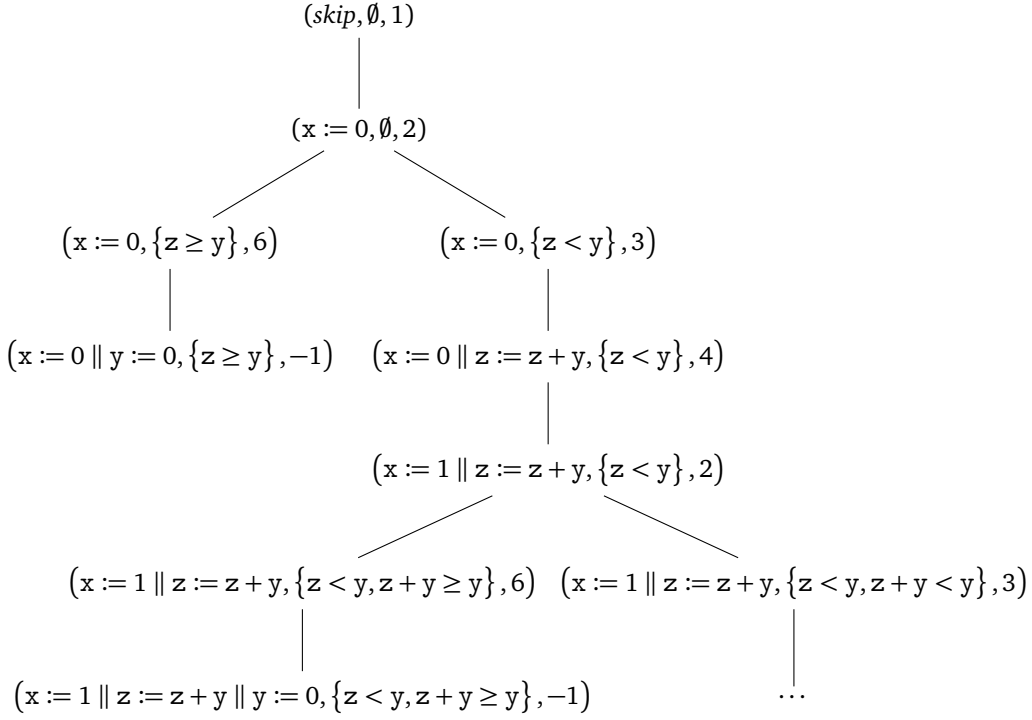


Figure 2.3: Partial Symbolic Execution Tree for Listing 2.1.

dedicated to the treatment of modalities, that is of Java code. Ultimately, KeY transforms modalities into updates, and thus may be seen as a symbolic interpreter of Java code. Figure 2.4 shows some example calculus rules for symbolic execution. A sequent containing at least one modality corresponds to an SE state: Consider the sequent $\Gamma \Longrightarrow \{U\} \varphi, \Delta$. We transform the sequent into the canonical form $\Gamma \cup \{\neg\psi : \psi \in \Delta\} \Longrightarrow \{U\} \varphi$ by shifting all formulae of the succedent except for $\{U\} \varphi$ to the antecedent; the resulting sequent is logically equivalent to the original one. This sequent uniquely corresponds to the SE state $(U, \Gamma \cup \{\neg\psi : \psi \in \Delta\}, \varphi)$. Note that in principle, there could be other formulae in Δ that we could have chosen instead of $\{U\} \varphi$ as a source for the extraction of the symbolic state and program counter; $\neg\{U\} \varphi$ would then become a part of the path condition. Thus, a sequent may be mapped to several different SE states. However, for most cases resulting from correctness proofs of Java programs, the desired mapping can be uniquely determined, since there is only one formula in the succedent which contains a modality.

Notation 2.15. We write $\vdash \varphi$ to express that the sequent $\Longrightarrow \varphi$ is provable within the sequent calculus of KeY. For sets of formulae Δ , $\vdash \Delta$ means that $\Longrightarrow \bigwedge \Delta$ is provable, where

$$\bigwedge \{\varphi_1, \varphi_2, \dots, \varphi_n\} := \varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_n.$$

Equivalently, $\bigvee \Delta$ represents the formula resulting from a disjunction of the contained elements.

Furthermore, we write $\varphi \equiv \psi$ to express that φ is logically equivalent to ψ , i.e. φ is true in a model $(K_\Sigma, \sigma, \beta)$ iff ψ is true in $(K_\Sigma, \sigma, \beta)$. Again, $\Gamma \equiv \Delta$ means $\bigwedge \Gamma \equiv \bigwedge \Delta$ for sets of formulae Γ, Δ . \diamond

$$\begin{array}{c}
\text{assignLocal} \frac{\Gamma \Longrightarrow \{U\} \{a := t\} [\pi \ \omega] \varphi, \Delta}{\Gamma \Longrightarrow \{U\} [\pi \ a = t; \ \omega] \varphi, \Delta} \\
\\
\text{assignField} \frac{\Gamma \Longrightarrow \{U\} \{\text{heap} := \text{store}(\text{heap}, o, f, t)\} [\pi \ \omega] \varphi, \Delta}{\Gamma \Longrightarrow \{U\} [\pi \ o.f = t; \ \omega] \varphi, \Delta} \\
\\
\text{conditional} \frac{\Gamma \Longrightarrow \{U\} \text{if} (exp \doteq true) \text{then} ([\pi \ p_1 \ \omega] \varphi) \text{else} ([\pi \ p_2 \ \omega] \varphi), \Delta}{\Gamma \Longrightarrow \{U\} [\pi \ \text{if} (exp) \ p_1 \ \text{else} \ p_2; \ \omega] \varphi, \Delta} \\
\\
\text{unwindLoop} \frac{\Gamma \Longrightarrow \{U\} [\pi \ \text{if} (exp) \{p'; \ \text{while} (exp) \ p\} \ \omega] \varphi, \Delta}{\Gamma \Longrightarrow \{U\} [\pi \ \text{while} (exp) \ p; \ \omega] \varphi, \Delta} \\
\\
\text{expandMethod} \frac{\begin{array}{c} \Gamma \Longrightarrow \{U\} [\pi \ \text{method-frame}(\text{result} = r, \text{this} = o) : \\ \quad \{ \text{body}(m, A) \} \ \omega] \varphi, \Delta \\ \Gamma \Longrightarrow \{U\} \text{exactInstance}_A(o), \Delta \end{array}}{\Gamma \Longrightarrow \{U\} [\pi \ r = o.m(); \ \omega] \varphi, \Delta} \\
\\
\text{emptyModality} \frac{\Gamma \Longrightarrow \{U\} \varphi, \Delta}{\Gamma \Longrightarrow \{U\} [] \varphi, \Delta}
\end{array}$$

Figure 2.4: Selected rules of the KeY calculus for symbolic execution

3 A Lattice Model for Symbolic Execution

The problem to be solved in this thesis and particularly in this chapter is the merging of two branches in a symbolic execution tree, the last states of which have the same program counter. Figure 3.1 illustrates this situation, Figure 3.2 shows an example in KeY/Java DL syntax for joining two nodes after an if-statement. Our goal is to join SE nodes (U_1, C_1, φ) and (U_2, C_2, φ) with the same program counter φ to a new state (U^*, C^*, φ) that we call *join state*. This gives rise to two orthogonal questions:

- (1) How and when during symbolic execution of a program can we detect suitable branches to join?
- (2) What are the characteristics of sensible instantiations for U^* and C^* , and how can we construct them?

Question (1) addresses the integration of our techniques into the symbolic execution process; it would be desirable to automate the joining of branches such that a user presses “play” whereupon KeY outputs a DAG with suitable branches having been joined. A complete generation of the SET with subsequent pruning and joining steps is undesirable: amongst the disadvantages of this naive approach is the obvious performance overhead. Thus, branch joining should ideally be incorporated into the proof generation process. We refer to this question in Section 4.4. Question (2) concerns the actual joining of two branches, the computation of join states from two parents. In particular, we propose a general lattice framework for symbolic execution, with the property that join techniques conforming with our formal framework preserve the soundness of the KeY calculus. Subsequently, we fix the foundations of our framework by narrowing the gap between Symbolic Execution and Abstract Interpretation.

3.1 Concretization and Weakening

Symbolic Execution can be regarded, at least to some extent, as a case of Abstract Interpretation [CC77]. Each SE state describes a potentially infinite set of *concrete states*; only if all locations are set to concrete values, i.e. do not depend on symbolic input values, is the set of described concrete states a singleton set. However, abstract interpretation demands a complete semilattice with join operation, partial order, least and top element, which is usually not defined for SE states. Subsequently, we define a concretization function from SE states to concrete states, as well as a partial order relation between SE states. In Section 3.3, we furthermore define join operations on SE states, which allows us to stipulate lattice structures also for symbolic execution.

Definition 3.1 (Concrete Execution States). A *concrete execution state* is a pair (σ, φ) consisting of (i) a Kripke state $\sigma : PV \rightarrow D$ s.th. if $\sigma(\mathbf{a}) = x$ and $\alpha(\mathbf{a}) = A$, it holds that $\delta(x) = A$, and (ii) a Java DL formula φ , the program counter, usually containing a modality with the program that remains to be executed p . We denote the set of all concrete execution states for a program Prg by $ConcrStates_{Prg}$. \diamond

We now can introduce a concretization function from SE states to concrete states based on the valuation function val (\rightarrow Definition 2.10).

Definition 3.2 (Concretization Function). Let $s_{SE} = (U, C, \varphi) \in SEStates_{Prg}$. The *concretization function* $concr$ maps s_{SE} to a set of concrete states in $2^{ConcrStates_{Prg}}$ where

$$concr(s_{SE}) := \left\{ (\sigma', \varphi) : \sigma' = val_{(K_\Sigma, \sigma)}(U)(\sigma) \right. \\ \left. \wedge K_\Sigma = (D, \delta, I, S, \rho) \text{ is a Kripke structure } \wedge (K_\Sigma, \sigma) \models C \right\}$$

If the program counter φ is clear from the context, we also write $\sigma' \in concr(s_{SE})$ for $(\sigma', \varphi) \in concr(s_{SE})$. \diamond

For each possible value that a term, that is the right side in the symbolic state for a program variable x , can attain in any Kripke structure under the given constraints C , the concretization as defined above contains an assignment function mapping x to exactly this value. Thus, the set $concr(s_{SE})$ contains exactly the concrete states that are described by the SE state s_{SE} .

Definition 3.2 facilitates the natural definition of a partial order relation between SE states: A *weakening relation* expressing that one state describes more concrete states than another one.

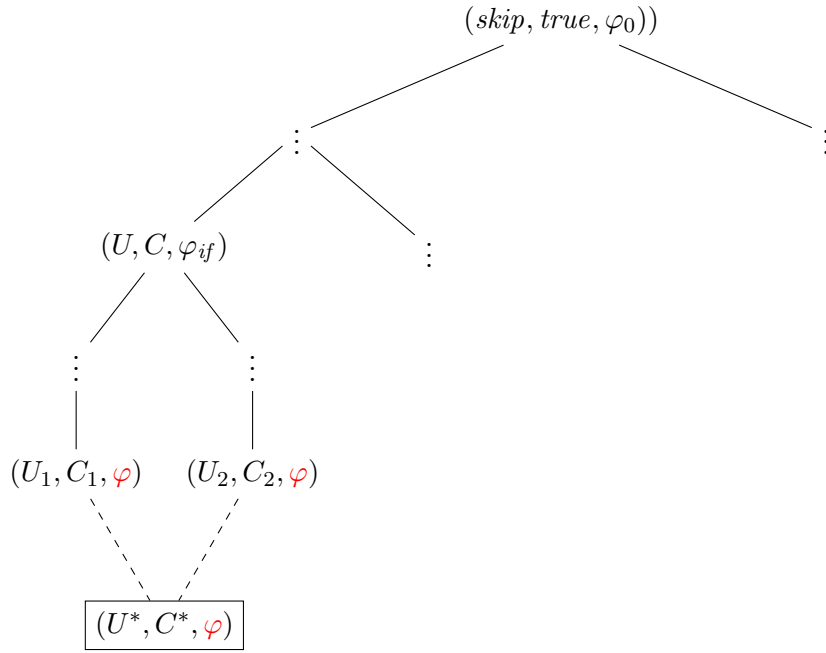


Figure 3.1: Joining two branches in an abstract SET after an if statement.

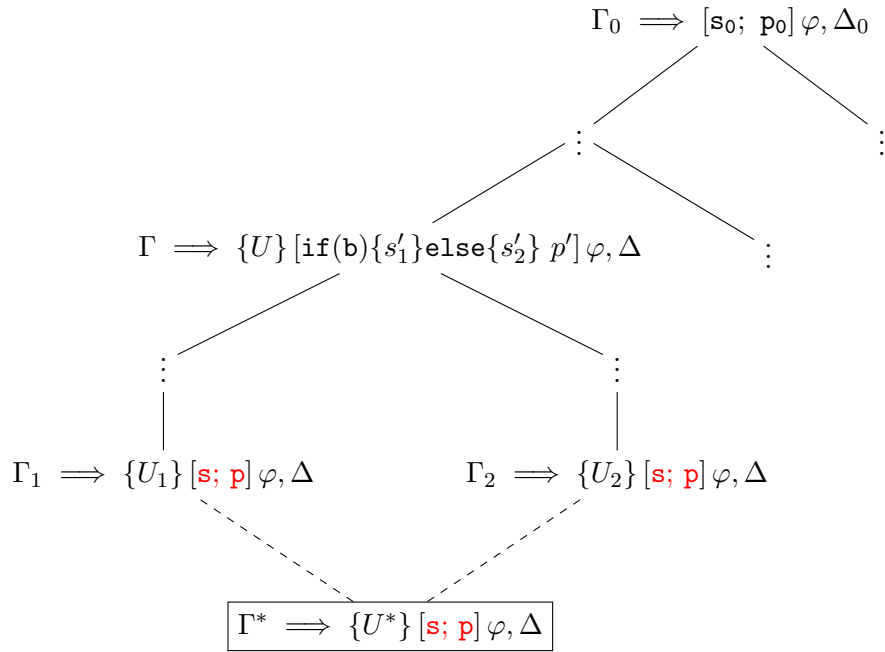


Figure 3.2: Joining two branches in a KeY proof tree after an if statement.

Definition 3.3 (Weakening Relation). Let $s_1, s_2 \in SEStates_{prg}$ be two SE states. We say that s_2 is *weaker* than, or a *weakening* of, s_1 and write $s_1 \lesssim s_2$ if and only if $concr(s_1) \subseteq concr(s_2)$. \diamond

Note that Definition 3.3, along with Definition 3.2, implies that a state s_2 can only be a weakening of a state with *satisfiable path condition* s_1 if they have the same program counters. If the path condition of s_1 is unsatisfiable, $concr(s_1)$ equals the empty set; therefore, any state will be a weakening of a state with unsatisfiable path condition, in particular including states with a different program counter. Usually, we assume that all path conditions are satisfiable. The following example illustrates this concept.

Example 3.4. Consider the SE state $s_{SE} := (x := y, y > 0, \varphi)$. The set of concretizations for s_{SE} is $concr(s) = \{(\sigma, \varphi) : \sigma(z) > 0\}$. One intuitive weakening of s_{SE} is obtained by weakening the constraint on the program variable y such that the value 0 is also covered, resulting in $s'_{SE} = (x := y, y \geq 0, \varphi)$. Indeed it holds for the set of concretizations of s'_{SE} , $concr(s'_{SE}) = \{(\sigma, p) : \sigma(z) \geq 0\}$, that $concr(s'_{SE}) \supseteq concr(s_{SE})$, so s'_{SE} is a weakening of s_{SE} . Note that $s''_{SE} = (x := y, y \geq 0, \varphi')$, for $\varphi \neq \varphi'$, is not weaker than s_{SE} . \diamond

Lemma 3.6 shows that the relation \lesssim is actually a partial order relation. We slightly generalized the antisymmetry condition in the lemma, since strict syntactical equality on SE states is too strong. Assume that the symbolic state of a state s_1 contains the elementary update $x := t$, whereas the symbolic state s_2 contains the elementary update $x := \text{if}(true) \text{ then}(t) \text{ else}(t')$. s_1 and s_2 are obviously syntactically different, whereas all Kripke models map the program variable x to the value of the same term t . Therefore, we use the equality of the *concretizations* of the states, as defined below, rather than the syntactical equality.

Definition 3.5 (Equality of Concretizations). The equivalence relation $\stackrel{concr}{=} \subseteq SEStates_{prg} \times SEStates_{prg}$ is defined by $s_1 \stackrel{concr}{=} s_2 \iff concr(s_1) = concr(s_2)$. \diamond

It is obvious that $\stackrel{concr}{=}$ is an equivalence relation, since it employs the usual equality $=$ on sets in a straightforward manner. Using $\stackrel{concr}{=}$, we can formulate the subsequent lemma.

Lemma 3.6. *The relation $\lesssim \in SEStates_{prg} \times SEStates_{prg}$ is a partial order relation.* \diamond

Proof. We have to show the following properties of \lesssim , for $s, s_1, s_2, s_3 \in SEStates_{prg}$:

- (1) reflexivity: $s \lesssim s$
- (2) antisymmetry: $s_1 \lesssim s_2 \wedge s_2 \lesssim s_1 \rightarrow s_1 \stackrel{concr}{=} s_2$
- (3) transitivity: $s_1 \lesssim s_2 \wedge s_2 \lesssim s_3 \rightarrow s_1 \lesssim s_3$

All properties follow from ordinary set theory. Property (1) follows from $concr(s) = concr(s)$ and, therefore, $concr(s) \subseteq concr(s)$. For (2), assume that $s_1 \lesssim s_2$ and $s_2 \lesssim s_1$, i.e. $concr(s_1) \subseteq concr(s_2)$ and $concr(s_2) \subseteq concr(s_1)$. From that, we obtain $concr(s_1) = concr(s_2)$. For (3), assume that $s_1 \lesssim s_2$ and $s_2 \lesssim s_3$, i.e. $concr(s_1) \subseteq concr(s_2)$ and $concr(s_2) \subseteq concr(s_3)$. Then $concr(s_1) \subseteq concr(s_3)$ follows from the transitivity of the subset relation. \square

Having defined a *semantic* condition for two SE states being in the weakening relation, we now aim for a corresponding *logical representation*. This condition is then employed in the implementation part (\rightarrow Section 4.3) to facilitate automatic proofs of the weakening relation between two SE states with support of the KeY system.

Definition 3.7 (Logical Representation of Weakening). Let $s_1 = (U_1, C_1, \varphi_1) \in SEStates_{prg}$ and $s_2 = (U_2, C_2, \varphi_2) \in SEStates_{prg}$ be two symbolic execution states. We say that s_2 is *logically weaker* than, or a *logical weakening* of, s_1 and write $s_1 \lesssim_{log} s_2$, if the following formula $\varphi_{\lesssim_{log}}$ is *provable* for s_1 and s_2 , where \bar{c} are the new constants introduced in s_2 (and not contained in s_1), \bar{x} are all program variables contained in U_1, U_2, C_1 and C_2 , \bar{v} is a tuple of fresh logical variables of the same length and types as \bar{c} , and P is a new predicate of suitable type:

$$\varphi_{\lesssim_{log}} := \begin{cases} \forall \bar{v}; ((\bigwedge C_2 \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}]) \rightarrow (\bigwedge C_1 \rightarrow \{U_1\} P(\bar{x})) & \varphi_1 = \varphi_2 \\ false & \text{otherwise} \end{cases} \quad \diamond$$

Note that we claim provability instead of the mere truth of the formula $\varphi_{\lesssim_{\log}}$. Since logical weakening is meant to be used in the KeY system, it is important that the formula can actually be proven.

Subsequently, we establish two lemmas as well as a proposition following from those, which provide affirmations about the relation between semantic and logical weakening. As we will see, logical weakening is an equivalent of semantic weakening. The conditions on the path condition of the second SE state s_2 in the lemmas correspond to the property (SEL5) introduced later in Definition 3.12.

Lemma 3.8. *Let $s_1 = (U_1, C_1, \varphi_1)$ and $s_2 = (U_2, C \wedge C_{ax}, \varphi_2) \in SEStates_{\text{prg}}$ be two SE states such that $\vdash \bigwedge C_1 \rightarrow C$ and $\vdash \exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ where \bar{c} is a tuple of Skolem constants introduced in s_2 (and not present in s_1 and C). Then, semantic weakening implies logical weakening, i.e., it holds that*

$$s_1 \lesssim s_2 \implies s_1 \lesssim_{\log} s_2 \quad \diamond$$

Proof. We have to show that $\varphi_{\lesssim_{\log}}$ is provable. Consider the following proof:

$$\frac{\frac{\forall \bar{v}; ((C \wedge C_{ax} \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}]), \bigwedge C_1 \implies \{U_1\} P(\bar{x})}{\forall \bar{v}; ((C \wedge C_{ax} \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}]) \implies \bigwedge C_1 \rightarrow \{U_1\} P(\bar{x})} \rightarrow R}{\implies \forall \bar{v}; ((C \wedge C_{ax} \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}]) \rightarrow (\bigwedge C_1 \rightarrow \{U_1\} P(\bar{x}))} \rightarrow R$$

To close the proof under the assumption that C_1 is satisfiable, we have to find suitable instantiations for the variables \bar{v} . Let $\Sigma_1 \subseteq \Sigma_2$ be two signatures, where Σ_2 results from Σ_1 by adding the constants \bar{c} , and Σ_1 is a suitable signature for s_1 without the constants \bar{c} . Furthermore, let (K_{Σ_1}, σ') be an arbitrary Java DL model satisfying C_1 . Obviously, $\sigma = \text{val}_{(K_{\Sigma_1}, \sigma')}(U_1)(\sigma') \in \text{concr}(s_1)$, and by $s_1 \lesssim s_2$, $\sigma \in \text{concr}(s_2)$. We expand K_{Σ_1} to a structure K_{Σ_2} in the signature Σ_2 by choosing the interpretation of the constants \bar{c} in K_{Σ_2} such that $\text{val}_{(K_{\Sigma_2}, \sigma')}(U_2)(\sigma') = \sigma$ and $(K_{\Sigma_2}, \sigma') \models C_{ax}$. Satisfying the condition $(K_{\Sigma_2}, \sigma') \models C_{ax}$ is possible due to $\vdash \exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$. The condition $\text{val}_{(K_{\Sigma_2}, \sigma')}(U_2)(\sigma') = \sigma$ can be satisfied since otherwise, $s_1 \lesssim s_2$ would not hold in general. Now let $\bar{d} := I_{K_{\Sigma_2}}(\bar{c})$. Then there are closed terms \bar{t} of the signature Σ_1 such that $\text{val}_{K_{\Sigma_2}}(\bar{t}) = \bar{d}$ due to $\text{val}_{(K_{\Sigma_2}, \sigma')}(U_2)(\sigma') = \sigma$. We now continue our proof (we omit $\forall \bar{v}; ((\bigwedge C_2 \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}])$ after the $\exists L$ application as well as $\{U_1\} P(\bar{x})$ in the first and second leaf branch, and implicitly eliminate the conjunction from $\bigwedge C_1$ in the antecedent):

$$\frac{\frac{C_1 \implies C \quad C_1 \implies C_{ax}[\bar{t}/\bar{c}]}{C_1 \implies (C \wedge C_{ax})[\bar{t}/\bar{c}], \{U_1\} P(\bar{x})} \wedge R \quad C_1, (\{U_2\} P(\bar{x}))[\bar{t}/\bar{c}] \implies \{U_1\} P(\bar{x})}{\frac{(C \wedge C_{ax} \rightarrow \{U_2\} P(\bar{x}))[\bar{t}/\bar{c}], \bigwedge C_1 \implies \{U_1\} P(\bar{x})}{\forall \bar{v}; ((C \wedge C_{ax} \rightarrow \{U_2\} P(\bar{x}))[\bar{v}/\bar{c}]), \bigwedge C_1 \implies \{U_1\} P(\bar{x})} \exists L} \rightarrow L} \vdots$$

The branch $C_1 \implies C$ can be closed since $\vdash \bigwedge C_1 \rightarrow C$. Furthermore, we can close the second branch $C_1 \implies C_{ax}[\bar{t}/\bar{c}]$ since, due to our reasoning above, any model satisfying C_1 also satisfies $C_{ax}[\bar{t}/\bar{c}]$ for the chosen terms \bar{t} . For the remaining branch, let us assume that the model (K_{Σ_2}, σ') interprets the predicate P such that $(\{U_2\} P(\bar{x}))[\bar{t}/\bar{c}]$ (otherwise, we could close the branch), i.e.

$$(K_{\Sigma_2}, \text{val}_{(K_{\Sigma_2}, \sigma')}(U_2)(\sigma')) \models P(\bar{x}).$$

Since we chose $\text{val}_{(K_{\Sigma_2}, \sigma')}(U_2)(\sigma') = \sigma$, it holds that

$$\begin{aligned} & (K_{\Sigma_2}, \text{val}_{(K_{\Sigma_1}, \sigma')}(U_1)(\sigma')) \models P(\bar{x}) \\ \iff & (K_{\Sigma_2}, \sigma') \models \{U_1\} P(\bar{x}) \end{aligned}$$

which is furthermore, since U_1 is in the old language Σ_1 , equivalent to $(K_{\Sigma_1}, \sigma') \models \{U_1\} P(\bar{x})$. Therefore, we can also close the third branch of the proof. \square

In the above proof, there are some potential incompleteness issues since we sometimes argue for the provability of a sequent using its truth. At these places, we make use of the *relative* completeness of the KeY calculus. In a complete calculus, all formulae that are *valid* can actually be *proven* within the calculus. Since the logic of KeY supports arithmetic, the calculus is inherently incomplete by Gödel's first incompleteness theorem. However, by the rules for arithmetic implemented in KeY, it is possible to prove almost any arithmetical statement that occurs *in practice* [BHS07].

Lemma 3.9. *Let $s_1 = (U_1, C_1, \varphi_1)$ and $s_2 = (U_2, C \wedge C_{ax}, \varphi_2) \in SEStates_{prg}$ be two SE states such that $\vdash \bigwedge C_1 \rightarrow C$ and $\vdash \exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ where \bar{c} is a tuple of Skolem constants introduced in s_2 (and not present in s_1 and C). Then, logical weakening implies semantic weakening, i.e. it holds that*

$$s_1 \lesssim_{log} s_2 \implies s_1 \lesssim s_2. \quad \diamond$$

Proof. We assume $s_1 \lesssim_{log} s_2$, i.e., the formula $\varphi_{\lesssim_{log}}$ is provable for s_1, s_2 according to Definition 3.7. Let $\sigma \in concr(s_1)$; we have to show that $\sigma \in concr(s_2)$. Since $\sigma \in concr(s_1)$, there is a model $(K_{\Sigma}, \sigma') \models C_1$. Furthermore, $\{U_1\}P(\bar{x})$ is satisfiable, but not valid. Any proof of $\varphi_{\lesssim_{log}}$ must therefore find instantiations \bar{t} for \bar{v} such that, after an application of the implication left rule, the proof

$$\rightarrow_L \frac{C_1 \implies (\bigwedge C_2)[\bar{t}/\bar{c}], \{U_1\}P(\bar{x}) \quad C_1, (\{U_2\}P(\bar{x}))[\bar{t}/\bar{c}] \implies \{U_1\}P(\bar{x})}{C_1, (\bigwedge C_2)[\bar{t}/\bar{c}] \rightarrow (\{U_2\}P(\bar{x}))[\bar{t}/\bar{c}] \implies \{U_1\}P(\bar{x})}$$

can be closed. Consider the right branch

$$C_1, (\{U_2\}P(\bar{x}))[\bar{t}/\bar{c}] \implies \{U_1\}P(\bar{x})$$

or equivalently

$$(\{U_2\}P(\bar{x}))[\bar{t}/\bar{c}] \implies \bigwedge C_1 \rightarrow \{U_1\}P(\bar{x})$$

Since the succedent is not valid, it must hold that for *any model* satisfying C_1 , there are instantiations \bar{t} for the new constants such that the right sides of U_1 , under the constraints C_1 , are instances of the right sides of U_2 , namely those resulting from the substitution $U_2[\bar{t}/\bar{c}]$. In particular, $\sigma = val_{(K_{\Sigma}, \sigma')}(U_1)(\sigma')$ must also be contained in $concr(s_2)$. \square

Remark 3.10. Consider the left branch $\rightarrow L$ application in Lemma 3.9:

$$C_1 \implies (\bigwedge C_2)[\bar{t}/\bar{c}], \{U_1\}P(\bar{x}).$$

Since C_1 is satisfiable, $\{U_1\}P(\bar{x})$ is not valid, and $\bigwedge C_2$ is equivalent to a formula $C \wedge C_{ax}$ such that $C_1 \rightarrow C$ and $\vdash \exists \bar{v}; (C_{ax}[\bar{t}/\bar{c}])$, the formula $\{U_1\}P(\bar{x})$ is irrelevant for a proof of the sequent and we obtain, since the \bar{c} do not occur in C_1 , the w.r.t. provability equivalent sequent

$$C_1 \implies C \wedge (C_{ax}[\bar{t}/\bar{c}]).$$

Therefore, either $(C_{ax})[\bar{t}/\bar{c}]$ must be provable or C_1 must imply $(C_{ax})[\bar{t}/\bar{c}]$. This characterizes both situations occurring in our join operations introducing new constants: One of those introduces axioms like $c \geq 0$ which are, after the substitution and instantiation, provable without C_1 . The other one constructs C_{ax} as a conjunction of formulae of the form

$$(\bigwedge C_1 \rightarrow c_i \doteq \{U_1\}x_i) \wedge (\bigwedge C' \rightarrow c_i \doteq \{U'\}x_i)$$

that are only provable based on the knowledge of C_1 (and C'). \diamond

The following proposition, an easy conclusion from the Lemmas 3.8 and 3.9, provides an equivalence result for semantic and logical weakening.

Proposition 3.11 (Weak Equivalence of Semantic and Logical Weakening). *Let $s_1 = (U_1, C_1, \varphi_1)$ and $s_2 = (U_2, C \wedge C_{ax}, \varphi_2) \in SEStates_{prg}$ be two SE states such that $\vdash \bigwedge C_1 \rightarrow C$ and $\vdash \exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ where \bar{c} is a tuple of Skolem constants introduced in s_2 (and not present in s_1 and C). Then, semantic and logical weakening coincide, i.e. it holds that*

$$s_1 \lesssim s_2 \iff s_1 \lesssim_{log} s_2. \quad \diamond$$

Proof. The proposition follows immediately from Lemmas 3.8 and 3.9. \square

Proposition 3.11 facilitates the implementation of checks in the KeY system verifying the soundness of the implemented join rules for every application. The concrete details of the implementation are outlined in Section 4.3.

3.2 The General Lattice Model

The core of our formal framework for embedding join operations into symbolic execution is a family of join-semilattices parametric in a join operation. The partial order induced by the join operation is constrained by the semantic weakening relation (\rightarrow Definition 3.3).

Definition 3.12 (Induced Join-Semilattices of SE States). Let $\dot{\cup} : SEStates_{prg} \times SEStates_{prg} \rightarrow SEStates_{prg}$ be an operation on SE states. Then, we call the family of structures

$$\mathfrak{L}_{prg} := \left\{ (S_\varphi, \dot{\cup})_\varphi \mid (U, C, \varphi') \in S_\varphi \iff ((U, C, \varphi') \in SEStates_{prg} \text{ and } \varphi' = \varphi) \right\}$$

the induced family of join-semilattices for $\dot{\cup}$, each structure $(S, \dot{\cup})_\varphi \in \mathfrak{L}_{prg}$ the induced join-semilattice for $\dot{\cup}$ and φ , and the relation

$$\begin{aligned} &\preceq \subseteq SEStates_{prg} \times SEStates_{prg} \text{ where} \\ a \preceq b &\iff \begin{cases} a \dot{\cup} b \stackrel{concr}{=} b \text{ and} \\ a, b \text{ have the same program counter} \end{cases} \end{aligned}$$

the induced partial order relation for $\dot{\cup}$, if the properties (SEL1) to (SEL5) are satisfied for $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U_3, C_3, \varphi) \in SEStates_{prg}$:

- (SEL1) *Idempotency*: $a \dot{\cup} a \stackrel{concr}{=} a$.
- (SEL2) *Commutativity*: $a \dot{\cup} b \stackrel{concr}{=} b \dot{\cup} a$
- (SEL3) *Associativity*: $(a \dot{\cup} b) \dot{\cup} c \stackrel{concr}{=} a \dot{\cup} (b \dot{\cup} c)$
- (SEL4) *Correctness*: $a \preceq b$ implies $a \lesssim b$
- (SEL5) *Conservativity*: $a \preceq b$ implies that C_2 is logically equivalent to a formula $C \wedge C_{ax}$, where $\vdash \bigwedge C_1 \rightarrow C$, C does not contain uninterpreted Skolem constants not occurring in a and, if \bar{c} are all uninterpreted Skolem constants in C_{ax} not contained in a , the formula $\exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ is provable. \diamond

Note that we could replace \lesssim in (SEL4) by \lesssim_{log} by Proposition 3.11. The properties (SEL1) to (SEL3) are obviously desirable for any join operation. If we join two states, we usually do not want lose precision by strictly weakening program variables that evaluate to the same value in both states; this motivates the idempotency property. Likewise, we do not want the order of the join to influence the result, which motivates the commutativity and associativity properties. However, these properties are also standards for the definition of a semilattice (\rightarrow Definition 2.3). Property (SEL4) makes use of the semantic weakening relation (\rightarrow Definition 3.3). We call this property “correctness” since it allows, along with (SEL5), for proving the correctness Theorem 3.16.

The subsequent lemma that is employed in the proof of the correctness theorem concerns the interplay between a join operation and its corresponding induced partial order relation.

Lemma 3.13. *Let $a, b \in SEStates_{prg}$ be two SE states with the same program counter, and let $\dot{\cup}$ be an idempotent and associative join operation with the induced partial order relation \preceq . Then, for any $a, b \in SEStates_{prg}$ with the same program counter, it holds that $a \preceq b$ if and only if there is a state c with the same program counter as a, b such that $b = a \dot{\cup} c$. \diamond*

Proof. The direction “ \implies ” follows from a choice of $c := b$. By definition of \preceq , it holds that $a \preceq b$ iff $a \dot{\cup} b \stackrel{concr}{=} b$. “ \impliedby ”: By the reflexivity of the equivalence relation $\stackrel{concr}{=}$, we have $a \dot{\cup} c \stackrel{concr}{=} a \dot{\cup} c$. Then by idempotency of $\dot{\cup}$, we obtain $(a \dot{\cup} a) \dot{\cup} c \stackrel{concr}{=} a \dot{\cup} c$, and then by associativity $a \dot{\cup} (a \dot{\cup} c) \stackrel{concr}{=} a \dot{\cup} c$. Since $a \dot{\cup} c = b$, we have $a \dot{\cup} b \stackrel{concr}{=} b$, and therefore $a \preceq b$. \square

Note that in the above proof, we made use of the idempotency and associativity properties of the induced join-semilattices. So, besides the intuitive motivation why those properties are desirable, those standard lattice properties also play a vital role for the proof of the central soundness theorem. The following lemma is needed in the proof of the soundness theorem.

Lemma 3.14. Let $\Sigma_1 \subseteq \Sigma_2$ be two signatures, where Σ_2 results from Σ_1 by the addition of new constants \bar{c} . Let $\varphi \in \text{Form}_{\Sigma_1}$ be a Java DL formula that is true in all models of the set

$$\mathcal{M}_1 := \{K_{\Sigma_2} : (K_{\Sigma_2}, s) \models C \wedge C_{ax}\} \times S,$$

where $C \in \text{Form}_{\Sigma_1}$, $C_{ax} \in \text{Form}_{\Sigma_2}$ such that $\vdash \exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ and S is a set of Kripke states in the signature Σ_1 . Then φ is also true in all models in the set

$$\mathcal{M}_2 := \{K_{\Sigma_1} : (K_{\Sigma_1}, s) \models C\} \times S \quad \diamond$$

Proof. Assume that φ does hold in the set \mathcal{M}_1 , but does not hold for all formulae in \mathcal{M}_2 . Let $(K_{\Sigma_1}, \sigma) \in \mathcal{M}_2$ be a model such that $(K_{\Sigma_1}, \sigma) \not\models \varphi$. From a proof of $\exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$, we obtain a tuple \bar{t} as instantiations of \bar{v} for which $C_{ax}[\bar{t}/\bar{c}]$ holds in all models in Σ_1 . Then we can expand K_{Σ_1} to a structure K_{Σ_2} in the signature Σ_2 where the constants \bar{c} are interpreted according to the terms \bar{t} . Since (K_{Σ_2}, σ) still satisfies C , it holds that $(K_{\Sigma_2}, \sigma) \models C \wedge C_{ax}$; since however φ is in the old language of Σ_1 , and therefore its interpretation in (K_{Σ_2}, σ) equals that of (K_{Σ_1}, σ) , we still have $(K_{\Sigma_2}, \sigma) \not\models \varphi$, which is a contradiction to $(K_{\Sigma_2}, \sigma) \in \mathcal{M}_1$. \square

Our proposed join rule, having one premise and more than one conclusion, is a *defocusing rule* in the sense of [CS00; Fin05]. Its application on two sequents therefore effectively renders a proof tree into a DAG and thereby realizes the main goal of this thesis. Subsequently, before introducing the main theorem, we provide a definition for defocusing rules.

Definition 3.15 (Defocusing Rule). Let S, S_1, \dots, S_n be sequents in the language of Java DL. A *defocusing rule* is a rule of the form

$$\frac{S}{\frac{S_1 \quad S_2 \quad \dots \quad S_n}{}}$$

We designate defocusing rules by the double bar separating the premise from the conclusions. \diamond

Theorem 3.16 (Correctness of Joins with Induced Join-Semilattices). Let \mathcal{L}_{prg} be an induced family of join-semilattices for a join operation $\dot{\sqcup}$ with the induced partial order relation \preceq . The (defocusing) join rule

$$\text{join by } \dot{\sqcup} \frac{C^* \implies \{U^*\} \varphi}{\frac{\Gamma_1 \implies \{U_1\} \varphi, \Delta_1 \quad \Gamma_2 \implies \{U_2\} \varphi, \Delta_2}{((1), (2) \text{ hold})}}$$

where

$$(1) (U^*, C^*, \varphi) = s_1 \dot{\sqcup} s_2$$

$$(2) s_1 = (U_1, \Gamma_1 \cup \{\neg\psi : \psi \in \Delta_1\}, \varphi), s_2 = (U_2, \Gamma_2 \cup \{\neg\psi : \psi \in \Delta_2\}, \varphi) \in \text{SEStates}_{\text{prg}}$$

is sound, i.e. if $C^* \implies \{U^*\} \varphi$ is valid, then both $\Gamma_1 \implies \{U_1\} \varphi, \Delta_1$ and $\Gamma_2 \implies \{U_2\} \varphi, \Delta_2$ are valid. \diamond

Proof. Under the assumption of the validity of $C^* \implies \{U^*\} \varphi$, we have to show the validity of $\Gamma_1 \implies \{U_1\} \varphi, \Delta_1$ and $\Gamma_2 \implies \{U_2\} \varphi, \Delta_2$, respectively (we consider from now on only the first conclusion of the rule; the proof for the second one is analogous). From

$$s^* := (U^*, C^*, \varphi) = s_1 \dot{\sqcup} s_2$$

we obtain $s_1 \preceq s^*$ by Lemma 3.13. Therefore, by property (SEL4) of Definition 3.12, we also obtain $s_1 \lesssim s^*$, i.e. $\text{concr}(s_1) \subseteq \text{concr}(s^*)$. Let (K_{Σ}, σ) be an arbitrary model, and let

$$C_1 := \Gamma_1 \cup \{\neg\psi : \psi \in \Delta_1\}.$$

We assume that $(K_{\Sigma}, \sigma) \models C_1$; otherwise, we are done. We therefore need to show that

$$(K_{\Sigma}, \sigma) \models \{U_1\} \varphi.$$

Due to the validity of $C^* \implies \{U^*\} \varphi$, we know that φ holds for all models in

$$\left\{ (K'_{\Sigma'}, \sigma'') : \sigma'' = \text{val}_{(K'_{\Sigma'}, \sigma')} (U^*) (\sigma') \wedge (K'_{\Sigma'}, \sigma') \models C^* \right\} = \\ \{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C^*\} \times \text{proj}_1(\text{concr}(s^*))$$

where $\Sigma' \supseteq \Sigma$ results from Σ by adding the new constants introduced in s^* . Since $\text{concr}(s_1) \subseteq \text{concr}(s^*)$, we particularly know that φ is true in all models contained in the set

$$\{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C^*\} \times \text{proj}_1(\text{concr}(s_1)).$$

Due to (SEL5), C^* is logically equivalent to a formula $C \wedge C_{ax}$, where $C_1 \rightarrow C$. Therefore,

$$\{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C^*\} = \{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C \wedge C_{ax}\} \supseteq \{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C_1 \wedge C_{ax}\},$$

so φ is true in all models of the set

$$\{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C_1 \wedge C_{ax}\} \times \text{proj}_1(\text{concr}(s_1))$$

From Lemma 3.14, we obtain that φ is also modeled by the elements of the seemingly bigger set

$$\{K'_{\Sigma'} : (K'_{\Sigma'}, \sigma') \models C_1\} \times \text{proj}_1(\text{concr}(s_1)) =: \mathcal{M}$$

Obviously, it holds that $(K_{\Sigma}, \text{val}_{(K_{\Sigma}, \sigma)}(U_1)(\sigma)) \in \mathcal{M}$. Therefore,

$$(K_{\Sigma}, \text{val}_{(K_{\Sigma}, \sigma)}(U_1)(\sigma)) \models \varphi$$

which is equivalent to $(K_{\Sigma}, \sigma) \models \{U_1\} \varphi$. □

Remark 3.17 (Considerations on the Join Rule). An integration of the join rule defined in Theorem 3.16 into an existing sequent calculus bears the problem that existing soundness proofs would no longer be valid, considering that they are based on the assumption of rules having only one conclusion. Usually, this only affects the top level part of the soundness proof, since the proofs for individual rules are constructed independently. Still, we circumvent this complication by proposing the following *weakening rule* “weaken by join” with only one conclusion:

$$\text{weaken by join } \frac{C^* \implies \{U^*\} \varphi}{\Gamma_1 \implies \{U_1\} \varphi, \Delta_1}$$

where C^* and U^* are computed including a sequent $\Gamma_2 \implies \{U_2\} \varphi, \Delta_2$ as in Theorem 3.16. The “weaken by join” rule is accordingly also sound. When the rule is applied to two leaves $\Gamma_1 \implies \{U_1\} \varphi, \Delta_1$ and $\Gamma_2 \implies \{U_2\} \varphi, \Delta_2$ of a proof tree, we can append the subtree T following $C^* \implies \{U^*\} \varphi$ in the first branch unchanged to the second branch. Figure 3.3 visualizes this scenario. The effect is similar to an employment of the actual defocusing join rule: We only have to *construct* one proof subtree for the such “joined” branches, although we use this tree at two places. In the implementation of our framework, we actually close the second branch to avoid redundancy (\rightarrow Section 4.3). \diamond

To simplify proofs of property (SEL4) in the subsequent section, we introduce the following simple lemma, which allows us to show property (SEL4) by proving the implication $a \dot{\sqcup} b = c \implies a \lesssim c$.

Lemma 3.18. *Let $a, b, c \in \text{SEStates}_{\text{prg}}$ be two SE states with the same program counter, and let $\dot{\sqcup}$ be an idempotent (SEL1) and associative (SEL3) join operation with the corresponding partial order relation \lesssim . Then, if $a \dot{\sqcup} b = c \implies a \lesssim c$, it holds that $a \leq c \implies a \lesssim c$. \diamond*

Proof. We obtain the equivalence $a \dot{\sqcup} b = c \iff a \leq c$ from Lemma 3.13. Therefore, it follows from $a \dot{\sqcup} b = c \implies a \lesssim c$ that $a \leq c \implies a \lesssim c$. □




Figure 3.3: Double application of the weakening rule




Figure 3.4: Design space for join operators

3.3 Constructing Join Nodes

In this section, we study instantiations of operators $\dot{\sqcup} : SEStates_{prg} \times SEStates_{prg} \rightarrow SEStates_{prg}$ inducing families of join-semilattices.

Figure 3.4 illustrates the design space for the operators $\dot{\sqcup}$. The mentioned subsumption techniques are referred to in Section 6.1. Subsequently, we discuss techniques ranging from zero precision (forgetting the symbolic states) to full precision (remembering all values in the symbolic states), also comprising methods in between those border cases.

Remark 3.19. For the join techniques presented subsequently, we assume that the considered program variables are deterministically ordered according to some fixed, e.g. alphabetic, total order. Furthermore, we also assume that the generation of fresh constants for given program variables is deterministic. This may be accomplished by appending a numeric postfix to the name of the given variable and incrementing this number until the such generated symbol does not occur in the SE states that are joined. This property ensures that the names of fresh constants do not depend on the order of the input states, and thus simplifies some of our proofs. \diamond

3.3.1 Full Anonymization

The most radical way of constructing a join node inducing a family of join-semilattices of SE states is to completely “forget” the symbolic state of the parent nodes that are to be joined for variables that are assigned different values. In this technique, the value of a program location is set to a fixed, but unknown value. Technically, this is realized by the introduction of a fresh constant (of suitable type). The technique marks the zero-precision border of the design space for join operations.

Technique 3.20 (Full Anonymization Join Method). *Given (U_1, C_1, φ) and $(U_2, C_2, \varphi) \in SEStates_{prg}$ for any program Prg with program variables $x_1, x_2, \dots, x_n \in PV$, the join operation $\dot{\sqcup}_{fresh} : SEStates_{prg} \times SEStates_{prg} \rightarrow SEStates_{prg}$ is defined by*

$$(U_1, C_1, \varphi) \dot{\sqcup}_{fresh} (U_2, C_2, \varphi) := ((U_1, C_1) \oplus_{fresh} (U_2, C_2), (U_1, C_1) \odot_{fresh} (U_2, C_2), \varphi)$$

where $(U_1, C_1) \dot{\cup}_{fresh} (U_2, C_2) := C_1 \vee C_2$ and $(U_1, C_1) \oplus_{fresh} (U_2, C_2) := \{x_1 := t_1 \parallel x_2 := t_2 \parallel \dots \parallel x_n := t_n\}$ with

$$t_i := \begin{cases} \{U_1\} x_i & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(x_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(x_i)) \quad (\star) \\ c_i & \text{otherwise} \end{cases}$$

for fresh constant symbols $c_i \in \text{Func}$ of suitable types that are not contained in $U_1, U_2, C_1, C_2, \varphi$, and a fresh, uninterpreted predicate symbol P . \diamond

The following proposition assures that joins using Technique 3.20 are sound in the KeY calculus by Theorem 3.16.

Proposition 3.21. *Technique 3.20 induces a family of join-semilattices of SE states, i.e. the operation $\dot{\cup}_{fresh}$ and its associated partial order relation \preceq satisfy the axioms (SEL1) to (SEL5) of Definition 3.12.* \diamond

Proof. (SEL1) Let $a = (U, C, \varphi) \in \text{SEStates}_{prg}$. Then, $a \dot{\cup}_{fresh} a = (U^*, C \vee C, \varphi)$. In U^* , each program variable x_i is set to the term to which it evaluates in U , since $(C \rightarrow \{U\} P(x_i)) \leftrightarrow (C \rightarrow \{U\} P(x_i))$ is always provable. Because of this and idempotency of disjunction, we have $a \dot{\cup}_{fresh} a \stackrel{concr}{=} a$ in both directions of $\stackrel{concr}{=}$.

(SEL2) Let $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi) \in \text{SEStates}_{prg}$. We prove

$$a \dot{\cup}_{fresh} b =: (U_1^*, C_1^*, \varphi) \stackrel{concr}{=} (U_2^*, C_2^*, \varphi) := b \dot{\cup}_{fresh} a$$

by showing

$$\sigma \in \text{concr}(a \dot{\cup}_{fresh} b) \iff \sigma \in \text{concr}(b \dot{\cup}_{fresh} a).$$

It holds that

$$\begin{aligned} & \sigma \in \text{concr}(a \dot{\cup}_{fresh} b) \\ \iff & \text{there is } (K_\Sigma, \sigma'), (K_\Sigma, \sigma') \models C_1 \vee C_2 \text{ and } \sigma = \text{val}_{(K_\Sigma, \sigma')}(U_1^*)(\sigma') \\ \iff & (K_\Sigma, \sigma') \models C_2 \vee C_1 \\ \iff & \sigma'' := \text{val}_{(K_\Sigma, \sigma')}(U_2^*)(\sigma') \in \text{concr}(b \dot{\cup}_{fresh} a). \end{aligned}$$

Therefore it suffices to show that for every x_i :

$$(K_\Sigma, \sigma') \models \{U_1^*\} P(x_i) \leftrightarrow \{U_2^*\} P(x_i),$$

for a suitable new predicate P , i.e. that the right sides for x_i are equivalent in (K_Σ, σ') . This is trivially the case for equal new constants c_i as right sides. Note that a right side in U_1^* is a fresh constant iff the corresponding right side in U_2^* is a fresh constant, since the bi-implication in (\star) is commutative; furthermore, new constants are assumed to be deterministically generated (\rightarrow Remark 3.19). In the other case, that is if (\star) is provable, the equivalence obviously also holds.

(SEL3) Let $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U_3, C_3, \varphi) \in \text{SEStates}_{prg}$. We have to show that

$$(a \dot{\cup}_{fresh} b) \dot{\cup}_{fresh} c \stackrel{concr}{=} a \dot{\cup}_{fresh} (b \dot{\cup}_{fresh} c),$$

that is

$$\sigma \in \text{concr}((a \dot{\cup}_{fresh} b) \dot{\cup}_{fresh} c) \iff \sigma \in \text{concr}(a \dot{\cup}_{fresh} (b \dot{\cup}_{fresh} c)).$$

Let

$$\begin{aligned} a \dot{\cup}_{fresh} b & := (U_1^*, C_1 \vee C_2, \varphi) \\ b \dot{\cup}_{fresh} c & := (U_2^*, C_2 \vee C_3, \varphi) \\ (a \dot{\cup}_{fresh} b) \dot{\cup}_{fresh} c & := (U_1^{**}, (C_1 \vee C_2) \vee C_3, \varphi) \\ a \dot{\cup}_{fresh} (b \dot{\cup}_{fresh} c) & := (U_2^{**}, C_1 \vee (C_2 \vee C_3), \varphi) \end{aligned}$$

and $\sigma \in \text{concr}((a \dot{\downarrow}_{\text{fresh}} b) \dot{\downarrow}_{\text{fresh}} c)$. Therefore, there is a model (K_Σ, σ') such that $(K_\Sigma, \sigma') \models (C_1 \vee C_2) \vee C_3$ and $\sigma = \text{val}_{(K_\Sigma, \sigma')} (U_1^{**})(\sigma')$. Obviously, it holds that $(K_\Sigma, \sigma') \models C_1 \vee (C_2 \vee C_3)$ due to associativity of disjunction, thus it holds that

$$\text{val}_{(K_\Sigma, \sigma')} (U_2^{**})(\sigma') \in \text{concr}(a \dot{\downarrow}_{\text{fresh}} (b \dot{\downarrow}_{\text{fresh}} c)).$$

It remains to show that $\text{val}_{(K_\Sigma, \sigma')} (U_2^{**})(\sigma') = \text{val}_{(K_\Sigma, \sigma')} (U_1^{**})(\sigma')$. Let x be in x_1, x_2, \dots, x_n . Consider two cases: (i) The right side for x in U_1^{**} is $\{U_1^*\}x$. Then, the terms $\{U_1\}x$, $\{U_2\}x$, $\{U_3\}x$ and also $\{U_1^*\}x$ must be equivalent in every model due to the definition of t_i . The right side in U_1^{**} is therefore $\{U_1\}x$ which we know is equivalent to $\{U_1^*\}x$. (ii) The right side for x in U_1^{**} is a Skolem constant. Then, by (\star) , $\{U_1^*\}x$ and $\{U_3\}x$ are not logically equivalent terms. Thus, either $\{U_1\}x$ and $\{U_2\}x$ are not equivalent, or both are not equivalent to $\{U_3\}x$. In both of these cases, $\{U_1\}x$ and $\{U_2^*\}x$ cannot be equivalent, and therefore the right side for x in U_2^{**} is also the same constant.

(SEL4) Since we already proved idempotency and associativity, it suffices to show

$$a \dot{\downarrow}_{\text{fresh}} b = c \implies a \lesssim c$$

for any $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U^*, C_1 \vee C_2, \varphi) \in \text{SEStates}_{\text{prg}}$ by Lemma 3.18. Assuming $a \dot{\downarrow}_{\text{fresh}} b = c$, we thus have to show that $\text{concr}(a) \subseteq \text{concr}(c)$. Let $\sigma \in \text{concr}(a)$, i.e. there is a model $(K_\Sigma, \sigma') \models C_1$ such that $\sigma = \text{val}_{(K_\Sigma, \sigma')} (U_1)(\sigma')$. Let K'_Σ be an expansion of K_Σ by the new constants c_i such that $I_{K'_\Sigma}(c_i) = \sigma(x_i)$. Since C_1 is in the old language, $(K'_\Sigma, \sigma') \models C_1$ and thus $(K'_\Sigma, \sigma') \models C_1 \vee C_2$. Therefore, we have

$$\text{val}_{(K'_\Sigma, \sigma')} (U^*)(\sigma') \in \text{concr}(c).$$

It suffices to show that, for all $i = 1, 2, \dots, n$:

$$\text{val}_{(K'_\Sigma, \sigma')} (U^*)(\sigma')(x_i) = \text{val}_{(K_\Sigma, \sigma')} (U_1)(\sigma')(x_i).$$

We distinguish the following two cases:

Case “ (\star) is provable”. In this case,

$$\begin{aligned} \text{val}_{(K'_\Sigma, \sigma')} (U^*)(\sigma')(x_i) &= \text{val}_{(K'_\Sigma, \sigma')} (x_i := \{U_1\}x_i)(\sigma')(x_i) \\ &= \left(x \mapsto \begin{cases} \sigma'(x) & \text{if } x \neq x_i \\ \text{val}_{(K'_\Sigma, \sigma')} (\{U_1\}x_i) & \text{otherwise} \end{cases} \right)(x_i) \\ &= \text{val}_{(K'_\Sigma, \sigma')} (\{U_1\}x_i) \\ &= \text{val}_{(K'_\Sigma, \sigma')} (U_1)(\sigma')(x_i) \\ &= \text{val}_{(K_\Sigma, \sigma')} (U_1)(\sigma')(x_i). \end{aligned}$$

where the last equation is true since U_1 is in the language without the constants.

Case “otherwise”. Then,

$$\begin{aligned} \text{val}_{(K'_\Sigma, \sigma')} (U^*)(\sigma')(x_i) &= \text{val}_{(K'_\Sigma, \sigma')} (x_i := c_i)(\sigma')(x_i) \\ &= \left(x \mapsto \begin{cases} \sigma'(x) & \text{if } x \neq x_i \\ \text{val}_{(K'_\Sigma, \sigma')} (c_i) & \text{otherwise} \end{cases} \right)(x_i) \\ &= \text{val}_{(K'_\Sigma, \sigma')} (c_i) = I_{K'_\Sigma}(c_i) \stackrel{\text{def.}}{=} \sigma(x_i) \\ &= \text{val}_{(K_\Sigma, \sigma')} (U_1)(\sigma')(x_i). \end{aligned}$$

(SEL5) We have to show that $C^* \leftrightarrow (C \wedge C_{\text{ax}})$, where $\bigwedge C_1 \rightarrow C$ and $\exists \bar{v}; (C_{\text{ax}}[\bar{v}/\bar{c}])$, for $\bar{c} = c_1, c_2, \dots, c_n$ being the constants introduced in c and \bar{v} being a tuple of suitable fresh logical variables, is provable. By definition, $C^* = (\bigwedge C_1 \vee \bigwedge C_2)$. We choose $C := \bigwedge C_1 \vee \bigwedge C_2$ and $C_{\text{ax}} := \text{true}$. Obviously, $\bigwedge C_1 \rightarrow (\bigwedge C_1 \vee \bigwedge C_2)$, and true is trivially provable. \square

3.3.2 If-Then-Else

The “If-Then-Else” technique is a “classic” of state joining for symbolic execution (see, e.g., [HSS09; Kuz+12; Sen+14]) that retains the full precision of the analysis. We realize the technique by employing the if-the-else construct of Java DL (\rightarrow Definition 2.6). An alternative way is the introduction of fresh Skolem constants for the symbolic states that are constrained by implications in the path condition. We discuss this approach in the subsequent subsection.

Technique 3.22 (If-Then-Else Join Method). *Given (U_1, C_1, φ) and $(U_2, C_2, \varphi) \in SEStates_{Pr_g}$ for any program Pr_g with program variables $x_1, x_2, \dots, x_n \in PV$, the join operation $\dot{\sqcup}_{ite} : SEStates_{Pr_g} \times SEStates_{Pr_g} \rightarrow SEStates_{Pr_g}$ is defined by*

$$(U_1, C_1, \varphi) \dot{\sqcup}_{ite} (U_2, C_2, \varphi) := ((U_1, C_1) \oplus_{ite} (U_2, C_2), (U_1, C_1) \odot_{ite} (U_2, C_2), \varphi)$$

where $(U_1, C_1) \odot_{ite} (U_2, C_2) := (\bigwedge C_1) \vee (\bigwedge C_2)$ and $(U_1, C_1) \oplus_{ite} (U_2, C_2) := \{x_1 := t_1 \parallel \dots \parallel x_n := t_n\}$ with

$$t_i := \begin{cases} \{U_1\} x_i & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(x_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(x_i)) \quad (\star) \\ \text{if } (\bigwedge C_1) \text{ then } (\{U_1\} x_i) \text{ else } (\{U_2\} x_i) & \text{otherwise} \end{cases}$$

for a fresh, uninterpreted predicate symbol P . \diamond

Remark 3.23 (Incompatible Path Conditions). For our join techniques, and in particular for the techniques 3.22 and 3.25, we assume that two path conditions C_1 and C_2 of different branches in a tree are generally incompatible. This assumption is reasonable since splits in symbolic execution should only occur when the execution depends on the concrete value of a variable; in this case, a case distinction takes place. If C_1 and C_2 are the path conditions for those case distinction branches, it holds that $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$, as well as the contraposition $\vdash \bigwedge C_2 \rightarrow \neg \bigwedge C_1$. The assumption *could* be problematic when If-Then-Else joins are applied to an SET where other branches are joined with rules violating these constraints. Since our proposed join rules either use the canonical disjunction of path conditions for the path conditions of join nodes, or a strengthening of that (in particular, they respect the property (SEL5)), we are confident that this issue can be excluded. \diamond

Proposition 3.24. *Technique 3.22 induces a family of join-semilattices of SE states, i.e. the operation $\dot{\sqcup}_{ite}$ and its associated partial order relation \preceq satisfy the axioms (SEL1) to (SEL5) of Definition 3.12. \diamond*

Proof. (SEL1) The proof for (SEL1) is similar to the corresponding case in the proof of Proposition 3.21.

(SEL2) Under the assumption $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$, the terms $\text{if } (\bigwedge C_1) \text{ then } (\{U_1\} x) \text{ else } (\{U_2\} x)$ and $\text{if } (\bigwedge C_2) \text{ then } (\{U_2\} x) \text{ else } (\{U_1\} x)$ are logically equivalent. Based on this observation, the proof of (SEL2) is similar to the corresponding case in the proof of Proposition 3.21.

(SEL3) Under the assumption $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$, the proof of (SEL3) is similar to the corresponding case in the proof of Proposition 3.21.

(SEL4) Since we already proved idempotency and associativity, it suffices to show

$$a \dot{\sqcup}_{ite} b = c \implies a \preceq c$$

for any $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U^*, C_1 \vee C_2, \varphi) \in SEStates_{Pr_g}$ by Lemma 3.18. Assuming $a \dot{\sqcup}_{ite} b = c$, we thus have to show that $\text{concr}(a) \subseteq \text{concr}(c)$. Let $\sigma \in \text{concr}(a)$, i.e. there is a model $(K_\Sigma, \sigma') \models C_1$ such that $\sigma = \text{val}_{(K_\Sigma, \sigma')}(U_1)(\sigma')$. Since furthermore, $(K_\Sigma, \sigma') \models C_1 \vee C_2$, it holds that

$$\text{val}_{(K_\Sigma, \sigma')}(U^*)(\sigma') \in \text{concr}(c).$$

It suffices to show that, for all $i = 1, 2, \dots, n$:

$$\text{val}_{(K_\Sigma, \sigma')}(U^*)(\sigma')(x_i) = \text{val}_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i).$$

We distinguish the following two cases:

Case “(*) is provable”. In this case,

$$\begin{aligned}
\text{val}_{(K_\Sigma, \sigma')} (U^*) (\sigma') (x_i) &= \text{val}_{(K_\Sigma, \sigma')} (x_i := \{U_1\} x_i) (\sigma') (x_i) \\
&= \left(x \mapsto \begin{cases} \sigma' (x) & \text{if } x \neq x_i \\ \text{val}_{(K_\Sigma, \sigma')} (\{U_1\} x_i) & \text{otherwise} \end{cases} \right) (x_i) \\
&= \text{val}_{(K_\Sigma, \sigma')} (\{U_1\} x_i) \\
&= \text{val}_{(K_\Sigma, \sigma')} (U_1) (\sigma') (x_i)
\end{aligned}$$

Case “otherwise”. Then,

$$\begin{aligned}
&\text{val}_{(K_\Sigma, \sigma')} (U^*) (\sigma') (x_i) = \\
&\text{val}_{(K_\Sigma, \sigma')} (x_i := \text{if} (\bigwedge C_1) \text{then} (\{U_1\} x_i) \text{else} (\{U_2\} x_i)) (\sigma') (x_i) = \\
&\left(x \mapsto \begin{cases} \sigma' (x) & \text{if } x \neq x_i \\ \text{val}_{(K_\Sigma, \sigma')} (\text{if} (\bigwedge C_1) \text{then} (\{U_1\} x_i) \text{else} (\{U_2\} x_i)) & \text{otherwise} \end{cases} \right) (x_i) = \\
&\left(x \mapsto \begin{cases} \sigma' (x) & \text{if } x \neq x_i \\ \text{val}_{(K_\Sigma, \sigma')} (\{U_1\} x_i) & x = x_i \wedge \text{val}_{(K_\Sigma, \sigma')} (\bigwedge C_1) = tt \\ \text{val}_{(K_\Sigma, \sigma')} (\{U_2\} x_i) & \text{otherwise} \end{cases} \right) (x_i) = \\
&\text{val}_{(K_\Sigma, \sigma')} (\{U_1\} x_i) = \\
&\text{val}_{(K_\Sigma, \sigma')} (U_1) (\sigma') (x_i).
\end{aligned}$$

(SEL5) We have to show that $C^* \leftrightarrow (C \wedge C_{ax})$, where $\bigwedge C_1 \rightarrow C$ and $\exists \bar{v}; (C_{ax} [\bar{v} / \bar{c}])$, is provable, which is equivalent to C_{ax} being provable since \downarrow_{ite} does not introduce new constants. By definition, $C^* = (\bigwedge C_1 \vee \bigwedge C_2)$. We choose $C := \bigwedge C_1 \vee \bigwedge C_2$ and $C_{ax} := true$. Obviously, $\bigwedge C_1 \rightarrow (\bigwedge C_1 \vee \bigwedge C_2)$, and $true$ is trivially provable. \square

3.3.3 If-Then-Else by additional Path Condition Constraints

An alternative to the If-Then-Else approach discussed in Section 3.3.2 is the realization of constraints of program variables by the introduction of fresh Skolem constants along with additional constraints in the path conditions of join states. While being logically equivalent (under the assumption $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$ of Remark 3.23), this technique has a different runtime behavior when implemented in KeY (\rightarrow Section 5.1). Since the path condition of a state is usually contained in the antecedent of the corresponding sequent in the SET, we also refer to this technique as “If-Then-Else-Antecedent”.

Technique 3.25 (If-Then-Else with PC Constraints Join Method). *Given (U_1, C_1, φ) and $(U_2, C_2, \varphi) \in SEStates_{Prg}$ for any program Prg with program variables $x_1, x_2, \dots, x_n \in PV$, the join operation $\downarrow_{ite2} : SEStates_{Prg} \times SEStates_{Prg} \rightarrow SEStates_{Prg}$ is defined by*

$$(U_1, C_1, \varphi) \downarrow_{ite2} (U_2, C_2, \varphi) := ((U_1, C_1) \oplus_{ite2} (U_2, C_2), (U_1, C_1) \odot_{ite2} (U_2, C_2), \varphi)$$

where $(U_1, C_1) \oplus_{ite2} (U_2, C_2) := \{x_1 := t_1 \parallel x_2 := t_2 \parallel \dots \parallel x_n := t_n\}$ with

$$t_i := \begin{cases} \{U_1\} x_i & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(x_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(x_i)) \\ c_i & \text{otherwise} \end{cases}$$

and $(U_1, C_1) \odot_{ite2} (U_2, C_2) := (\bigwedge C_1 \vee \bigwedge C_2) \wedge \bigwedge C_i^{ite}$ where

$$C_i^{ite} := \begin{cases} true & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(x_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(x_i)) \\ (\bigwedge C_1 \rightarrow c_i \doteq \{U_1\} x_i) \wedge & \text{otherwise} \\ (\bigwedge C_2 \rightarrow c_i \doteq \{U_2\} x_i) & \end{cases}$$

for a fresh, uninterpreted predicate symbol P and fresh Skolem constants c_1, c_2, \dots, c_n of appropriate types that are not contained in $U_1, U_2, C_1, C_2, \varphi$. For the “otherwise” case, we assume $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$ (\rightarrow Remark 3.23). \diamond

Proposition 3.26. *Technique 3.25 induces a family of join-semilattices of SE states, i.e. the operation $\dot{\sqcup}_{ite2}$ and its associated partial order relation \preceq satisfy the axioms (SEL1) to (SEL5) of Definition 3.12. \diamond*

Proof. We omit the proofs of (SEL1) to (SEL3) which are basically obvious when following the arguments in the corresponding proofs of the techniques discussed above.

For (SEL4), it suffices to show

$$a \dot{\sqcup}_{ite2} b = c \implies a \preceq c$$

for any $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U^*, C_1 \vee C_2, \varphi) \in SEStates_{prg}$ by Lemma 3.18. Assuming $a \dot{\sqcup}_{ite2} b = c$, we thus have to show that $concr(a) \subseteq concr(c)$. We again assume that $\vdash \bigwedge C_1 \rightarrow \neg \bigwedge C_2$. Let $\sigma \in concr(a)$, i.e. there is a model $(K_\Sigma, \sigma') \models C_1$ such that $\sigma = val_{(K_\Sigma, \sigma')}(U_1)(\sigma')$. Let K'_Σ be an expansion of K_Σ by the new constants c_i such that $I_{K'_\Sigma}(c_i) = \sigma(x_i)$. Since C_1 is in the old language, $(K'_\Sigma, \sigma') \models C_1$ (which is a direct consequence of Herbrand's theorem), $(K'_\Sigma, \sigma') \models C_1 \vee C_2$ and $(K'_\Sigma, \sigma') \models \bigwedge C_i^{ite}$. Therefore, we have

$$val_{(K'_\Sigma, \sigma')}(U^*)(\sigma') \in concr(c).$$

It suffices to show that, for all $i = 1, 2, \dots, n$:

$$val_{(K'_\Sigma, \sigma')}(U^*)(\sigma')(x_i) = val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i).$$

We distinguish the following two cases:

Case “ (\star) is provable”. In this case,

$$\begin{aligned} val_{(K'_\Sigma, \sigma')}(U^*)(\sigma')(x_i) &= val_{(K'_\Sigma, \sigma')}(x_i := \{U_1\} x_i)(\sigma')(x_i) \\ &= \left(x \mapsto \begin{cases} \sigma'(x) & \text{if } x \neq x_i \\ val_{(K'_\Sigma, \sigma')}(\{U_1\} x_i) & \text{otherwise} \end{cases} \right)(x_i) \\ &= val_{(K'_\Sigma, \sigma')}(\{U_1\} x_i) \\ &= val_{(K'_\Sigma, \sigma')}(U_1)(\sigma')(x_i) \\ &\stackrel{(\dagger)}{=} val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i). \end{aligned}$$

where (\dagger) is true since U_1 is in the language without the constants.

Case “otherwise”. Then,

$$\begin{aligned} val_{(K'_\Sigma, \sigma')}(U^*)(\sigma')(x_i) &= val_{(K'_\Sigma, \sigma')}(x_i := c_i)(\sigma')(x_i) \\ &= \left(x \mapsto \begin{cases} \sigma'(x) & \text{if } x \neq x_i \\ val_{(K'_\Sigma, \sigma')}(c_i) & \text{otherwise} \end{cases} \right)(x_i) \\ &= val_{(K'_\Sigma, \sigma')}(c_i) = I_{K'_\Sigma}(c_i) \stackrel{\text{def.}}{=} \sigma(x_i) \\ &= val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i). \end{aligned}$$

For (SEL5), we have to show that $C^* \leftrightarrow (C \wedge C_{ax})$, where $\bigwedge C_1 \rightarrow C$ and $\exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$, for $\bar{c} = c_1, c_2, \dots, c_n$ and \bar{v} is a tuple of suitable fresh logical variables, is provable. By definition, $C^* = (\bigwedge C_1 \vee \bigwedge C_2) \wedge \bigwedge C_i^{ite}$. We choose $C := \bigwedge C_1 \vee \bigwedge C_2$ and $C_{ax} := \bigwedge C_i^{ite}$. Obviously, $\bigwedge C_1 \rightarrow (\bigwedge C_1 \vee \bigwedge C_2)$. $\exists \bar{v}; (C_{ax}[\bar{v}/\bar{c}])$ equals

$$\exists \bar{v}; \left(\bigwedge \{ (\bigwedge C_1 \rightarrow v_i \doteq \{U_1\} x_i) \wedge (\bigwedge C_2 \rightarrow v_i \doteq \{U_2\} x_i) \} \right).$$

This formula is easily provable for the instantiation

$$v_i := \begin{cases} \{U_1\} x_i & \text{if } C_1 \text{ is provable} \\ \{U_2\} x_i & \text{otherwise} \end{cases}$$

\square




Figure 3.5: Abstract domain for sign analysis

3.3.4 Abstract Weakening

Abstract Interpretation [CC77] is a static analysis technique that allows for the construction of fully automatic proof methods [Cou+05]. Our General Lattice Framework, along with the technique proposed subsequently, at least partly closes the gap between symbolic execution and abstract interpretation by facilitating joins based on abstract domain lattices. We first define the notion of abstract domain elements.

Definition 3.27 (Abstract Domain Element). An *Abstract Domain Element* is a function $defAx : \text{Terms}_\Sigma \rightarrow \text{Form}_\Sigma$ mapping Java DL terms of appropriate types to *closed* formulae. \diamond

Intuitively, an abstract domain element models an infinite set of *defining axioms* for Java DL terms. If an axiom is true for a given term, then this term is described by the corresponding abstract domain element. This rather technical, syntactical definition is beneficial for the application in branch joining. Note that we restrict abstract elements / domains to those that can be characterized in Java DL.

Definition 3.28 (Abstract Domain Lattice). An *Abstract Domain Lattice* is a join-semilattice $\mathcal{A}_T = (A_T, \sqcup)$ with the induced partial order relation \sqsubseteq for a countable set A_T of abstract domain elements accepting terms of some fixed type $T \in \mathcal{T}$ as arguments. We impose the following requirements on A_T and \sqsubseteq :

- (1) A_T includes two elements $\perp(t) = \text{false}$ and $\top(t) = \text{true}$
- (2) For $a, b \in A_T$ with $a \sqsubseteq b$, it holds that $\vdash a(t) \rightarrow b(t)$ for any term t of type T .
- (3) For all $a \in A_T$ except for \perp , it holds that $\vdash \exists v; a(v)$. \diamond

Example 3.29. A well-known example for an abstract domain lattice also used in [BHW09] is the lattice $\mathcal{A}_{\text{int}}^{\text{sgn}} = (A_{\text{int}}^{\text{sgn}}, \sqcup)$ for sign analysis visualized in the Hasse diagram of Figure 3.5. $A_{\text{int}}^{\text{sgn}}$ is defined as $A_{\text{int}}^{\text{sgn}} = \{\perp, \text{neg}, \text{zero}, \text{pos}, \text{leq}, \text{geq}, \top\}$ where $\perp(t) := \text{false}$, $\text{neg}(t) := t < 0$, $\text{zero}(t) := t \doteq 0$, $\text{pos}(t) := t > 0$, $\text{leq}(t) := t \leq 0$, $\text{geq}(t) := t \geq 0$ and $\top(t) := \text{true}$. The join operator \sqcup and its induced partial order relation \sqsubseteq are defined like depicted in the Hasse diagram. Intuitively, geq , for instance, represents the positive integers including zero, while pos represents the positive integers without zero; \top represents all integers, whilst \perp does not represent any integer. \diamond

For the definition of the following technique, we assume that there are only program variables of the same type $T \in \mathcal{T}$. However, the technique is easily generalized to multiple types by demanding the existence of one abstract domain lattice per type; moreover, it is possible to define “hybrid” approaches using, for instance, the “If-Then-Else” join technique whenever there is no abstract domain lattice for a given type.

Technique 3.30 (Abstract Weakening Join Method). Let $\mathcal{A}_T = (A_T, \sqcup)$ be an abstract domain lattice. Given (U_1, C_1, φ) and $(U_2, C_2, \varphi) \in \text{SEStates}_{\text{prg}}$ for any program Prg with program variables $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \text{PV}$ of type T , the join operation $\dot{\sqcup}_{\text{abstr}} : \text{SEStates}_{\text{prg}} \times \text{SEStates}_{\text{prg}} \rightarrow \text{SEStates}_{\text{prg}}$ is defined by

$$(U_1, C_1, \varphi) \dot{\sqcup}_{\text{abstr}} (U_2, C_2, \varphi) := ((U_1, C_1) \oplus_{\text{abstr}} (U_2, C_2), (U_1, C_1) \otimes_{\text{abstr}} (U_2, C_2), \varphi)$$

where $(U_1, C_1) \oplus_{\text{abstr}} (U_2, C_2) := \{\mathbf{x}_1 := t_1 \parallel \mathbf{x}_2 := t_2 \parallel \dots \parallel \mathbf{x}_n := t_n\}$ with, for a fresh, uninterpreted predicate symbol P and fresh Skolem constants c_1, c_2, \dots, c_n of appropriate types that are not contained in $U_1, U_2, C_1, C_2, \varphi$,

$$t_i := \begin{cases} \{U_1\} \mathbf{x}_i & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(\mathbf{x}_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(\mathbf{x}_i)) \\ c_i & \text{otherwise} \end{cases}$$

and $(U_1, C_1) \dot{\sqcup}_{abstr} (U_2, C_2) := (\bigwedge C_1 \vee \bigwedge C_2) \wedge \bigwedge C_i^{abs}$ where

$$C_i^{abs} := \begin{cases} true & \text{if } \vdash (C_1 \rightarrow \{U_1\} P(x_i)) \leftrightarrow (C_2 \rightarrow \{U_2\} P(x_i)) \\ (defAx_1 \sqcup defAx_2)(c_i) & \text{otherwise} \end{cases}$$

and, for $k \in 1, 2$, $defAx_k \in A_T$ is an abstract domain element such that $\vdash C_k \rightarrow defAx_k(\{U_k\} \mathbf{x}_i)$ is provable and there is no element $defAx'_k \in A_T$ with $defAx'_k \neq defAx_k$ and $defAx'_k \sqsubseteq defAx_k$. \diamond

The constraints on $defAx_k$ state that those elements must be least ones in the abstract domain lattice. Note that there does not necessarily exist a *unique* such element; in this case, any element for which there is no suitable strictly smaller one suffices. For countable lattices with a given enumeration $e_1, e_2, \dots, e_n, \dots$ of the domain elements such that for $0 \leq i < l$, it does not hold that $e_i \sqsubseteq e_l$, $defAx_k$ is computable. Such an enumeration can always be constructed for finite domains by a BFS traversal of the lattice starting from the smallest element. For the sign analysis domain, the enumeration is $\perp, \text{neg}, \text{zero}, \text{pos}, \text{leq}, \text{geq}, \top$. Generally, infinite domains should support *widening* [CC77] to ensure that suitable abstractions can be computed.

Proposition 3.31. *Technique 3.30 induces a family of join-semilattices of SE states, i.e. the operation $\dot{\sqcup}_{abstr}$ and its associated partial order relation \preceq satisfy the axioms (SEL1) to (SEL5) of Definition 3.12.* \diamond

Proof. We omit the proofs of the rather obvious properties (SEL1) to (SEL3).

For (SEL4), it suffices to show

$$a \dot{\sqcup}_{abstr} b = c \implies a \preceq c$$

for any $a = (U_1, C_1, \varphi)$, $b = (U_2, C_2, \varphi)$, $c = (U^*, C_1 \vee C_2, \varphi) \in SEStates_{prg}$ by Lemma 3.18. Assuming $a \dot{\sqcup}_{abstr} b = c$, we thus have to show that $concr(a) \subseteq concr(c)$. Let $\sigma \in concr(a)$, i.e. there is a model $(K_\Sigma, \sigma') \models C_1$ such that $\sigma = val_{(K_\Sigma, \sigma')}(U_1)(\sigma')$. Let K'_Σ be an expansion of K_Σ by the new constants c_i such that $I_{K'_\Sigma}(c_i) = \sigma(x_i)$. Since C_1 is in the old language, $(K'_\Sigma, \sigma') \models C_1$ and $(K'_\Sigma, \sigma') \models C_1 \vee C_2$. By $\vdash C_k \rightarrow defAx_k(\{U_k\} \mathbf{x}_i)$, $k = 1, 2$, we know $(K'_\Sigma, \sigma') \models defAx_1(\{U_1\} \mathbf{x}_i)$ and, by

$$val_{(K'_\Sigma, \sigma')}(\{U_1\} \mathbf{x}_i) = val_{(K'_\Sigma, \sigma')}(U_1)(\sigma')(x_i) = val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i) = \sigma(x_i) = I_{K'_\Sigma}(c_i),$$

$(K'_\Sigma, \sigma') \models defAx_1(c_i)$. Due to Definition 3.28, it thus holds that $(K'_\Sigma, \sigma') \models (defAx_1 \sqcup defAx_2)(c_i)$ and so $(K'_\Sigma, \sigma') \models \bigwedge C_i^{abs}$. Therefore, we have

$$val_{(K'_\Sigma, \sigma')}(U^*)(\sigma') \in concr(c).$$

It suffices to show that, for all $i = 1, 2, \dots, n$:

$$val_{(K'_\Sigma, \sigma')}(U^*)(\sigma')(x_i) = val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i).$$

We distinguish the following two cases:

Case “(*) is provable”. In this case,

$$\begin{aligned} val_{(K'_\Sigma, \sigma')}(U^*)(\sigma')(x_i) &= val_{(K'_\Sigma, \sigma')}(x_i := \{U_1\} \mathbf{x}_i)(\sigma')(x_i) \\ &= \left(x \mapsto \begin{cases} \sigma'(x) & \text{if } x \neq x_i \\ val_{(K'_\Sigma, \sigma')}(\{U_1\} \mathbf{x}_i) & \text{otherwise} \end{cases} \right)(x_i) \\ &= val_{(K'_\Sigma, \sigma')}(\{U_1\} \mathbf{x}_i) \\ &= val_{(K'_\Sigma, \sigma')}(U_1)(\sigma')(x_i) \\ &\stackrel{(\dagger)}{=} val_{(K_\Sigma, \sigma')}(U_1)(\sigma')(x_i). \end{aligned}$$

where (\dagger) is true since U_1 is in the language without the constants.

Case “otherwise”. Then,

$$\begin{aligned}
\text{val}_{(K'_\Sigma, \sigma')} (U^*) (\sigma') (x_i) &= \text{val}_{(K'_\Sigma, \sigma')} (x_i := c_i) (\sigma') (x_i) \\
&= \left(x \mapsto \begin{cases} \sigma' (x) & \text{if } x \neq x_i \\ \text{val}_{(K'_\Sigma, \sigma')} (c_i) & \text{otherwise} \end{cases} \right) (x_i) \\
&= \text{val}_{(K'_\Sigma, \sigma')} (c_i) = I_{K'_\Sigma} (c_i) \stackrel{\text{def.}}{=} \sigma (x_i) \\
&= \text{val}_{(K_\Sigma, \sigma')} (U_1) (\sigma') (x_i).
\end{aligned}$$

For (SEL5), we have to show that $C^* \leftrightarrow (C \wedge C_{ax})$, where $\bigwedge C_1 \rightarrow C$ and $\exists \bar{v}; (C_{ax} [\bar{v}/\bar{c}])$, for $\bar{c} = c_1, c_2, \dots, c_n$ and \bar{v} is a tuple of suitable fresh logical variables, is provable. By definition, $C^* = (\bigwedge C_1 \vee \bigwedge C_2) \wedge \bigwedge C_i^{abs}$. We choose $C := \bigwedge C_1 \vee \bigwedge C_2$ and $C_{ax} := \bigwedge C_i^{abs}$. Obviously, $\bigwedge C_1 \rightarrow (\bigwedge C_1 \vee \bigwedge C_2) \cdot \exists \bar{v}; (C_{ax} [\bar{v}/\bar{c}])$ equals

$$\exists \bar{v}; \left(\bigwedge \{ (\text{defAx}_1 \sqcup \text{defAx}_2) (v_i) : i = 1, 2, \dots, n \} \right).$$

This formula is provable due to (2) of Definition 3.28, since neither defAx_1 nor defAx_2 are \perp , because otherwise $\vdash C_k \rightarrow \text{defAx}_k (\{U_k\} x_i)$ would not hold for satisfiable C_k . \square

3.3.5 Heap Treatment

All of the above discussed join technique are in principle able to deal with heap structures. Since heaps are modeled in KeY as program variables of the special type *Heap*, the techniques 3.20, 3.22 and 3.25 are able to treat heaps without special adaptations; Technique 3.30 would depend on a domain for the type *Heap*. However, the joining of whole heap expressions to, e.g., if-then-else expressions, leads to long terms and potential redundancies which may result in a suboptimal treatment of those by the automatic strategies of the KeY system. We therefore treat heaps in a “zip” procedure: Assume, for example, that we want to join two heap expressions

`store(store(create(heap, s_2), s_2, <initialized>, TRUE), s_2, num, mul(num_0, -1))`

and

`store(store(create(heap, s_2), s_2, <initialized>, TRUE), s_2, num, num_0),`

where `num` is negative in the corresponding first state and positive in the second. Instead of just connecting the whole terms by if-then-else, we simultaneously go deeper inside the expressions, as long as the top functions are equal (e.g., two “store” applications) and have the syntactically same expressions for objects / fields as targets, and apply join operators on stored values that differ. When two different top level functions are reached (e.g., a “create” and a “store” occurrence) or (syntactically) different objects / fields are targeted, we use if-then-else as a fallback. For the above example and a join with the sign lattice abstraction, the resulting heap expression is

`store(store(create(heap, s_2), s_2, <initialized>, TRUE), s_2, num, geq),`

where the constant `geq` is constrained by the formula `geq >= 0` in the antecedent. We define this procedure formally in Definition 3.32.

Definition 3.32 (General Heap Treatment). The function

$$\text{joinHeaps} : \text{Terms}_\Sigma^{\text{Heap}} \times \text{Terms}_\Sigma^{\text{Heap}} \times \text{SEStates}_{\text{prg}} \times \text{SEStates}_{\text{prg}} \rightarrow 2^{\text{Form}_\Sigma} \times \text{Terms}_\Sigma^{\text{Heap}}$$

is defined inductively over the structure of heap terms as follows:

$$\text{joinHeaps}(h_1, h_2, s_1, s_2) := \begin{cases} (\emptyset, h_1) & \text{if } h_1 = h_2 \\ (C \cup ((s_1, v_1) \otimes (s_2, v_2)), & \text{if } h_1 = \text{store}(h'_1, o, f, v_1) \text{ and} \\ \quad \text{store}(h', o, f, (s_1, v_1) \oplus (s_2, v_2))) & \quad h_2 = \text{store}(h'_2, o, f, v_2) \\ (C, \text{create}(h', o)) & \text{if } h_1 = \text{create}(h'_1, o) \text{ and} \\ & \quad h_2 = \text{create}(h'_2, o) \\ (\emptyset, \text{if}(C_1) \text{ then}(h_1) \text{ else}(h_2)) & \text{otherwise} \end{cases}$$

where

- $\text{joinHeaps}(h'_1, h'_2, s_1, s_2) = (C, h')$ – the recursive step,

- C_1 is the path condition of s_1 , and
- $\oplus : SEStates_{prg} \times Terms_{\Sigma}^T \times SEStates_{prg} \times Terms_{\Sigma}^T \rightarrow Terms_{\Sigma}^T$ and
 $\otimes : SEStates_{prg} \times Terms_{\Sigma}^T \times SEStates_{prg} \times Terms_{\Sigma}^T \rightarrow 2^{\text{Form}_{\Sigma}}$ are families of join operators for computing join values resp. additional path condition constraints for terms of types T . \diamond

The second component of the result of *joinHeaps* is a term of type *Heap*, the joined heap, while the first component contains the set of formulae constraining introduced Skolem constants, if any. Instances of operators \otimes and \oplus in the above definition could return, for the example of abstract weakening, a defining axiom for a fresh Skolem constant and the constant itself, respectively. The difference to the operators \otimes_{abstr} and \oplus_{abstr} defined on SE states consists in the restriction to consider exactly one term instead of all right sides of contained program variables; otherwise, they work analogously. For the “otherwise” case, we also could have chosen the introduction of a fresh Skolem constant of type *Heap* as a basis for joining heaps by techniques 3.20 (full anonymization) or 3.25 (if-then-else-antecedent).

While this general approach does not really consider special characteristics of objects, there are methods in literature focusing on heap structures. Anand et al. [APV06] propose a technique for lists and arrays building upon shape analysis techniques [SRW02]. They contract list elements to so-called “summary nodes”, which facilitates the joining of branches whenever such an abstraction “subsumes” another leaf in the SET. We refer to this in Section 6.1.

4 Implementation

Based on our theoretical framework presented in Chapter 3, we implemented support for joining branches of proof trees in the KeY system. The implementation consists of five components: (1) An abstract class `JoinRule` following the template method pattern and a class `JoinRuleUtils` comprising a set of commonly used static methods for join operations, (2) a framework for abstraction based on finite lattices, (3) a class `CloseAfterJoin` for closing partner goals after a join, (4) a macro for running proof steps on a sequent containing Java code until a potential join point is reached, and (5) an implementation of four sample join rules specializing `JoinRule`. Subsequently, we elaborate on these components.

4.1 JoinRule and JoinRuleUtils

The class `JoinRule` is responsible for all tasks accumulated during the joining of branches in KeY proof except for the actual computation of a join state for two given symbolic execution states. Its implementation follows the template method pattern: Concrete join rules must implement the abstract method `joinValuesInStates(LocationVariable, SymbolicExecutionState, Term, SymbolicExecutionState, Term, Services)` for joining two terms in their respective SE states. `JoinRuleUtils` comprises a quite large set of static methods that are used by join classes. These methods can be categorized into simple auxiliary methods, methods related to general logic (syntax, provability, simplification, and calculus-related), and methods closely related to join operations. The UML class diagram in Figure 6.1 visualizes the dependencies between the two general join classes and the concrete join rules; Figure 6.2 offers a more detailed view into the structure of `JoinRule` and `JoinRuleUtils`. Listing 4.1 shows an excerpt of the implementation of `JoinRule`. In particular, `JoinRule` checks whether joining is applicable for a given position in a sequent; for investigating whether this is the case, the method `JoinRule.findPotentialJoinPartners(Goal, PosInOccurrence)` searches the goals of the proof tree for potential partner goals. For enabling the user to choose partners among the candidates, the dialog class `JoinPartnerSelectionDialog` is used. After the join partner selection, the node on which the rule is being applied is joined with the partners in a loop (lines 27 – 33 in Listing 4.1). Since join operations are required to be commutative and associative (\rightarrow Definition 3.12), the order of the join does not matter; therefore, concrete join classes inheriting from `JoinRule` only need to specify a method for joining values in *two* SE states. After the join, the rule `CloseAfterJoin` (\rightarrow Section 4.3) is applied on the partner goals. The method `joinStates(...)` may be overridden for join operations requiring special behavior; the same holds for the method `joinHeaps(...)` realizing the procedure described in Definition 3.32.

Listing 4.1: Excerpt of the `JoinRule` implementation

```
1 public abstract class JoinRule extends JoinRuleUtils implements BuiltInRule {
2     // ...
3
4     @Override
5     public final ImmutableList<Goal> apply(Goal goal, final Services services,
6         RuleApp ruleApp) throws RuleAbortException {
7
8         // ...
9         ImmutableList<Pair<Goal, PosInOccurrence>> joinPartners =
10             findJoinPartners(newGoal, pio);
11
12         // ...
13         // Convert sequents to SE states
14         ImmutableList<SymbolicExecutionState> joinPartnerStates = ImmutableList.nil();
15         for (Pair<Goal, PosInOccurrence> joinPartner : joinPartners) {
16             Triple<Term, Term, Term> partnerSEState =
17                 sequentToSETriple(joinPartner.first, joinPartner.second, services);
18
19             joinPartnerStates = joinPartnerStates.prepend(new SymbolicExecutionState(
20                 partnerSEState.first, partnerSEState.second, joinPartner.first.node());
21         }
22
```

```

23     SymbolicExecutionStateWithProgCnt thisSEState =
24         sequentToSETriple(newGoal, pio, services);
25
26     // The join loop
27     SymbolicExecutionState joinedState = new SymbolicExecutionState(
28         thisSEState.first, thisSEState.second, goal.node());
29
30     for (SymbolicExecutionState state : joinPartnerStates) {
31         joinedState = joinStates(joinedState, state, thisSEState.third, services);
32         joinedState.setCorrespondingNode(goal.node());
33     }
34
35     Term resultPathCondition = joinedState.second;
36     resultPathCondition = trySimplify(services.getProof(), resultPathCondition, true);
37
38     // ...
39     // Close partner goals
40     for (Pair<Goal, PosInOccurrence> joinPartner : joinPartners) {
41         closeJoinPartnerGoal(
42             newGoal.node(),
43             joinPartner.first,
44             joinedState,
45             sequentToSEPair(joinPartner.first, joinPartner.second, services),
46             thisSEState.third);
47     }
48
49     // ...
50     return newGoals;
51 }
52
53 protected SymbolicExecutionState joinStates(
54     SymbolicExecutionState state1,
55     SymbolicExecutionState state2,
56     Term programCounter,
57     Services services) {
58
59     // ...
60     for (LocationVariable v : progVars) {
61         // ...
62
63         Pair<HashSet<Term>, Term> joinedVal =
64             joinValuesInStates(v, state1, rightSide1, state2, rightSide2, services);
65
66         newElementaryUpdates = newElementaryUpdates.prepend(
67             tb.elementary(
68                 v,
69                 joinedVal.second));
70
71         newPathCondition = tb.and(
72             newPathCondition,
73             tb.and(joinedVal.first));
74
75         // ...
76     }
77     // ...
78 }
79
80 // ...
81
82
83 protected abstract Pair<HashSet<Term>, Term> joinValuesInStates(
84     LocationVariable v,
85     SymbolicExecutionState state1,
86     Term valueInState1,
87     SymbolicExecutionState state2,
88     Term valueInState2,
89     Services services);
90
91 // ...
92 }

```

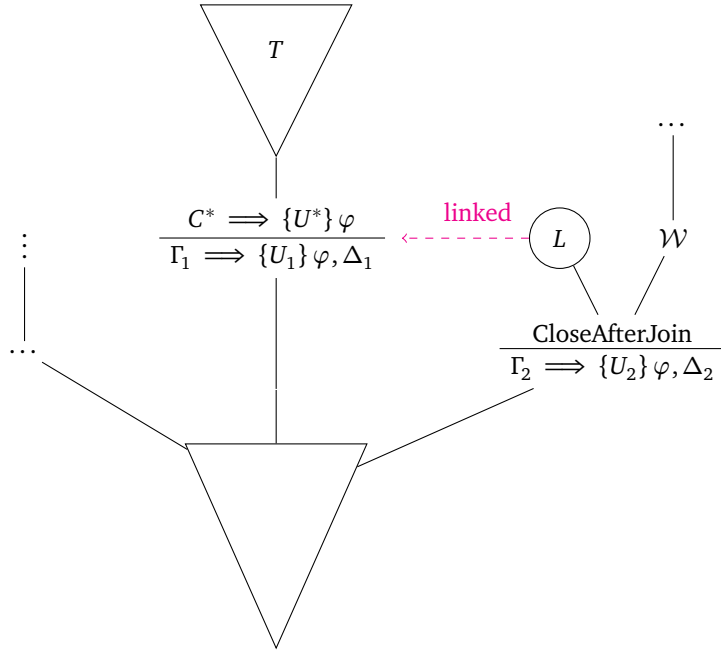


Figure 4.1: Visualization of `CloseAfterJoin` rule. L is the linking node, \mathcal{W} contains the proof obligation corresponding to the logical weakening formula.

4.2 The Abstraction Framework

Our implemented lattice abstraction framework (UML diagram in Figure 6.3) centers around three classes. The abstract class `AbstractDomainElement` represents an element of an abstract lattice domain. Basically, an abstract domain element encapsulates a defining axiom that can be obtained by calling the method `getDefiningAxiom(Term, Services)`. An abstract domain inherits from the class `AbstractDomainLattice`: Most importantly, it has to implement the methods `join(AbstractDomainElement, AbstractDomainElement)` and `iterator()`. The abstract join rule `JoinWithLatticeAbstraction` uses `iterator()` to obtain an ordered enumeration of abstract domain elements; for each element, it checks whether its defining axiom holds for the concrete term that is to be abstracted. The such determined abstract elements are joined with the method `join(AbstractDomainElement, AbstractDomainElement)`. We implemented two abstract domains based on this framework: the sign lattice domain for integers and a simple domain for booleans. Note that this framework is designed to cope with finite and quite small domains only. It probably has to be extended to support more complicated abstract domains (\rightarrow Section 6.1).

4.3 The Partner Goals Closing Rule

As we pointed out in Remark 3.17, the defocusing join rule with two conclusions bears theoretical complications. Since furthermore the tree structure is strongly incorporated in the KeY system, a real DAG representation is definitely difficult to achieve. Instead of applying the weakening rule of Remark 3.17 twice and re-using the computed proof subtree of one of the branches (\rightarrow Figure 3.3), we apply the special rule `CloseAfterJoin` to partner goals of a join operation (UML diagram: Figure 6.4). The rule adds two sub branches to a partner goal. One of these branches contains a single node linking the partner goal to the corresponding join node. We added a field `linkedNode` to the `Node` class for the purpose of establishing this connection. If the join node is closed, an action is triggered that closes all of its linked nodes. The second branch added after the partner goal contains a proof obligation expressing that the join node is logically weaker than its partner node (\rightarrow Definition 3.7). Recall that by Proposition 3.11, logical weakening is equivalent to semantic weakening. Thus, this measure, assuming that the quite simple condition (SEL5) of Definition 3.12 is obeyed, ensures that proofs are not rendered unsound even for join methods with implementation failures, or such methods for which the conformity with the axioms of Definition 3.12 has not been theoretically proven. A visualization of the result of a `CloseAfterJoin` application, based upon Figure 3.3, is shown in Figure 4.1. The node after the sequent

$\Gamma_1 \implies \{U_1\} \varphi, \Delta_1$ is the *join node*, the node corresponding to the sequent $\Gamma_2 \implies \{U_2\} \varphi, \Delta_2$ is the *partner node*. The logical weakening formula is added to the node \mathcal{W} , whereas the node L links to the join node. The subtree T does appear only once and is not, as in Figure 3.3, appended to the partner node.

Listing 4.2 shows the part of the method `CloseAfterJoin.apply(Goal, Services, RuleApp)` which is responsible for linking partner nodes to their join nodes. Whenever the proof tree is extended or pruned, the state of the join node is checked. In the case that the join node has been deleted from the tree, the linked node is marked as independent again. Thus, the proof is still sound, which it would not necessarily be if we did ultimately close partner goals after a join. If the branch of the join node has been closed, the linked goal after the partner node is also being closed.

Listing 4.2: Linking of partner nodes to join nodes

```
services.getProof().addProofTreeListener(new ProofTreeAdapter() {
    @Override
    public void proofGoalsChanged(ProofTreeEvent e) {
        if (joinNode.isClosed()) {
            // The joined node has been closed; now also close this node.
            services.getProof().closeGoal(linkedGoal);
        }
    }

    @Override
    public void proofPruned(ProofTreeEvent e) {
        if (!proofContainsNode(e.getSource(), joinNodeF)) {
            // The joined node has been pruned; now mark this node
            // as not linked and set it to automatic again.
            linkedGoal.node().setLinkedNode(null);
            linkedGoal.setEnabled(true);
        }
    }
});
```

4.4 Macro for Execution until Join Points

For applying join rules on suitable nodes in a proof, it is clearly undesirable to generate proofs completely by hand or to prune automatically generated proofs with the intention to create a situation where joining is possible. In our implementation, we did not establish an inclusion of our join rules into the automatic Java DL strategies; this could be an object to future work. However, we created the macro `FinishSymbolicExecutionUntilJoinPointMacro` (\rightarrow Figure 6.5) which, after started on a sequent containing Java code, allows the execution of a parent strategy only until a previously registered break point has been reached. Whenever a Java block is seen that contains either an if block, a while loop with a break, or a try-catch block, the respective next statement afterward is added to the set of break points. Thus, the macro behaves like the given parent strategy, but might stop the proof procedure at points facilitating an application of join rules. This induces a semi-automatic proof strategy consisting of a repeated execution of the macro interleaved with manual applications of join rules on the generated proof nodes (\rightarrow Section 5.1). The creation of a fully automatic strategy / macro based on this procedure could be a first step into the direction of proof generation with automatic join rule applications in KeY.

Listing 4.3 contains an excerpt of the method `isApprovedApp(RuleApp, PosInOccurrence, Goal)` which decides whether a proposed rule application is permissible or not. In line 4, execution is stopped whenever a breakpoint is reached; in line 10, breakpoints are searched for and registered.

4.5 Sample Join Rule Implementations

We implemented concrete join rules based on our implemented framework and the techniques 3.20, 3.22, 3.25, and 3.30. The techniques for heap joining follow Definition 3.32. As depicted in the UML diagram of Figure 6.1, all join classes depend on `JoinIfThenElse`, in particular on the methods for creating if-then-else terms (\rightarrow Figure 6.6). The classes `JoinWithLatticeAbstraction` and `JoinIfThenElseAntecedent` use if-then-else as fall back for heap joining; `JoinIfThenElseAntecedent` further employs the method `createDistFormAndRightSidesForITEUpd(LocationVariable, ...)` for

Listing 4.3: FinishSymbolicExecutionUntilJoinPointMacro: Registering join points and deciding about continuation of execution

```
1 @Override
2 public boolean isApprovedApp(RuleApp app, PosInOccurrence pio, Goal goal) {
3     if (!hasModality(goal.node())) { return false; }
4     if (hasBreakPoint(goal.sequent().succedent())) { return false; }
5
6     if (pio != null) {
7         JavaBlock theJavaBlock = getJavaBlockRecursive(pio.subTerm());
8
9         //...
10        breakpoints.addAll(findJoinPoints((StatementBlock) theJavaBlock.program()));
11
12        //...
13    }
14
15    return super.isApprovedApp(app, pio, goal);
16 }
```

Listing 4.4: Step two of distinguishing formulae simplification (Part 1)

```
1 public static Term trySimplify(
2     final Proof parentProof, final Term term, boolean countDisjunctions) {
3     try {
4         Term simplified = simplify(parentProof, term);
5
6         if (countAtoms(simplified) < countAtoms(term) &&
7             (countDisjunctions(simplified, false) < countDisjunctions(term, false))) {
8             return simplified;
9         }
10    } catch (ProofInputException e) {}
11
12    return term;
13 }
```

obtaining the shortest distinguishing formula given two path conditions. We briefly elaborate on this simplification and another one for optimizing the structural complexity of disjunctive path conditions.

Simplified Distinguishing Formulae

This is an optimization of Technique 3.22 aimed to increase the proof performance by generating shorter if-then-else terms. Given two SE states (U_1, C_1, φ) and (U_2, C_2, φ) , and assuming that C_1 and C_2 are contradicting, the terms

$$\begin{aligned} & \text{if}(\bigwedge C_1) \text{ then}(\{U_1\} \mathbf{x}) \text{ else}(\{U_2\} \mathbf{x}) \quad \text{and} \\ & \text{if}(\bigwedge C_2) \text{ then}(\{U_2\} \mathbf{x}) \text{ else}(\{U_1\} \mathbf{x}) \end{aligned}$$

are equivalent. The implementation chooses the shortest of those after having applied two steps of simplifications on the path conditions. The first step includes removing all common atoms from the path conditions, since they are not needed for distinguishing the states. In the second step, a side proof is run on the path conditions. If the conjunction of the open goals in the resulting proof is simpler regarding the number of atoms and disjunctions, this conjunction is chosen as distinguishing formula. Otherwise, the result of the first step is returned. The listings 4.4 and 4.5 show the corresponding code for this simplification step.

Optimized Path Conditions

Further optimizations take place in the creation of the disjunctive path condition. Instead of choosing the canonical $\bigwedge C_1 \vee \bigwedge C_2$ for the path condition of a join node, the method `JoinRuleUtils.createSimplifiedDisjunctivePathCondition(Term, Term, Services)` performs a two-step simplification to generate a more concise path condition. In step one, complementary elements

Listing 4.5: Step two of distinguishing formulae simplification (Part 2)

```

1 private static Term simplify(Proof parentProof, Term term)
2     throws ProofInputException {
3     final Services services = parentProof.getServices();
4     final ApplyStrategyInfo info = tryToProve(term, services, true);
5
6     try {
7         // The simplified formula is the conjunction of all open goals
8         ImmutableList<Goal> openGoals = info.getProof().openEnabledGoals();
9         final TermBuilder tb = services.getTermBuilder();
10        if (openGoals.isEmpty()) { return tb.tt(); }
11        else {
12            ImmutableList<Term> goalImplications = ImmutableList.nil();
13            for (Goal goal : openGoals) {
14                Term goalImplication = sequentToFormula(goal.sequent(), services);
15                goalImplications = goalImplications.append(goalImplication);
16            }
17
18            return tb.and(goalImplications);
19        }
20    } finally {
21        SideProofUtil.disposeOrStore(
22            "Simplification_of_"
23            + ProofSaver.printAnything(term,
24            parentProof.getServices()), info);
25    }
26 }

```

of $\bigwedge C_1$ and $\bigwedge C_2$ are removed, i.e. if $C_1 \equiv \psi_1 \wedge \varphi$ and $C_2 \equiv \psi_2 \wedge \neg\varphi$, the result will be $\psi_1 \vee \psi_2$. This simplification is admissible since we assume that $\bigwedge C_1 \leftrightarrow \neg \bigwedge C_2$, and the sequent

$$(\psi_1 \wedge \varphi) \leftrightarrow \neg(\psi_2 \wedge \neg\varphi) \implies ((\psi_1 \wedge \varphi) \vee (\psi_2 \wedge \neg\varphi)) \leftrightarrow (\psi_1 \vee \psi_2)$$

formalizing this simplification is valid (which can be proven by hand or, for instance, in KeY). Since we also assume $\bigwedge C_1 \leftrightarrow \neg \bigwedge C_2$ in the proof of Proposition 3.24, which is justified for the reason that it holds in symbolic execution without joins and our join rules do not violate it, the assumption is also applicable at this place. Intuitively, the removal of contradicting conjuncts in joined path conditions is sensible since the condition that made symbolic execution split is no longer relevant after all of the such generated branches are joined again. The measure of removing it from the resulting path condition by syntactic checks spares KeY having to apply a potentially large number of rules for obtaining the same result. Finally, in the second step, we apply a distributivity law to further simplify the structure of the formula. We decided not to include the automatic simplification described in the above paragraph for optimized path conditions, since our experiments showed that the time overhead induced by this process is out of proportion to the achieved results.

Sign Lattice and Boolean Abstraction

For showing the feasibility of performing an abstract interpretation with KeY by the means of our theoretical and implemented framework, we implemented a join rule [JoinWithSignLattice](#) for joining states based on the sign domain (\rightarrow Example 3.29). The rule itself is very simple (\rightarrow Figure 6.3): It just returns an instance of [SignAnalysisLattice](#) for the Integer sort, and an instance of [BooleanLattice](#) for the Boolean sort. For other sorts, it returns null; in this case, the super class [JoinWithLatticeAbstraction](#) generates an if-then-else expression for the corresponding values. The core of the implementation are the classes realizing the abstract domains and their elements. Figure 6.7 shows the UML diagram for the sign analysis abstract domain lattice, Figure 6.8 for the Boolean domain. We created abstract classes representing the respective domain element types, and one concrete class for each domain element. Note that our framework does not require one class per abstract element; it would also be possible to define a single concrete class the objects of which are capable of remembering the elements they represent and of returning the right defining axioms. Since the domain element classes like [Neg](#) and [Leq](#) are all singletons, there is no substantial difference; this choice is just a matter of style. For extending our framework to more complex abstract domains (see, e.g., [Cou01; Cou+05]), we suppose that our abstraction framework will have to be extended to allow for more flexibility.

5 Evaluation and Case Study

Incorporating joins during symbolic execution can *increase* as well as *decrease* the length of proofs based on chosen parameters for proof generation and the concrete scenario. In Section 5.1, we investigate this impact for a set of example programs. Section 5.2 concerns the benefits of if-then-else joins for increasing the precision of an information flow analysis.

5.1 Performance Evaluation

We evaluate the performance of our implementation using four simple and two more complex Java programs: (1) A method containing an if-then-else block for computing the absolute of an integer (\rightarrow Listing 5.1), (2) a method computing the absolute of an integer inside an object, returning a new object (\rightarrow Listing 5.2), (3) a method computing the quotient of two integers, catching a potential exception arising from a division by zero (\rightarrow Listing 5.3), (4) a method comprising a while loop with a break statement for finding an element in an integer array (\rightarrow Listing 5.4), (5) an integer multiplication method (\rightarrow Listing 5.5) also used, in a slightly different form, in [HSS09], and (6) a method computing the greatest common divisor of two integers (\rightarrow Listing 5.6). The examples (1) to (4) serve as basic representatives for the cases where merges typically may take place, the most obvious amongst those being if-then-else and try-catch blocks. Example (4) allows for a merge of the “break” sub branch in the “preserves invariant” branch with the “use case” branch; those branches arise after an application of the loop invariant rule. Example (2) demonstrates that the implementation is also capable of handling heaps. The contract of Example (5) is proven using loop unwinding; for this reason, we restrict the range of the input values. The chosen bound 5 is rather arbitrary and quite low, but that is just for convenience of the proof; any concrete value at this place would allow for proving the contract with unwinding. The chosen join points, for this example, reside at the beginning of the while loop. Execution is split after the if statement in the body of the loop, therefore an alternative, earlier join point would be at line 16 in the listing. We chose the point two statements later to minimize the number of manual interactions with the proof. The gcd Example (6) involves a normalization before calling a helper method: the input variables are converted into their absolutes in lines 8 and 9 in the listing. In the proof, the method `gcdHelp` is integrated by its method contract.

Table 5.2 shows the experimental results for the example programs under chosen execution strategies. The recorded parameters comprise the number of nodes in the resulting proof of the respective method contract as well as the corresponding numbers of branches (excluding the branches / nodes for the logical weakening check after joins), rule and one step simplifier applications, joins, and the necessary number of manual user interactions for the proof. In the last column, the improvement of the respective strategy over the baseline strategy (fully automatic proof with/without one step simplification) is listed. Improvements of at least 2% are highlighted in **green**, declines of at least 2% in **red**; small changes are of gray color. The employed strategies are explained in Table 5.3. The “canonical” strategies, `join-ite`, `join-ite2` and `join-sign`, involve a repeated symbolic execution until the next join point with the macro we developed for this purpose (\rightarrow Section 4.3), followed by a join of applicable leaf nodes; it should be possible to derive fully automatic strategies from this procedure in a quite straightforward manner. For the multiplication and gcd examples, we make use of specialized, less canonical strategies that are better suited for those two examples.

For the simple examples (1) to (4) with activated one step simplifier, joining branches with if-then-else techniques is clearly not beneficial. Those examples highly benefit from one step simplifier applications which are not applicable in that way after a join operation. Consider the abs example (2) with the `join-ite2` strategy: In this case, the proof with join is about 23% longer than the corresponding automatic proof. The proof splits into two branches corresponding to the sign of the input variable. In the proof with join, the “`num \geq 0`” branch is 11 nodes shorter. However, the “`num $<$ 0`” branch in the automatic proof can be closed quite quickly due to some very efficient one step simplifier applications: The goal $\Delta, \text{num} \leq -1 \implies \Gamma, \text{num} * -1 \geq 0$ in the automatic proof corresponds to the more complicated goal $\Delta', \text{num} \leq -1 \rightarrow \text{result_1} = \text{num} * -1, \text{num} \leq -1 \vee \text{result_1} = \text{num} \implies \Gamma, \text{result_1} \geq 0$ in the proof with join. Therefore, the automatic proof can be closed in one step after the usage of the class invariant axiom for SimpleMath, whereas its counterpart with join needs another 37 steps including a cut over the sign of num. These observations also explain why joining with if-then-else methods is less disadvantageous, and sometimes even beneficial, for the very same examples with deactivated one step

Example	Strategy	# Nodes / # Branches	# Rule Apps / # OSS Apps	# Joins	# User Inter- actions	Improvement of # Nodes to baseline
div	auto	185 / 9	415 / 52	0	0	–
div	join-ite	196 / 9	394 / 51	1	1	-5.95%
div	join-ite2	193 / 8	394 / 51	1	1	-4.32%
div	auto-no-OSS	401 / 9	400 / 0	0	0	–
div	join-ite-no-OSS	380 / 9	379 / 0	1	1	5.24%
div	join-ite2-no-OSS	374 / 6	374 / 0	1	1	6.73%
abs	auto	78 / 5	155 / 18	0	0	–
abs	join-ite	79 / 5	136 / 16	1	1	-1.28%
abs	join-ite2	96 / 5	153 / 21	1	1	-23.08%
abs	join-sign	63 / 4	122 / 14	1	1	19.23%
abs	auto-no-OSS	158 / 5	157 / 0	0	0	–
abs	join-ite-no-OSS	138 / 5	137 / 0	1	1	12.66%
abs	join-ite2-no-OSS	155 / 5	154 / 0	1	1	1.90%
absObj	auto	211 / 10	451 / 52	0	0	–
absObj	join-ite	219 / 11	418 / 45	5	2	-3.79%
absObj	join-ite2	249 / 11	448 / 51	5	2	-18.01%
absObj	join-sign	190 / 9	389 / 43	5	2	9.95%
absObj	auto-no-OSS	444 / 10	443 / 0	0	0	–
absObj	join-ite-no-OSS	396 / 11	395 / 0	5	2	10.81%
absObj	join-ite2-no-OSS	426 / 11	425 / 0	5	2	4.05%
find	auto	695 / 21	1496 / 114	0	0	–
find	join-ite	702 / 19	1466 / 107	1	1	-1.01%
find	join-ite2	702 / 19	1466 / 107	1	1	-1.01%
find	auto-no-OSS	1438 / 21	1437 / 0	0	0	–
find	join-ite-no-OSS	1427 / 19	1426 / 0	1	1	0.76%
find	join-ite2-no-OSS	1421 / 19	1420 / 0	1	1	1.18%
multiply	auto	1062 / 28	1576 / 189	0	0	–
multiply	mult1	888 / 25	1284 / 158	1	2	16.38%
multiply	mult1-ite2	915 / 25	1308 / 158	1	2	13.84%
multiply	mult2	(timeout)		1	2	N/A
multiply	mult3	1044 / 26	1426 / 178	2	4	1.69%
gcd	auto	8371 / 37	10312 / 675	0	0	–
gcd	gcd-ite	8048 / 32	8923 / 609	2	2	3.86%
gcd	gcd-ite2	7315 / 31	8084 / 624	2	2	12.61%
gcd	auto-no-OSS	9103 / 34	9102 / 0	0	0	–
gcd	gcd-ite2-no-OSS	7197 / 29	7196 / 0	2	2	20.94%

Table 5.2: Performance evaluation for the example programs div, abs, absObj, find, multiply, and gcd

Strategy Name	Description
auto	Full usage of KeY's automated proof search function.
join-ite	Repeatedly run symbolic execution until join point using the <code>FinishSymbolicExecutionUntilJoinPointMacro</code> and join applicable branches by if-then-else until symbolic execution is finished; then proceed automatically.
join-ite2	Like join-ite, but with joins using the if-then-else-antecedent rule.
join-sign	Like join-ite, but with joins using sign lattice abstraction.
mult1	Deactivate loop treatment and start automatic strategy. Unwind the while loop once, and use "Full Auto Pilot" macro (applied on the sequent, with deactivated loop treatment). Two occurrences of while loops occur. Join the two branches by if-then-else join, automatically proceed from there.
mult1-ite2	Like mult1, but with joins using the if-then-else-antecedent rule.
mult2	Like mult1, but with using the automatic strategy instead of the "Full Auto Pilot" macro.
mult3	Like mult1, but with one additional unwinding and join step.
gcd-ite	Two times execution until join point with <code>FinishSymbolicExecutionUntilJoinPointMacro</code> followed by an if-then-else join, then start of automatic proof engine. The method <code>gcdHelp</code> is included by contract.
gcd-ite2	Like gcd-ite, but with joins using the if-then-else-antecedent rule.
-no-OSS	Strategy "" with deactivated <code>OneStepSimplifier</code> .

Table 5.3: Strategies used in performance evaluation.

simplification. An investigation of the abs example without one step simplification shows that in the " $\text{num} \geq 0$ " branch, 31 nodes are saved, while in the " $\text{num} < 0$ " branch the additional nodes for case distinctions are furthermore partly compensated due to some rule applications that previously appeared redundantly. Another reason for the negative results of the simple examples is the late occurrence of the join. When joining directly before the last statement, there is not much code left for which a redundant execution can be avoided; the price for the more complicated expressions after the join cannot be compensated by saving redundancy.

The more extensive examples (5) and (6) show that it is possible to achieve shorter proofs even with the general if-then-else / if-then-else-antecedent join rules and usage of one step simplification. The gcd example is furthermore the only case where if-then-else-antecedent shows better results than the if-then-else method. The largest improvement in our measurements for the if-then-else methods could be measured in the gcd example (6) for the difference between automatic execution without one step simplifier and the gcd-ite2-no-OSS strategy: An improvement of 20.94% in the number of nodes of the proofs. This even exceeds the abs example (2) with the specialized join method by sign lattice abstraction.

For the multiply example (5), we discovered that joins with if-then-else might disturb the automatic strategies of KeY such that the proof procedure does not terminate; KeY was, in the case of the strategy mult2 making use of the "Play" button starting automatic execution, not able to derive the terminating condition for the while loop, resulting in endless unwinding operations. However, when using the macro "Full Auto Pilot" on the whole sequent, the proof did indeed terminate, in this case even with an improvement over the number of nodes and branches when performing one join operation. We discovered that the sequent before the join is more complicated when using the automatic strategy than when using the "Full Auto Pilot" macro. A closer evaluation of this problem might be in the scope of future work.

5.2 Case Study: Information Flow Analysis

Information flow between variables in a program is the transmission of some kind of information from one program variable to another one. Certain types of flows may be undesirable, an observation that motivates imposing *information flow policies* on programs. For instance, information about a secret password should not be leaked to a program variable that is communicated to the outside, maybe over a network connection. We investigate a simple class of policies that involve a partitioning of

Listing 5.1: The abs example – If block

```
1 /*@ public normal_behavior
2   @ ensures \result >= 0;
3   @*/
4 public int abs(int num) {
5     int result;
6
7     if (num < 0) {
8         result = -num;
9     } else {
10        result = num;
11    }
12
13    return result; // join point
14 }
```

Listing 5.2: The absObject example – Join with heaps

```
1 /*@ public normal_behavior
2   @ ensures \result.num >= 0;
3   @*/
4 public SimpleMath absObject() {
5     SimpleMath result = new SimpleMath();
6
7     if (num < 0) {
8         result.num = -num;
9     } else {
10        result.num = num;
11    }
12
13    return result; // join point
14 }
```

Listing 5.3: The div example – Exceptional control flow

```
1 /*@ public normal_behavior
2   @ ensures divisor != 0 ==> \result == dividend / divisor;
3   @ ensures divisor == 0 ==> \result == Integer.MAX_VALUE;
4   @*/
5 public int div(int dividend, int divisor) {
6     int result;
7
8     try {
9         result = dividend / divisor;
10    } catch (ArithmeticException e) {
11        result = Integer.MAX_VALUE;
12    }
13
14    return result; // join point
15 }
```

Listing 5.4: The find example – While loop with break statement

```
1  /*@ public normal_behavior
2  @
3  @ ensures ((\exists int i; i >= 0 && i < arr.length; arr[i] == toFind)
4  @          ==> arr[\result] == toFind) &&
5  @          ((\forall int i; i >= 0 && i < arr.length; arr[i] != toFind)
6  @          ==> \result == -1);
7  @ assignable \nothing;
8  @*/
9  public int find(int[] arr, int toFind) {
10     int curPos = 0;
11
12     /*@ loop_invariant
13     @ (\forall int i; i >= 0 && i < curPos; arr[i] != toFind) &&
14     @   curPos >= 0;
15     @
16     @ decreases arr.length - curPos;
17     @ assignable curPos;
18     @*/
19     while (curPos < arr.length) {
20         if (arr[curPos] == toFind) {
21             break;
22         }
23
24         curPos = curPos + 1;
25     }
26
27     if (curPos < arr.length) { // join point
28         return curPos;
29     } else {
30         return -1;
31     }
32 }
```

Listing 5.5: The multiply example – Loop unrolling

```
1  /*@ public normal_behavior
2  @ requires x0 >= 0 && y0 >= 0;
3  @ requires x0 < 5 && y0 < 5; // bounds for successful unwinding
4  @ ensures \result == x0 * y0;
5  @*/
6  int multiply(int x0, int y0) {
7     int x = x0;
8     int y = y0;
9     int z = 0;
10
11     while (x != 0) { // join point
12         if (x % 2 != 0) {
13             z += y;
14         }
15
16         x = x / 2;
17         y = y * 2;
18     }
19
20     return z;
21 }
```

Listing 5.6: The gcd example – Normalization before method call

```
1 /*@ public normal_behavior
2   @ ensures (a != 0 || b != 0) ==>
3   @       (a % \result == 0 && b % \result == 0 &&
4   @       (\forall int x; x > 0 && a % x == 0 && b % x == 0;
5   @       \result % x == 0));
6   @*/
7 public static int gcd(int a, int b) {
8   if (a < 0) a = -a;
9   if (b < 0) b = -b; // join point
10
11   int big, small; // join point
12   if (a > b) {
13     big = a;
14     small = b;
15   } else {
16     big = b;
17     small = a;
18   }
19
20   return gcdHelp(big, small);
21 }
```

program variables into two classes “high” and “low”, disallowing a flow from variables classified “high” to variables classified “low”, such that a “low” observer cannot infer information about the initial values of the “high” variables.

A common technique for a language-based analysis of information flow is based on specialized type systems [SM06]. Such systems should necessarily be sound, i.e. they never classify insecure programs as secure. However, they will classify certain secure programs as insecure, due to a lack of sensitivity to the particular control-flow of a program. We subsequently propose a quite simple information flow analysis technique based on symbolic execution which is sound for our chosen examples, but classifies some secure examples as insecure. We then show how the if-then-else join method can be utilized to achieve a classification as “secure” for those intuitively secure examples. Our goal is not to devise a general sound and precise analysis technique, but to show how the precision of an example analysis can be improved using join methods. We believe that it should be possible to similarly extend existing approaches for information flow analysis based on symbolic execution, thereby retaining the soundness of these systems and increasing their precision.

Algorithm 1 Information flow analysis without joins

Require: p is a Java program with program variables \bar{h} classified as “high” and \bar{l} classified as “low”;
 $P(\dots)$ is a fresh predicate accepting \bar{l} as inputs.

Ensure: Returns true only if p does not contain a flow from a variable in \bar{h} to a variable in \bar{l} .

```
1:  $problem \leftarrow \langle \{ \text{try } \{ p \} \text{ catch (Exception } e) \{ \} \} \rangle P(\bar{l})$ 
2: Load  $problem$  into KeY and finish symbolic execution
3:  $G \leftarrow$  all leafs in the resulting SET containing  $P$ 
4: for all  $g \in G$  do
5:   Apply as many OneStepSimplifier applications on  $g$  as possible
6:   Transform  $g$  to the form  $\Gamma \implies P(\dots)$  by “negation right” applications
7:   if  $\Gamma$  contains an atomic formula containing a variable in  $\bar{h}$  or  
      $P(\dots)$  contains a term comprising a variable in  $\bar{h}$  then
8:     return false
9:   end if
10: end for
11:
12: return true
```

We investigate the behavior of Algorithm 1 based on five examples, where `l`, `res` are “low” variables and `h`, `secret` are “high” variables.

Listing 5.7: "Exceptional Flow" Example

```
public class ExcptFlow {  
    private /*@ spec_public @*/ int secret;  
  
    public ExcptFlow(int secret) {  
        this.secret = secret;  
    }  
  
    /* Roughly following "RIFLE: An Architectural  
       Framework for User-Centric Information-Flow Security"  
       by Vachharajani et al. */  
    public boolean insecureExceptional(int input) {  
        if (input == secret) {  
            work();  
        }  
  
        return true;  
    }  
  
    public boolean secureExceptional(int input) {  
        if (input == secret) {  
            try {  
                work();  
            } catch (Exception e) {}  
        }  
  
        return true;  
    }  
  
    /*@ public exceptional_behavior  
       @ signals_only RuntimeException;  
       @*/  
    public void work() {  
        throw new RuntimeException("Crashed.");  
    }  
}
```

- (1) `if (h>0) {l=1;} else {l=2;}.` This program is *insecure* since an attacker is able to derive from the final value of `l` whether the initial value of `h` was positive or not. Algorithm 1 outputs *false* for the program since the final goals contain formulae $h \geq 1$ and $h \leq 0$, respectively.
- (2) `if (h>0) {l=1;} else {l=2;} l=0;.` This program is *secure* since an attacker cannot draw any conclusions about the initial value of `h` since the final value of `l` is always 0. However, Algorithm 1 outputs *false* for the program since the final goals contain formulae $h \geq 1$ and $h \leq 0$, respectively, as in example (1).
- (3) `if (h>0) {l=2; h=1;} else {l=2; h=2;}.` This program is *secure* since an attacker cannot draw any conclusions about the initial value of `h` since the final value of `l` is always 2. However, Algorithm 1 outputs *false* for the program since the final goals contain formulae $h \geq 1$ and $h \leq 0$, respectively, as in example (1).
- (4) `ExcptFlow a = new ExcptFlow(secret); res = a.secureExceptional(input);` where the class `ExcptFlow` is defined according to Listing 5.7. The program is *secure* since the result of the function call is always *true*. However, Algorithm 1 outputs *false* for the program since the final goals contain formulae $secret = input$ and $!secret = input$.
- (5) `ExcptFlow a = new ExcptFlow(secret); res = a.insecureExceptional(input);` where the class `ExcptFlow` is defined according to Listing 5.7. The program is *insecure* since the result of the function call is *true* iff the input equals the secret. Algorithm 1 outputs *false* for the program since the final goals contain formulae $secret = input$ and $!secret = input$.

The examples show that Algorithm 1 is too restrictive. It would return *true* only for examples like $l=0$; that are very simple concerning their control-flow. We therefore refine the algorithm to Algorithm 2 using the if-then-else join technique.

Algorithm 2 Information flow analysis with joins

Require: p is a Java program with program variables \bar{h} classified as “high” and \bar{l} classified as “low”;
 $P(\dots)$ is a fresh predicate accepting \bar{l} as inputs.

Ensure: Returns true only if p does not contain a flow from a variable in \bar{h} to a variable in \bar{l} .

```
1:  $problem \leftarrow \langle \{ \text{try } \{ p \} \text{ catch (Exception } e) \{ \} \} \rangle P(\bar{l})$ 
2: Load  $problem$  into KeY and finish symbolic execution
3:  $G \leftarrow$  all leaves in the resulting SET containing  $P$ 
4: Apply, if possible, “if-then-else” join on  $P(\dots)$ , if necessary after OneStepSimplifier applications
5: for all  $g \in G$  do
6:   Apply as many OneStepSimplifier applications on  $g$  as possible
7:   Transform  $g$  to the form  $\Gamma \implies P(\dots)$  by “negation right” applications
8:   if  $\Gamma$  contains an atomic formula containing a variable in  $\bar{h}$  or
      $P(\dots)$  contains a term comprising a variable in  $\bar{h}$  then
9:     return false
10:  end if
11: end for
12:
13: return true
```

The only change in Algorithm 2 is the addition of line 4, where all branches containing the $P(\dots)$ formulae are joined. However, this change is strong enough to increase the precision of the results for our “secure” examples, while still classifying “insecure” examples correctly:

- (1) Algorithm 2 outputs *false* for the program since the final goals contain formulae $h \geq 1$ and $h \leq 0$, respectively.
- (2) Algorithm 2 outputs *true* for the program since the constraints $h \geq 1$ and $h \leq 0$ could be eliminated during the join.
- (3) Algorithm 2 outputs *true* for the program since the constraints $h \geq 1$ and $h \leq 0$ could be eliminated during the join.
- (4) Algorithm 2 outputs *true* for the program since the final goals no longer contain atomic formulae $secret = input$ and $!secret = input$. Instead, those atoms occur in sub formulae of a disjunction in the antecedent; it is not possible to directly derive information on the initial value of $secret$.
- (5) Algorithm 2 outputs *false* for the program since the final goals contain formulae $secret = input$ and $!secret = input$. After the join, the succedent of the join node consists of the interesting formula

$$\{result := \text{if}(\neg(secret = input)) \text{then} (true) \text{else} (result)\} P(result)$$

which concisely reflects the dependency between $secret$ and $result$.

6 Related Work and Conclusion

6.1 Related Work

This thesis may be seen as a contribution to research on the path explosion problem of symbolic execution as well as in relation to other static analysis techniques like abstract interpretation. Furthermore, joins in KeY proof trees might motivate the study of different kinds of information flow analyses based on the KeY system.

The Path Explosion Problem of Symbolic Execution

Existing work on the path explosion problem can be divided into two distinct lines of research. A popular approach, especially in the area of automatic test generation and bug detection, is the pruning of redundant paths in symbolic execution trees by subsumption checking (e.g. [APV06; BCE08; Jaf+12; JMN13; CJM14]). The second line of research applies state merging using if-then-else techniques (e.g. [HSS09; Kuz+12; Sen+14]).

Pruning by Subsumption Checks

When utilizing symbolic execution for the automatic generation of test cases, the repeated exploration of the same program statements is undesirable. A concept employed for subsumption checking that is used, for instance, in [Jaf+12; JMN13; CJM14], is the computation of *interpolants*: Given two formulae A and B (here: two path conditions) s.th. $A \wedge B$ is unsatisfiable, a Craig interpolant Ψ is a formula for which it holds that (1) $A \models \Psi$, (2) $\Psi \wedge B$ is unsatisfiable, and (3) all variables in Ψ are common to A and B . The Tracer system [Jaf+12] systematically underapproximates interpolants and stops the exploration of paths the path conditions of which are subsumed by the interpolants of previously explored paths. Chu et al. [CJM14] compute interpolants and continue the execution even of infeasible paths to obtain better interpolants which they also apply in subsumption checking. The RWset system [BCE08] prunes redundant paths by checking whether the remainder of a particular execution is capable of exploring new behavior. A work by Anand et al. [APV06] especially addresses subsumption in symbolic execution trees containing objects on a heap. They abstract from heap objects (in particular lists and arrays) by summarizing elements sharing common properties using shape analysis techniques and perform subsumption checking for the such abstracted objects. Their approach follows a trend in software model checking which proposes underapproximation based abstractions for the purpose of falsification. Since the intended scenario for our work is software verification, and thus proving the correctness (validity) of properties, underapproximation is not an option.

If-Then-Else State Merging

Existing approaches for joining states in symbolic execution are usually based upon if-then-else techniques similar to the techniques 3.22 and 3.25. Hansen et al. [HSS09] devise an algorithm that symbolically executes a program according to its control flow graph, thereby joining states with the same program counter using an if-then-else construct to create the new symbolic state and a disjunction of path conditions as new path condition for the join node. They use external tools to simplify the generated path conditions. An evaluation of their approach for a set of four examples shows mixed results: While joining is beneficial for three out of four examples, the execution time explodes for the fourth one. Kuznetsov et al. [Kuz+12] tackle this problem by investigating how state joining based on if-then-else constructs can be rendered “practical”: Their approach aims to automatically find an advantageous balance between exploring fewer complex (i.e., merged) states vs. more simpler (i.e. unmerged) states. Based on heuristics, they only merge states when this promises to reduce the exploration time. The central idea of the MultiSE system by Sen et al. [Sen+14] is the concept of so-called “value summaries” that map program variables to sets of guarded symbolic expressions. Value summaries are comparable to the if-then-else expressions of KeY, but without the else part; instead of this, another guarded expression is added to the summary. MultiSE explicitly represents the data structure induced by symbolic execution as a DAG. Join points are identified dynamically, the system does not rely on an explicit CFG of the program. For efficiency reasons, path conditions and guards in value summaries are implemented using binary decision diagrams. Experiments show significant speedup of this technique compared to standard (dynamic) SE in their experiments. Compared to our framework, all these systems / approaches

incorporate one chosen join technique based on an if-then-else construct, while our work permits a huge set of overapproximation techniques, also covering those based on if-then-else.

Abstract Interpretation

Abstract interpretation is a static analysis technique introduced in the 1970s by Cousot and Cousot [CC77]. The analysis is parametric in the chosen abstraction and proceeds fully automatically by the means of a fixpoint iteration for loops. State merging, following the CFG of the analyzed program, is inherent to the system. With ASTRÉE [Cou+05], there exists a mature system for the runtime error detection in C programs based on abstract interpretation which has been successfully applied in industry. Numerous abstract domains have been implemented for ASTRÉE. The system is sound and therefore necessarily incomplete, that is it may yield false alarms. Bubel et al. [BHW09] propose a dynamic logic with abstraction for the KeY calculus, mainly to facilitate the fully automatic discovery of loop invariants. The idea is based on the employment of a fixpoint algorithm in the spirit of abstract interpretation until an invariant is found. Their framework is based on partially interpreted constant symbols for each abstract value, accompanied by dedicated calculus rules for interpreting them. They use a join rule for while loops in side branches for the computation of invariants. Our work generalizes the abstraction and join aspects of Bubel et al. by devising a parametric framework with a join rule for arbitrary program statements. However, we do not cover the automatic discovery of loop invariants by fixpoint iteration. In contrast to abstract interpretation, our work supports fully precise (if-then-else) as well as abstraction-based state merging.

Information Flow Security

The problem of information flow security is the enforcement or assessment of information flow policies restricting the flow of information between locations in a program. Traditional techniques for the (language-based) enforcement of such policies are based on security type systems (see [SM06] for an overview). Those systems usually suffer from a lack of precision, that is they may classify many secure programs as insecure; however, they can reach full automation. The previously mentioned work by Bubel et al. [BHW09] tracks dependencies between variables to assess the validity of non-interference properties. Their employment of abstraction promises a higher degree of automation, but is likely to decrease the precision compared to deductive systems without abstraction (like, e.g., [BDR04; DHS05; JH14]). However, the approach is more precise than typical security type systems. A different line of research [SRK06; HKS06] establishes information flow security by analyzing “path conditions” computed from the *Program Dependence Graph (PDG)* [FOW84; OO84] of the investigated program. A path condition in this context is a generalization of the corresponding notion in symbolic execution: It is computed between two arbitrary basic blocks of a program by accumulating control and data conditions on all possible paths between the blocks. In symbolic execution, a path condition is computed from the root node on, and mostly captures control dependencies. Snelting et al. [SRK06] establish a straightforward theorem connecting PDGs to non-interference; Hammer et al. [HKS06] apply this result to medium-sized Java programs in a fully automatic approach. Generalized path conditions are precise necessary conditions for information flow between two program points. Our simple application to information flow security (\rightarrow Section 5.2) also takes path conditions (in the sense of SE) into account; it might be interesting to study the relation between SE path conditions and the generalized concept based on PDGs by Hammer et al. Definitely, KeY proof trees come closer to the CFG and therefore also to the PDG of programs by incorporating our join techniques.

6.2 Conclusion and Future Work

Conclusion

We extended symbolic execution by a parametric join rule which is sound for all join operations satisfying, besides three basic semilattice properties, an additional correctness property and a simple constraint on the generated path conditions. Using our join rule, symbolic execution trees are rendered into directed acyclic graphs. Thereby, we contribute to closing the gap between symbolic execution and abstract interpretation and to solving the path explosion problem of symbolic execution. Four example join rules based upon the general lattice framework, including two if-then-else based methods and one method based on abstract domains, are theoretically defined and practically implemented for the state-of-the-art deductive verification system KeY. We showed that the employment of join rules indeed reduces the number of nodes and branches in KeY proof trees for a set of example programs. Furthermore, we demonstrated that branch joining could be beneficial for analysis techniques in the area of

information flow security. We are confident that our work can support making symbolic execution more efficient, and give rise to new analysis methods based on symbolic execution DAGs with suitable join techniques. In addition to evaluating our join techniques on larger case studies, we identified four lines of future work for extending and improving our system.

Performance of the Implementation

We expect that there is a significant potential for improving the performance of our implementation. On the theoretical side, it might be rewarding to investigate whether Craig-Interpolants are employable for generating simplified distinguishing formulae in if-then-else constructs. The generation of those invariants is only easy for quantifier-free formulae [McM05; Bri+11]. Preconditions in sequents often contain quantifiers arising from method specifications, while “pure” path conditions usually are quantifier-free. If applicable, Craig-Interpolants could serve as efficiently computable and concise distinguishing formulae. Other implemented optimization methods, e.g. for the computation of optimized path conditions, also are likely to be improvable in terms of effectiveness and efficiency. In the case of abstraction-based join methods, the performance of validity checks for defining axioms of abstract domain values might be improvable by outsourcing the evaluation to a third-party Satisfiability Modulo Theories (SMT) solver. The last point concerning performance addresses the non-termination problem for the “multiply” example in Section 5.1. The question of why the results of the automatic strategy and the (automatic) macro differ and why this difference actually caused the non-determination issue should be studied closer to fix the problem and possibly generalize the gained experience to different scenarios.

Abstraction Framework

To catch up with abstract interpretation systems like ASTRÉE [Cou+05], many more abstract domains will have to be realized for KeY. For some of those, it might be necessary to extend the flexibility of our abstraction framework which for now is built to cope with quite small, static domains. For some applications, the implementation of abstractions specific to certain heap structures like lists or arrays (cf., e.g., [SRW02; APV06]) could be beneficial; so far, we only apply abstractions for integers or booleans to the elements of suitably uniformly shaped heaps. One of the strengths of abstract interpretation systems, the fixpoint iteration for unbounded loops, would, when implemented as a set of new rules or strategies for the KeY system, push forward the abstraction framework very much and help to further close the gap between symbolic execution and abstract interpretation. The work by Bubel et al. [BHW09] could serve as a guideline here.

Usability

The major starting point for making our implemented framework more usable is the creation of new (or the extension of existing) automatic strategies that are aware of join rules and join points. A very first and simple idea is the repeated execution of our macro (\rightarrow Section 4.4). However, a solution integrated into the main automatic strategies that identifies join points dynamically, such as [Sen+14], would be desirable. Such a solution could also take into account metrics for assessing the reward of a join operation, as in the approach by [Kuz+12]. A more technical improvement concerns the integration of the logical weakening proof goals into proofs with joins. At the moment, we append two nodes to partner nodes participating in a join operation (\rightarrow Section 4.3): One node linking to the join node, and one containing the logical weakening formula. As an alternative, we plan to add this side proof goal as a “contract target” to the meta information of a proof, in the fashion of the current treatment of method contracts. Thus, the size of a proof containing join rule applications would not be blown up by the proofs corresponding to logical weakening; instead, the validity of the main proof goal would be marked as relative to the validity of the logical weakening formulae.

New Analysis Techniques

We assume that the representation of symbolic execution as a DAG, being closer to the actual control flow of the considered programs than the tree representation, opens up symbolic execution to different kinds of static analysis techniques. In Section 5.2, we briefly studied the area of information flow analysis as one potential candidate for a class of such techniques. A closer investigation of the benefits of branch joining for an existing, provably sound information flow analysis based on symbolic execution trees would be interesting. Furthermore, we endorse the study of approaches based on generalized path conditions [SRK06; HKS06] with respect to their application on SE DAGs.

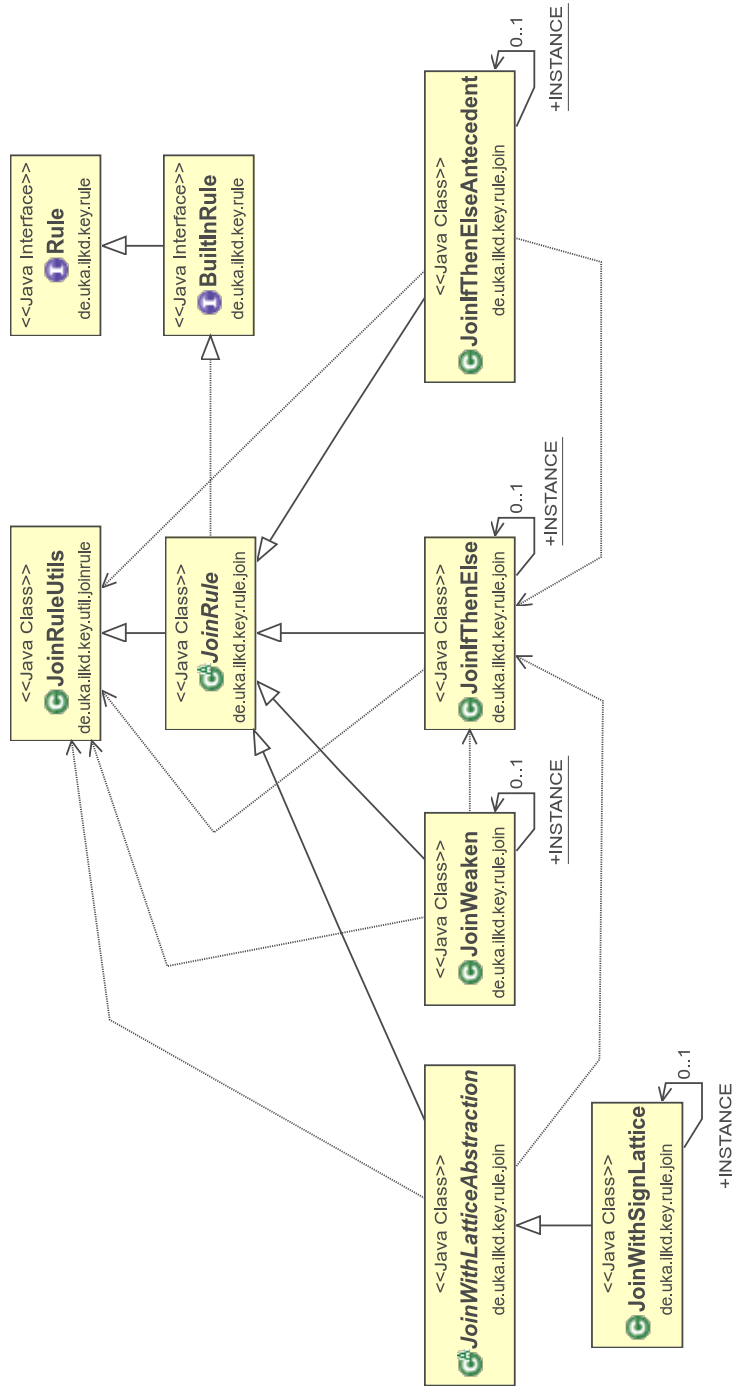


Figure 6.1: Overview of join classes

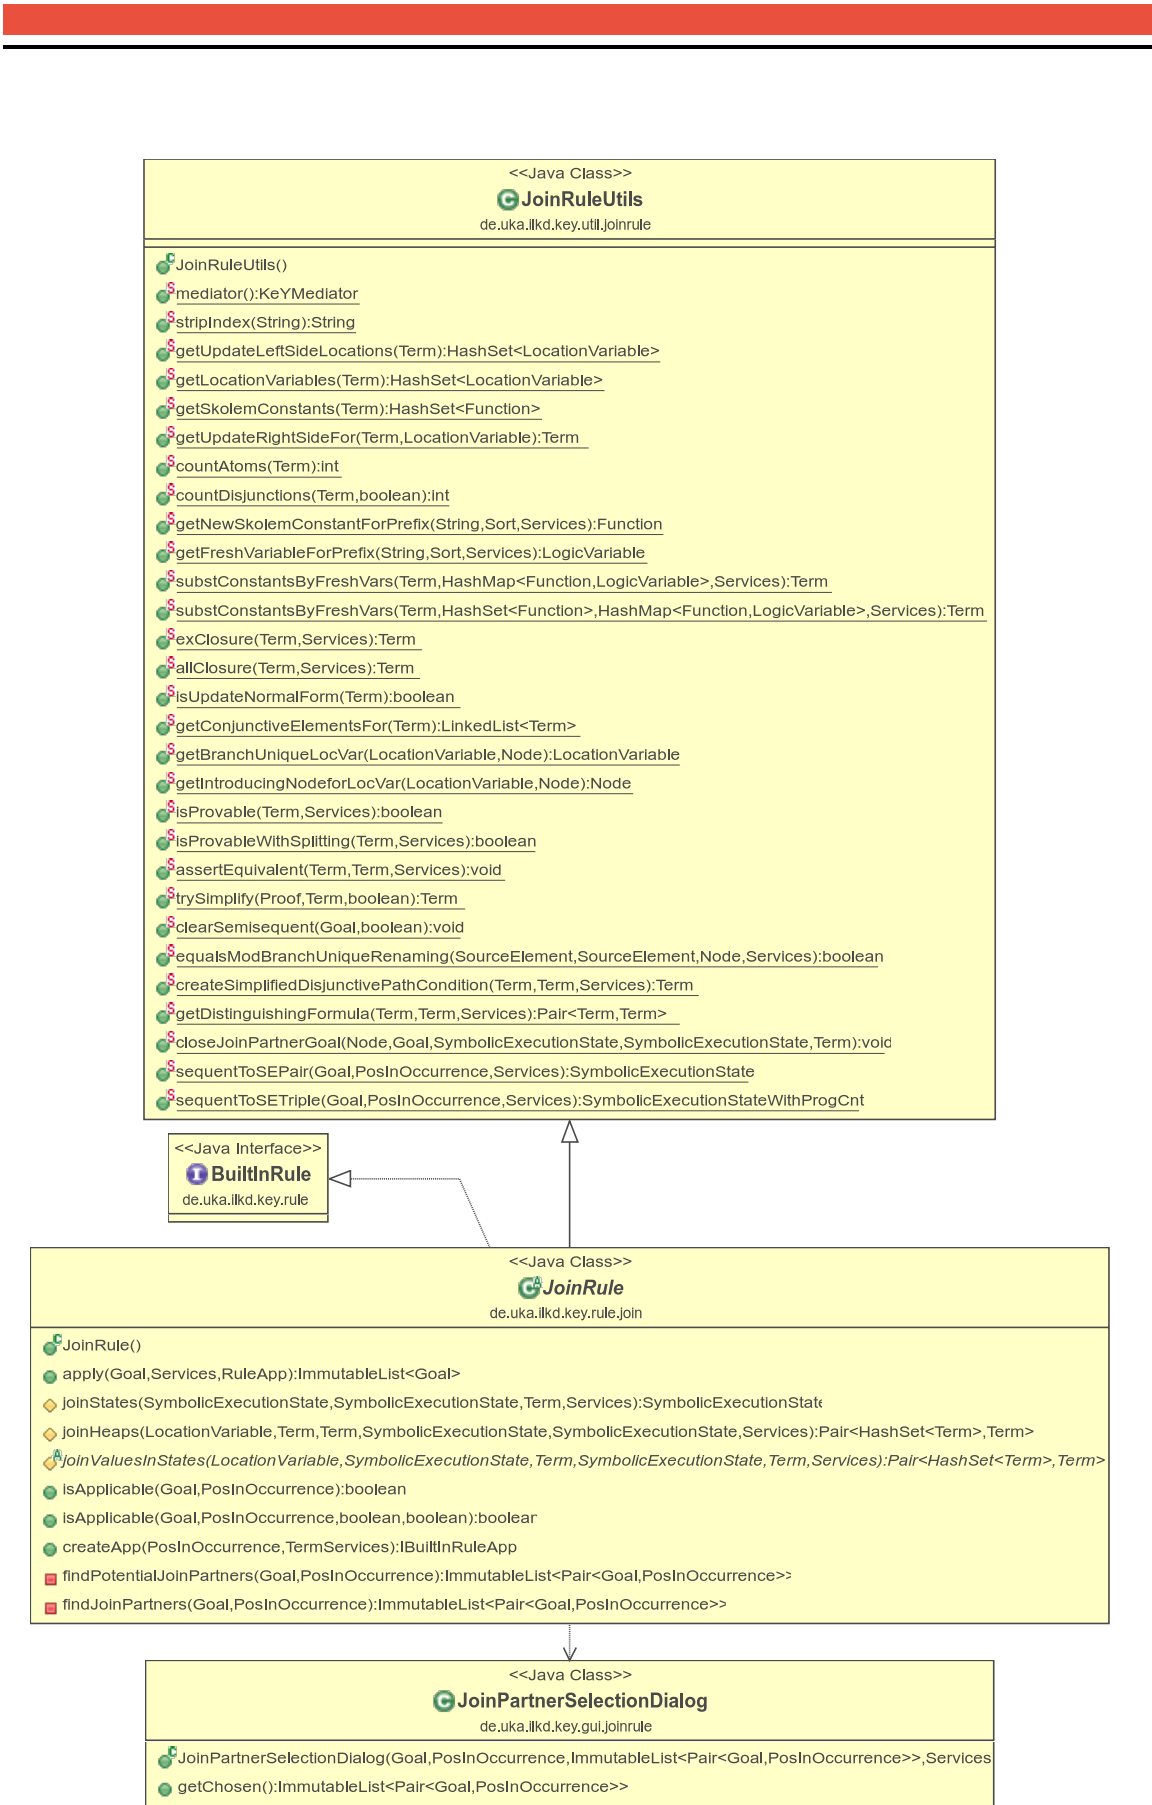


Figure 6.2: Join rule implementation

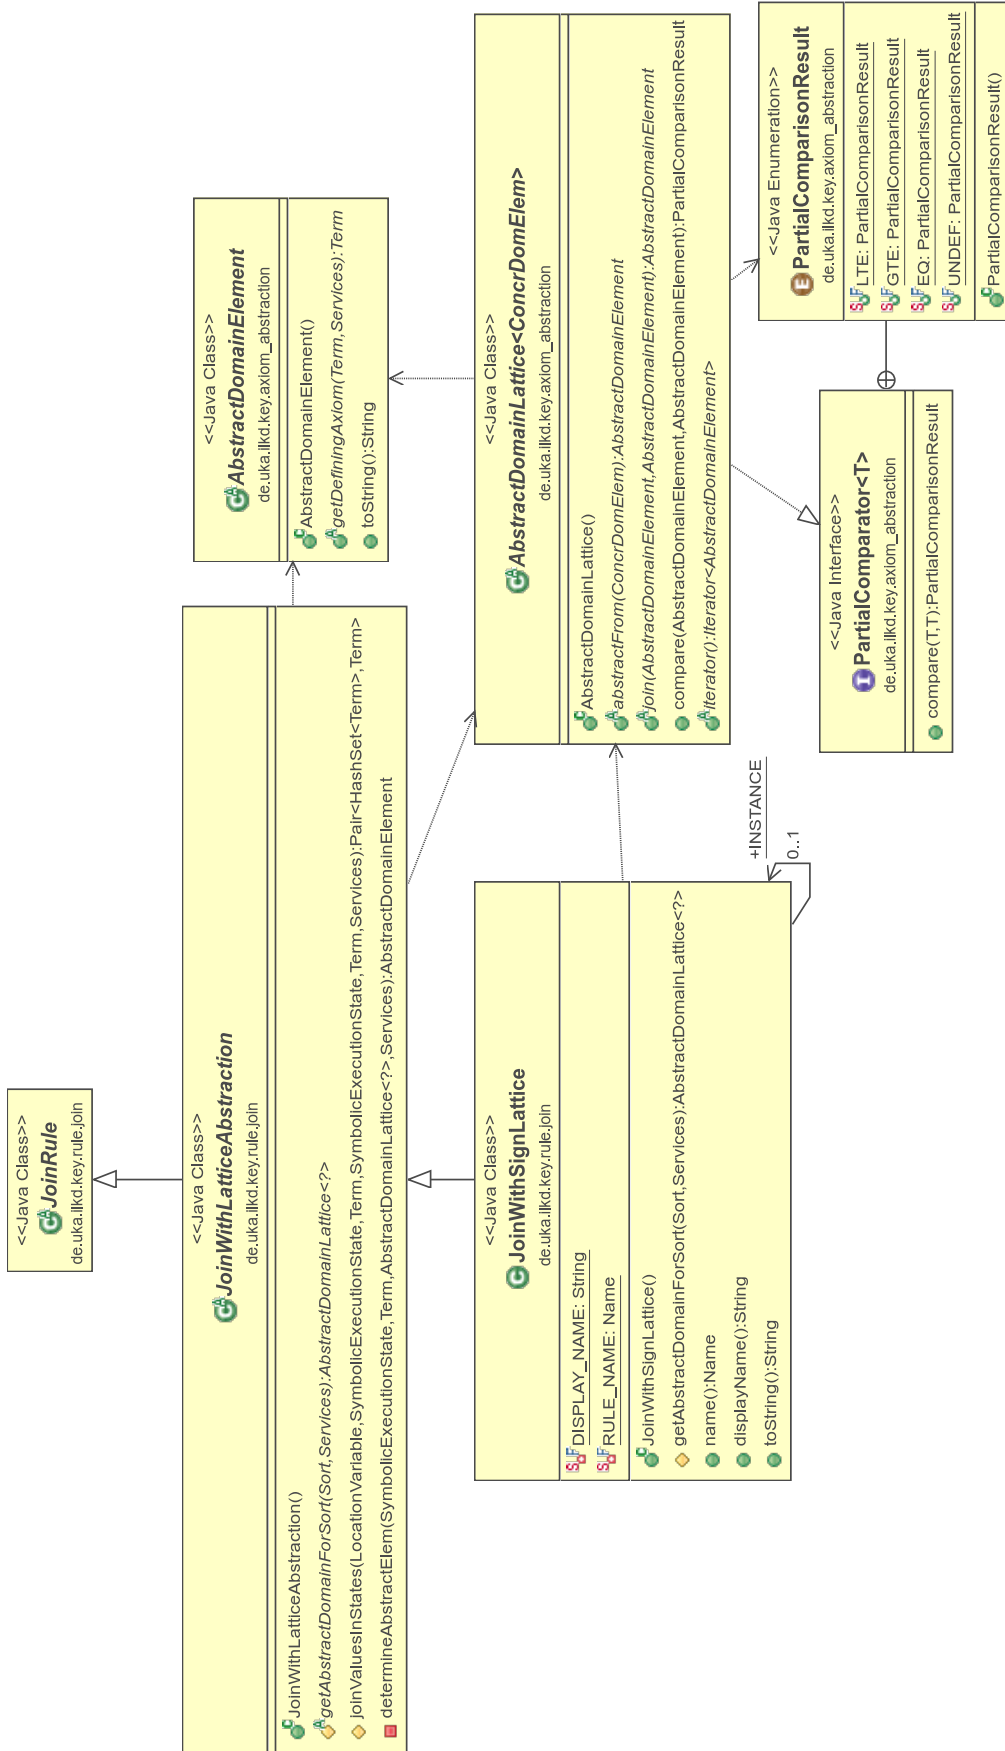


Figure 6.3: Lattice abstraction framework and rules

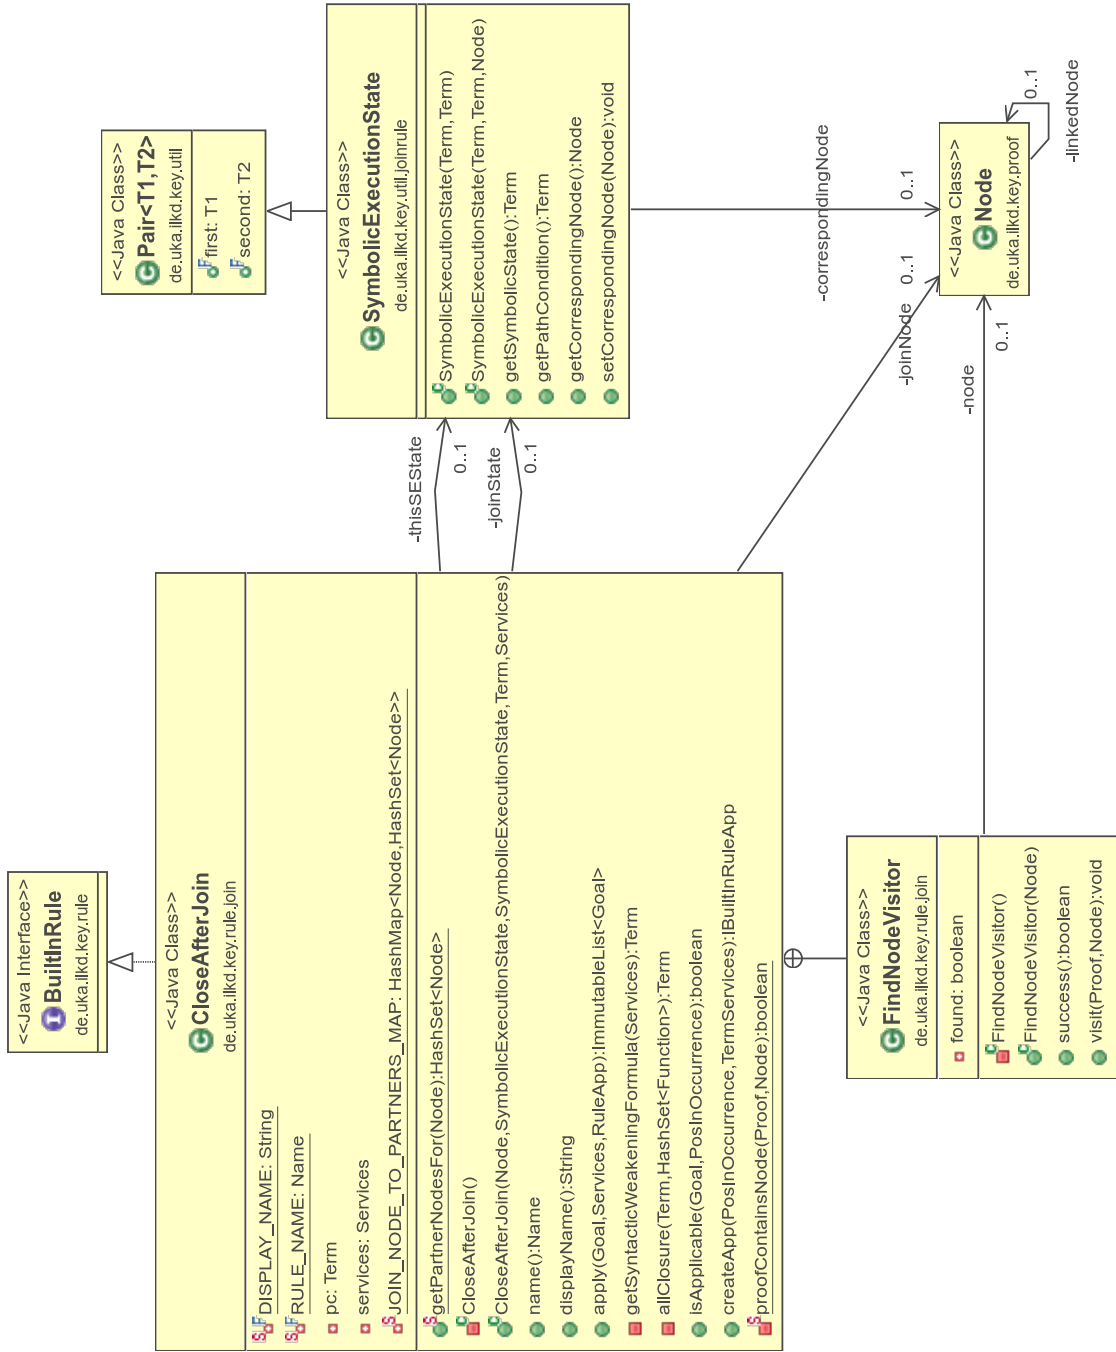


Figure 6.4: CloseAfterJoin rule



Figure 6.5: FinishSymbolicExecutionUntilJoinPointMacro

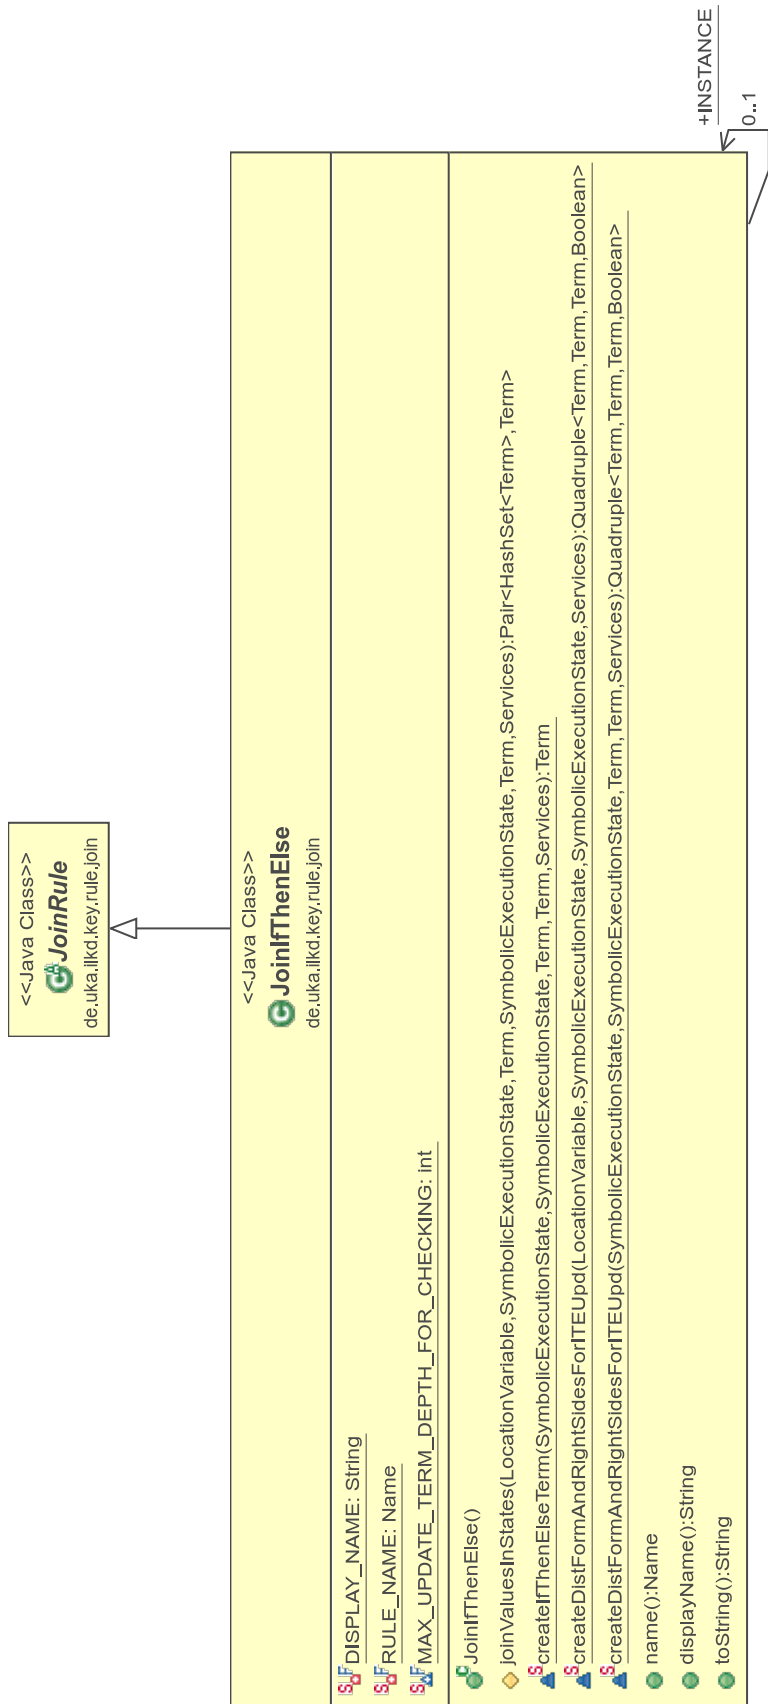


Figure 6.6: If-then-else rule implementation

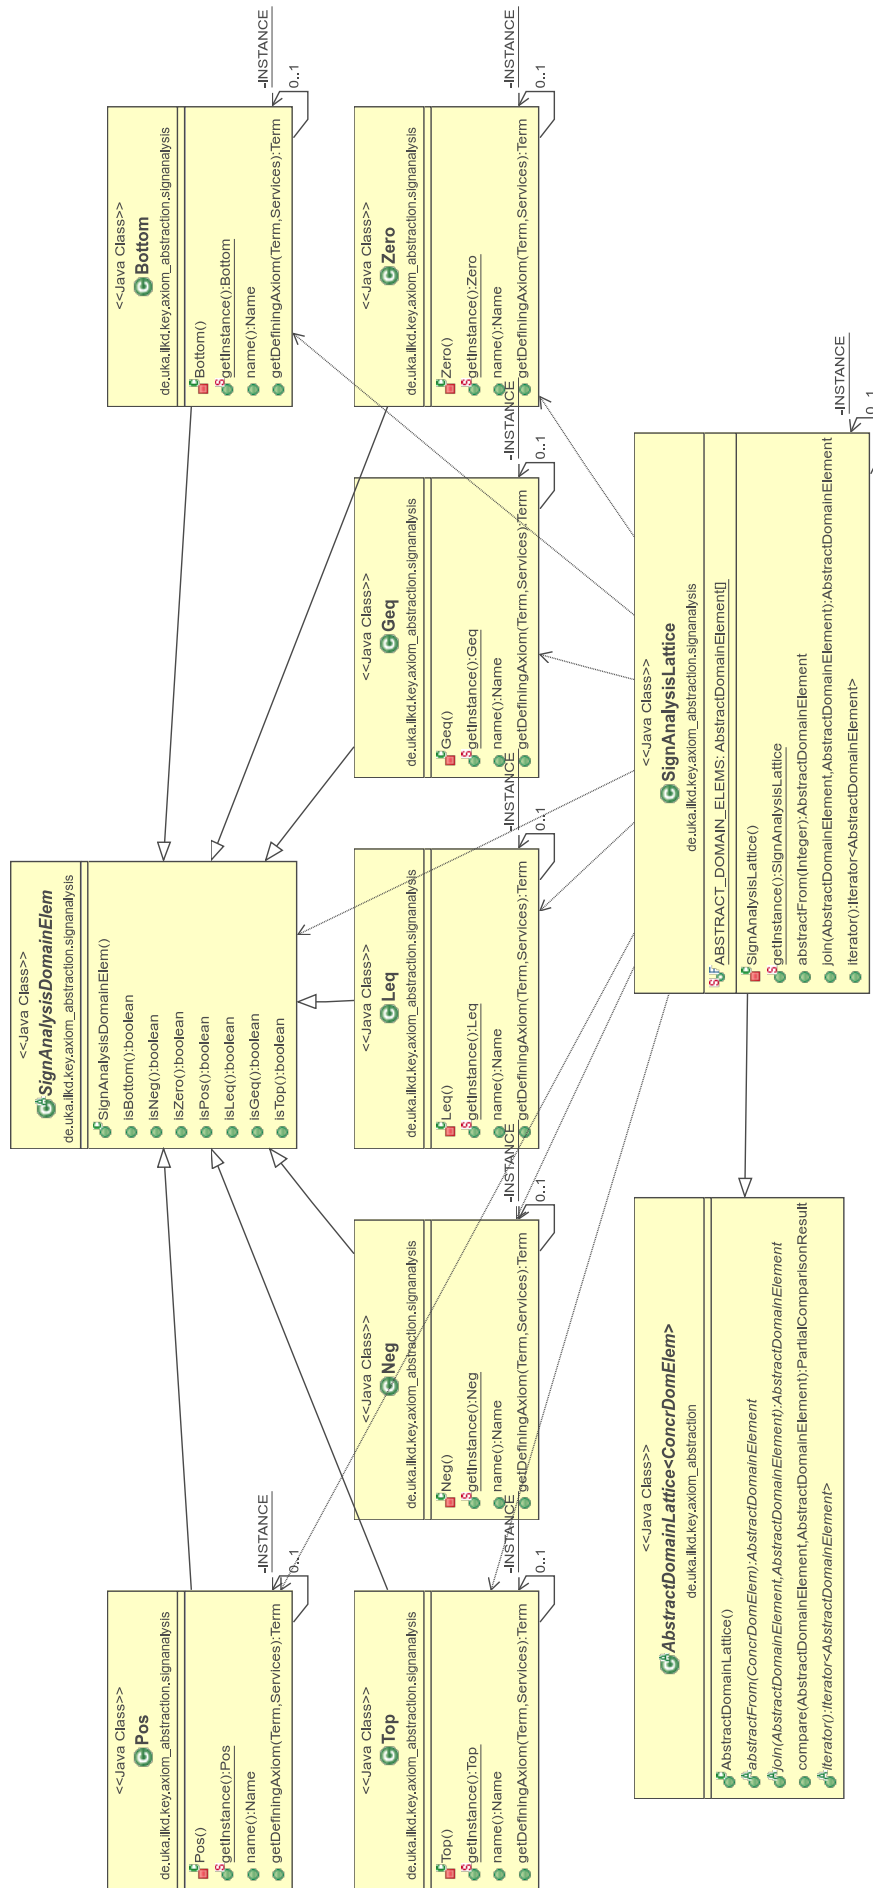


Figure 6.7: Sign Analysis Lattice

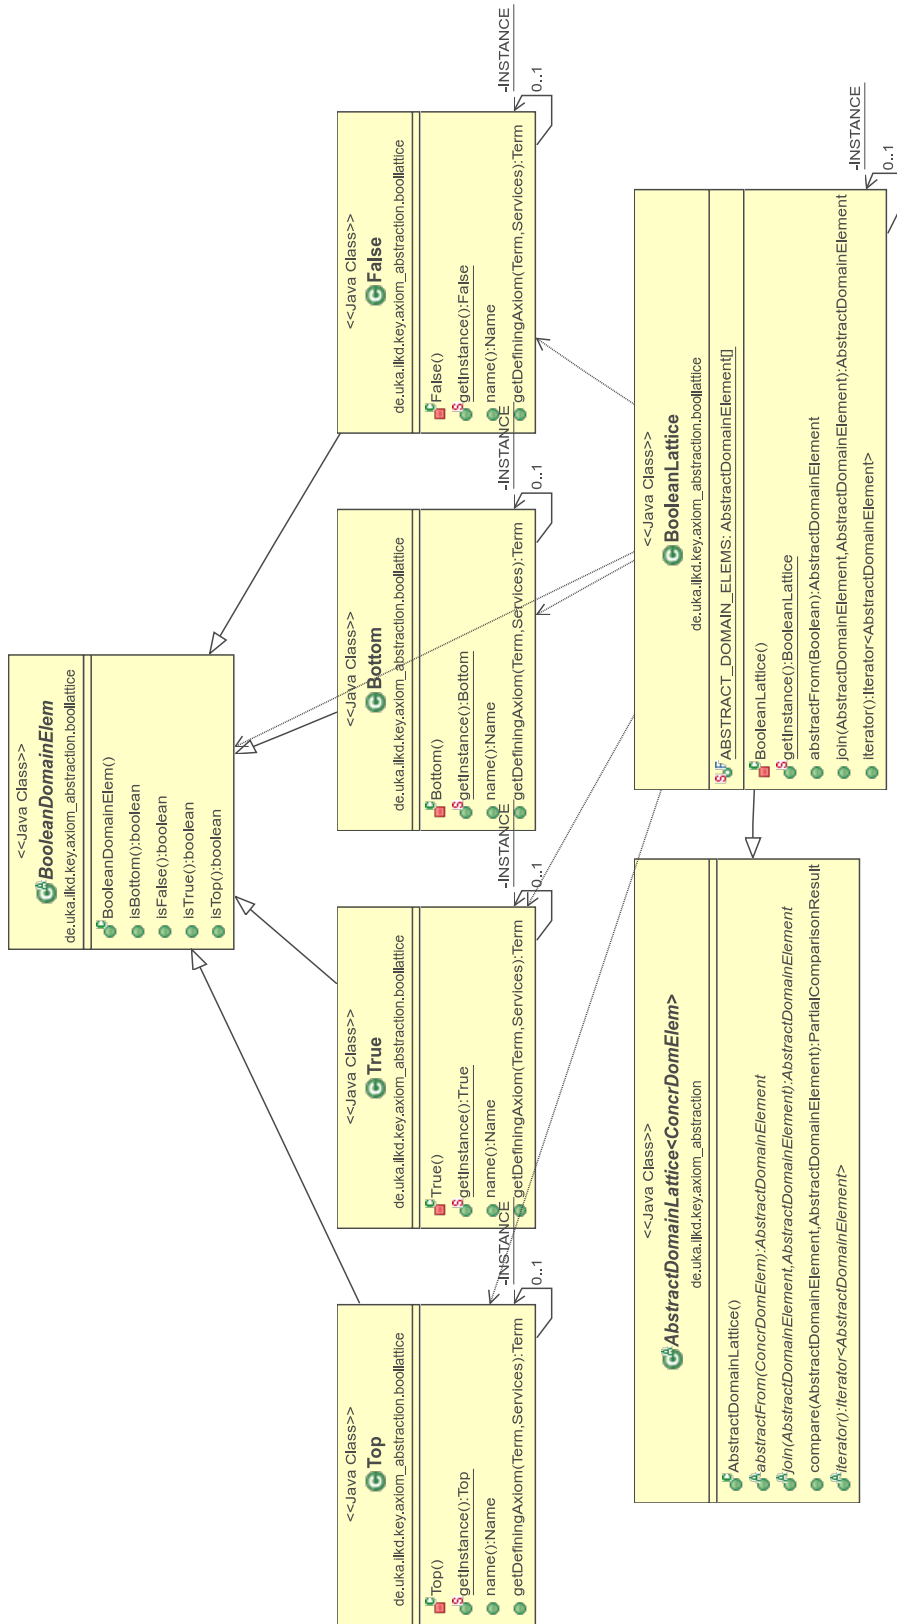


Figure 6.8: Boolean Lattice

List of Acronyms

CFG	Control Flow Graph
DAG	Directed Acyclic Graph
DL	Dynamic Logic
PDG	Program Dependence Graph
SE	Symbolic Execution
SET	Symbolic Execution Tree
SMT	Satisfiability Modulo Theories

List of Figures

1.1	Concrete Execution vs. Symbolic Execution	4
2.1	CFG for Listing 2.1.	7
2.2	Structure of Java DL type hierarchies	8
2.3	Partial Symbolic Execution Tree for Listing 2.1.	11
2.4	Selected rules of the KeY calculus for symbolic execution	12
3.1	Joining two branches in an abstract SET after an <code>if</code> statement.	14
3.2	Joining two branches in a KeY proof tree after an <code>if</code> statement.	14
3.3	Double application of the weakening rule	21
3.4	Design space for join operators	21
3.5	Abstract domain for sign analysis	27
4.1	Visualization of <code>CloseAfterJoin</code> rule	33
6.1	Overview of join classes	48
6.2	Join rule implementation	49
6.3	Lattice abstraction framework and rules	50
6.4	<code>CloseAfterJoin</code> rule	51
6.5	<code>FinishSymbolicExecutionUntilJoinPointMacro</code>	52
6.6	If-then-else rule implementation	53
6.7	Sign Analysis Lattice	54
6.8	Boolean Lattice	55

List of Tables

5.2	Performance evaluation for the example programs <code>div</code> , <code>abs</code> , <code>absObj</code> , <code>find</code> , <code>multiply</code> , and <code>gcd</code>	38
5.3	Strategies used in performance evaluation.	39

List of Algorithms

1	Information flow analysis without joins	42
2	Information flow analysis with joins	44

List of Listings

2.1	Example Java program	7
4.1	Excerpt of the <code>JoinRule</code> implementation	31
4.2	Linking of partner nodes to join nodes	34
4.3	<code>FinishSymbolicExecutionUntilJoinPointMacro</code> (excerpt)	35
4.4	Step two of distinguishing formulae simplification (Part 1)	35
4.5	Step two of distinguishing formulae simplification (Part 2)	36
5.1	The <code>abs</code> example – If block	40
5.2	The <code>absObject</code> example – Join with heaps	40
5.3	The <code>div</code> example – Exceptional control flow	40



5.4	The find example – While loop with break statement	41
5.5	The multiply example – Loop unrolling	41
5.6	The gcd example – Normalization before method call	42
5.7	“Exceptional Flow” Example	43

Bibliography

- [All70] Frances E. Allen. “Control Flow Analysis”. In: *Proceedings of a Symposium on Compiler Optimization*. New York, NY, USA: ACM, 1970, pp. 1–19.
- [APV06] Saswat Anand, Corina S. Pasareanu, and Willem Visser. “Symbolic Execution with Abstract Subsumption Checking”. In: *Model Checking Software*. Ed. by Antti Valmari. Lecture Notes in Computer Science 3925. Springer Berlin Heidelberg, Jan. 2006, pp. 163–181. ISBN: 978-3-540-33102-5, 978-3-540-33103-2.
- [BCE08] Peter Boonstoppel, Cristian Cadar, and Dawson Engler. “RWset: Attacking path explosion in constraint-based test generation”. In: *IN TACAS’08: INTERNATIONAL CONFERENCE ON TOOLS AND ALGORITHMS FOR THE CONSTRUCTIONS AND ANALYSIS OF SYSTEMS*. 2008.
- [BDR04] G. Barthe, PR. D’Argenio, and T. Rezk. “Secure Information Flow by Self-Composition”. In: *17th IEEE Computer Security Foundations Workshop, 2004. Proceedings*. June 2004, pp. 100–114.
- [BEL75] Robert S. Boyer, Bernard Elspas, and Karl N. Levitt. “SELECT – A Formal System for Testing and Debugging Programs by Symbolic Execution”. In: *Proceedings of the International Conference on Reliable Software*. New York, NY, USA: ACM, 1975, pp. 234–245.
- [Ben11] Benjamin Weiß. “Deductive Verification of Object-Oriented Software: Dynamic Frames, Dynamic Logic and Predicate Abstraction”. PhD thesis. Karlsruhe: Karlsruhe Institute of Technology, 2011.
- [BHS07] Bernhard Beckert, Reiner Hähnle, and Peter H. Schmitt. *Verification of Object-Oriented Software: The KeY Approach*. LNCS 4334. Berlin, Heidelberg: Springer-Verlag, 2007. ISBN: 3-540-68977-X, 978-3-540-68977-5.
- [BHW09] Richard Bubel, Reiner Hähnle, and Benjamin Weiß. “Abstract Interpretation of Symbolic Execution with Explicit State Updates”. In: *Formal Methods for Components and Objects*. Ed. by Frank S. de Boer, Marcello M. Bonsangue, and Eric Madelaine. Lecture Notes in Computer Science 5751. Springer Berlin Heidelberg, Jan. 2009, pp. 247–277. ISBN: 978-3-642-04166-2, 978-3-642-04167-9.
- [Bri+11] Angelo Brillout et al. “Beyond Quantifier-Free Interpolation in Extensions of Presburger Arithmetic”. In: *Verification, Model Checking, and Abstract Interpretation*. Ed. by Ranjit Jhala and David Schmidt. Lecture Notes in Computer Science 6538. Springer Berlin Heidelberg, 2011, pp. 88–102. ISBN: 978-3-642-18274-7, 978-3-642-18275-4.
- [BS08] J. Burnim and K. Sen. “Heuristics for Scalable Dynamic Test Generation”. In: *Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering. ASE ’08*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 443–446. ISBN: 978-1-4244-2187-9.
- [Bur74] Rodney Matineau Burstall. “Program Proving as Hand Simulation with a Little Induction”. In: *Information Processing*. Elsevier/North-Holland, 1974, pp. 308–312.
- [CC77] Patrick Cousot and Radhia Cousot. “Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints”. In: *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. POPL ’77. New York, NY, USA: ACM, 1977, pp. 238–252.
- [CJM14] Duc-Hiep Chu, Joxan Jaffar, and Vijayaraghavan Murali. “Lazy Symbolic Execution for Enhanced Learning”. In: *Runtime Verification*. Ed. by Borzoo Bonakdarpour and Scott A. Smolka. Lecture Notes in Computer Science 8734. Springer International Publishing, Sept. 2014, pp. 323–339. ISBN: 978-3-319-11163-6, 978-3-319-11164-3.
- [Cou+05] Patrick Cousot et al. “The ASTREÉ Analyzer”. In: *Programming Languages and Systems*. Ed. by Mooly Sagiv. Lecture Notes in Computer Science 3444. Springer Berlin Heidelberg, 2005, pp. 21–30. ISBN: 978-3-540-25435-5, 978-3-540-31987-0.
- [Cou01] Patrick Cousot. “Abstract Interpretation Based Formal Methods and Future Challenges”. In: *Informatics*. Ed. by Reinhard Wilhelm. Lecture Notes in Computer Science 2000. Springer Berlin Heidelberg, Jan. 2001, pp. 138–156. ISBN: 978-3-540-41635-7, 978-3-540-44577-7.

- [CS00] Alessandra Carbone and Stephen Semmes. *A Graphic Apology for Symmetry and Implicitness*. Oxford University Press, 2000. ISBN: 0198507291.
- [CS13] Cristian Cadar and Koushik Sen. “Symbolic Execution for Software Testing: Three Decades Later”. In: *Communications of the ACM* 56.2 (Feb. 2013), pp. 82–90.
- [DE82] R.B. Dannenberg and G.W. Ernst. “Formal Program Verification Using Symbolic Execution”. In: *IEEE Transactions on Software Engineering* SE-8.1 (Jan. 1982), pp. 43–52.
- [DHS05] Ádám Darvas, Reiner Hähnle, and David Sands. “A Theorem Proving Approach to Analysis of Secure Information Flow”. In: *Security in Pervasive Computing*. Ed. by Dieter Hutter and Markus Ullmann. Lecture Notes in Computer Science 3450. Springer Berlin Heidelberg, 2005, pp. 193–209. ISBN: 978-3-540-25521-5, 978-3-540-32004-3.
- [Fin05] Marcelo Finger. “DAG Sequent Proofs with a Substitution Rule”. In: *We will show Them - Essays in honour of Dov Gabbays 60th birthday, vol. 1*. Ed. by S. Artemov et al. London: Kings College Publications, 2005, pp. 671–686.
- [FOW84] Jeanne Ferrante, Karl J. Ottenstein, and Joe D. Warren. “The Program Dependence Graph and its Use in Optimization”. In: *International Symposium on Programming*. Ed. by M. Paul and B. Robinet. Lecture Notes in Computer Science 167. Springer Berlin Heidelberg, Jan. 1984, pp. 125–132. ISBN: 978-3-540-12925-7, 978-3-540-38809-8.
- [Gen64] Gerhard Gentzen. “Investigations into Logical Deduction”. In: *American philosophical quarterly* (1964), pp. 288–306.
- [God12] Patrice Godefroid. “Test Generation using Symbolic Execution”. In: *LIPICs-Leibniz International Proceedings in Informatics*. Vol. 18. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2012, pp. 24–33.
- [Gos+05] James Gosling et al. *The Java (TM) Language Specification*. 3rd. Addison-Wesley Professional, 2005. ISBN: 0321246780.
- [Grä78] Georg Grätzer. *Lattice Theory: Foundation*. Pure and Applied Mathematics 75. Academic Press, 1978. ISBN: 0-12-2957504.
- [Häh+10] Reiner Hähnle et al. “A Visual Interactive Debugger based on Symbolic Execution”. In: *Proceedings of the IEEE/ACM international conference on Automated software engineering*. Ed. by Jamie Andrews and Elisabetta Di Nitto. ACM Press, 2010, pp. 143–146.
- [HHB14] Martin Hentschel, Reiner Hähnle, and Richard Bubel. “Visualizing Unbounded Symbolic Execution”. In: *Tests and Proofs*. Ed. by Martina Seidl and Nikolai Tillmann. Lecture Notes in Computer Science 8570. Springer International Publishing, Jan. 2014, pp. 82–98. ISBN: 978-3-319-09098-6, 978-3-319-09099-3.
- [HKS06] Christian Hammer, Jens Krinke, and Gregor Snelting. “Information Flow Control for Java Based on Path Conditions in Dependence Graphs”. In: *Proceedings of the IEEE International Symposium on Secure Software Engineering (ISSSE 2006)*. Arlington, VA: IEEE, Mar. 2006, pp. 87–96.
- [HSS09] Trevor Hansen, Peter Schachte, and Harald Søndergaard. “State Joining and Splitting for the Symbolic Execution of Binaries”. In: *Runtime Verification*. Ed. by Saddek Bensalem and Doron A. Peled. Lecture Notes in Computer Science 5779. Springer Berlin Heidelberg, 2009, pp. 76–92. ISBN: 978-3-642-04693-3, 978-3-642-04694-0.
- [Jaf+12] Joxan Jaffar et al. “TRACER: A Symbolic Execution Tool for Verification”. In: *Proceedings of the 24th International Conference on Computer Aided Verification. CAV12*. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 758–766. ISBN: 978-3-642-31423-0.
- [JH14] Ran Ji and Reiner Hähnle. *Information Flow Analysis Based on Program Simplification*. Technical Report TUD-CS-2014-0877. Department of Computer Science, Technische Universität Darmstadt, 2014.
- [JMN13] Joxan Jaffar, Vijayaraghavan Murali, and Jorge A. Navas. “Boosting Concolic Testing via Interpolation”. In: *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering. ESEC/FSE 2013*. New York, NY, USA: ACM, 2013, pp. 48–58. ISBN: 978-1-4503-2237-9.
- [Kin76] James C. King. “Symbolic Execution and Program Testing”. In: *Communications of the ACM* 19.7 (July 1976), pp. 385–394.

-
- [Kuz+12] Volodymyr Kuznetsov et al. “Efficient State Merging in Symbolic Execution”. In: *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’12. New York, NY, USA: ACM, 2012, pp. 193–204. ISBN: 978-1-4503-1205-9.
- [McM05] K. L. McMillan. “Applications of Craig Interpolants in Model Checking”. In: *Tools and Algorithms for the Construction and Analysis of Systems*. Ed. by Nicolas Halbwachs and Lenore D. Zuck. Lecture Notes in Computer Science 3440. Springer Berlin Heidelberg, 2005, pp. 1–12. ISBN: 978-3-540-25333-4, 978-3-540-31980-1.
- [OO84] Karl J. Ottenstein and Linda M. Ottenstein. “The Program Dependence Graph in a Software Development Environment”. In: *Proceedings of the First ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments*. SDE 1. New York, NY, USA: ACM, 1984, pp. 177–184. ISBN: 0-89791-131-8.
- [RRR13] Ran Ji, Reiner Hähnle, and Richard Bubel. *Program Transformation Based on Symbolic Execution and Deduction*. Technical Report TUD-CS-2013-0348. Darmstadt: Department of Computer Science, Technische Universität Darmstadt, 2013.
- [Rüm03] Philipp Rümmer. *Ensuring the Soundness of Taclets – Constructing Proof Obligations for Java Card DL Taclets*. Minor Thesis, University of Karlsruhe, 2003.
- [Sen+14] Koushik Sen et al. *MultiSE: Multi-Path Symbolic Execution using Value Summaries*. Tech. rep. UCB/EECS-2014-173. EECS Department, University of California, Berkeley, Oct. 2014.
- [SM06] A. Sabelfeld and A. C. Myers. “Language-based Information-flow Security”. In: *IEEE J.Sel. A. Commun.* 21.1 (Sept. 2006), pp. 5–19.
- [SRK06] Gregor Snelting, Torsten Robschink, and Jens Krinke. “Efficient Path Conditions in Dependence Graphs for Software Safety Analysis”. In: *ACM Trans. Softw. Eng. Methodol.* 15.4 (2006), pp. 410–457.
- [SRW02] Mooly Sagiv, Thomas Reps, and Reinhard Wilhelm. “Parametric Shape Analysis via 3-valued Logic”. In: *ACM Trans. Program. Lang. Syst.* 24.3 (2002), pp. 217–298.
- [Xie+09] Tao Xie et al. “Fitness-guided path exploration in dynamic symbolic execution”. In: *IEEE/I-FIP International Conference on Dependable Systems Networks, 2009. DSN ’09*. June 2009, pp. 359–368.

Further Reading

- [Ahr+02] Wolfgang Ahrendt et al. “The Key System: Integrating Object-Oriented Design and Formal Methods”. In: *Fundamental Approaches to Software Engineering*. Ed. by Ralf-Detlef Kutsche and Herbert Weber. Lecture Notes in Computer Science 2306. Springer Berlin Heidelberg, Jan. 2002, pp. 327–330. ISBN: 978-3-540-43353-8, 978-3-540-45923-1.
- [Ahr+05] Wolfgang Ahrendt et al. “The KeY Tool”. In: *Software & Systems Modeling* 4.1 (Feb. 2005), pp. 32–54.
- [Ahr+06] Wolfgang Ahrendt et al. “Verifying Object-Oriented Programs with KeY: A Tutorial”. In: *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006, Amsterdam, The Netherlands, November 7-10, 2006, Revised Lectures*. Ed. by Frank S. de Boer et al. Vol. 4709. Lecture Notes in Computer Science. Springer, 2006, pp. 70–101. ISBN: 978-3-540-74791-8.
- [Ahr+14] Wolfgang Ahrendt et al. “The KeY Platform for Verification and Analysis of Java Programs”. In: *Verified Software: Theories, Tools and Experiments*. Ed. by Dimitra Giannakopoulou and Daniel Kroening. Lecture Notes in Computer Science. Springer International Publishing, Jan. 2014, pp. 55–71. ISBN: 978-3-319-12153-6, 978-3-319-12154-3.
- [Bal+00] Michael Balser et al. “Formal System Development with KIV”. In: *Fundamental Approaches to Software Engineering*. Ed. by Tom Maibaum. Lecture Notes in Computer Science 1783. Springer Berlin Heidelberg, Jan. 2000, pp. 363–366. ISBN: 978-3-540-67261-6, 978-3-540-46428-0.
- [BHH12] Richard Bubel, Reiner Hähnle, and Ran Ji. “Program Specialization via a Software Verification Tool”. In: *Formal Methods for Components and Objects*. Ed. by Bernhard K. Aichernig, Frank S. de Boer, and Marcello M. Bonsangue. Lecture Notes in Computer Science 6957. Springer Berlin Heidelberg, Jan. 2012, pp. 80–101. ISBN: 978-3-642-25270-9, 978-3-642-25271-6.
- [CDE08] Cristian Cadar, Daniel Dunbar, and Dawson Engler. “KLEE: Unassisted and Automatic Generation of High-coverage Tests for Complex Systems Programs”. In: *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation*. OSDI’08. Berkeley, CA, USA: USENIX Association, 2008, pp. 209–224.
- [Deu76] L. Peter Deutsch. *An Interactive Program Verifier*. Xerox Palo Alto Research Center, 1976.
- [DLR06] Xianghua Deng, Jooyong Lee, and Robby. “Bogor/Kiasan: A k-bounded Symbolic Execution for Checking Strong Heap Properties of Open Systems”. In: *21st IEEE/ACM International Conference on Automated Software Engineering, 2006. ASE ’06*. Sept. 2006, pp. 157–166.
- [EH07] Christian Engel and Reiner Hähnle. “Generating Unit Tests from Formal Proofs”. In: *Tests and Proofs*. Ed. by Yuri Gurevich and Bertrand Meyer. Lecture Notes in Computer Science 4454. Springer Berlin Heidelberg, Jan. 2007, pp. 169–188. ISBN: 978-3-540-73769-8, 978-3-540-73770-4.
- [FJM05] Jeffrey Fischer, Ranjit Jhala, and Rupak Majumdar. “Joining Dataflow with Predicates”. In: *Proceedings of the 10th European Software Engineering Conference Held Jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ESEC/FSE-13. New York, NY, USA: ACM, 2005, pp. 227–236. ISBN: 1-59593-014-0.
- [FQ02] Cormac Flanagan and Shaz Qadeer. “Predicate Abstraction for Software Verification”. In: *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’02. New York, NY, USA: ACM, 2002, pp. 191–202. ISBN: 1-58113-450-9.
- [GAP10] Miguel Gómez-Zamalloa, Elvira Albert, and Germán Puebla. “Test Case Generation for Object-oriented Imperative Languages in Clp^{*}”. In: *Theory Pract. Log. Program.* 10.4-6 (July 2010), pp. 659–674.
- [GL11] Patrice Godefroid and Daniel Luchaup. “Automatic Partial Loop Summarization in Dynamic Test Generation”. In: *Proceedings of the 2011 International Symposium on Software Testing and Analysis*. ACM, 2011, pp. 23–33.

-
- [GL75] S. Galand and G. Loncour. “Structured Implementation of Symbolic Execution: A First Part in a Program Verifier”. In: *Information Processing Letters* 3.4 (Mar. 1975), pp. 97–103.
- [Häh+86] R. Hähnle et al. “An Interactive Verification System based on Dynamic Logic”. In: *8th International Conference on Automated Deduction*. Ed. by Jörg H. Siekmann. Lecture Notes in Computer Science 230. Springer Berlin Heidelberg, Jan. 1986, pp. 306–315. ISBN: 978-3-540-16780-8, 978-3-540-39861-5.
- [HK76] Sidney L. Hantler and James C. King. “An Introduction to Proving the Correctness of Programs”. In: *ACM Comput. Surv.* 8.3 (Sept. 1976), pp. 331–353.
- [HRS87] M. Heisel, W. Reif, and W. Stephan. “Program Verification by Symbolic Execution and Induction”. In: *GWAI-87 11th German Workshop on Artificial Intelligence*. Ed. by Katharina Morik. Informatik-Fachberichte 152. Springer Berlin Heidelberg, Jan. 1987, pp. 201–210. ISBN: 978-3-540-18388-4, 978-3-642-73005-4.
- [HT08] Jonathan de Halleux and Nikolai Tillmann. “Parameterized Unit Testing with Pex”. In: *Tests and Proofs*. Ed. by Bernhard Beckert and Reiner Hähnle. Lecture Notes in Computer Science 4966. Springer Berlin Heidelberg, Jan. 2008, pp. 171–181. ISBN: 978-3-540-79123-2, 978-3-540-79124-9.
- [HTK00] David Harel, Jerzy Tiuryn, and Dexter Kozen. *Dynamic Logic*. Cambridge, MA, USA: MIT Press, 2000. ISBN: 0262082896.
- [Jam+12] Konrad Jamrozik et al. “Augmented Dynamic Symbolic Execution”. In: *Automated Software Engineering (ASE), 2012 Proceedings of the 27th IEEE/ACM International Conference on*. Ed. by Michael Goedicke, Tim Menzies, and Motoshi Saeki. Essen, Germany: IEEE, 2012, pp. 254–257.
- [JB12] Ran Ji and Richard Bubel. “PE-KeY: A Partial Evaluator for Java Programs”. In: *Integrated Formal Methods*. Ed. by John Derrick et al. Lecture Notes in Computer Science 7321. Springer Berlin Heidelberg, Jan. 2012, pp. 283–295. ISBN: 978-3-642-30728-7, 978-3-642-30729-4.
- [JHB13] Ran Ji, Reiner Hähnle, and Richard Bubel. “Program Transformation Based on Symbolic Execution and Deduction”. In: *Software Engineering and Formal Methods*. Ed. by Robert M. Hierons, Mercedes G. Merayo, and Mario Bravetti. Lecture Notes in Computer Science 8137. Springer Berlin Heidelberg, Jan. 2013, pp. 289–304. ISBN: 978-3-642-40560-0, 978-3-642-40561-7.
- [Ji14] Ran Ji. “Sound Program Transformation Based on Symbolic Execution and Deduction”. Dissertation. TU Darmstadt, June 2014.
- [JP08] Bart Jacobs and Frank Pessens. *The VeriFast Program Verifier*. Technical Report CW-520. Department of Computer Science, Katholieke Universiteit Leuven, Aug. 2008.
- [Kne91] Ralf Kneuper. “Symbolic Execution: A Semantic Approach”. In: *Science of Computer Programming* 16.3 (1991), pp. 207–249.
- [Lei05] K. Rustan M. Leino. “Efficient Weakest Preconditions”. In: *Inf. Process. Lett.* 93.6 (2005), pp. 281–288.
- [NNH99] Flemming Nielson, Hanne R. Nielson, and Chris Hankin. *Principles of Program Analysis*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1999. ISBN: 3540654100.
- [PV04] Corina S. Păsăreanu and Willem Visser. “Verification of Java Programs Using Symbolic Execution and Invariant Generation”. In: *Model Checking Software*. Ed. by Susanne Graf and Laurent Mounier. Lecture Notes in Computer Science 2989. Springer Berlin Heidelberg, 2004, pp. 164–181. ISBN: 978-3-540-21314-7, 978-3-540-24732-6.
- [SST13] Jiri Slaby, Jan Strejček, and Marek Trtík. “Compact Symbolic Execution”. In: *Automated Technology for Verification and Analysis*. Ed. by Dang Van Hung and Mizuhito Ogawa. Lecture Notes in Computer Science 8172. Springer International Publishing, 2013, pp. 193–207. ISBN: 978-3-319-02443-1, 978-3-319-02444-8.
- [VP11] Dries Vanoverberghe and Frank Piessens. “Theoretical Aspects of Compositional Symbolic Execution”. In: *Fundamental Approaches to Software Engineering*. Ed. by Dimitra Gianakopoulou and Fernando Orejas. Lecture Notes in Computer Science 6603. Springer Berlin Heidelberg, Jan. 2011, pp. 247–261. ISBN: 978-3-642-19810-6, 978-3-642-19811-3.