

A GENERIC FRAMEWORK FOR SYMBOLIC
EXECUTION:
THEORY AND APPLICATIONS

by

ANDREI ARUSOAIÉ

*A dissertation
submitted in partial fulfillment of the
requirements for the degree of*

DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE

Faculty of Computer Science
ALEXANDRU IOAN CUZA, UNIVERSITY OF IAȘI

*2014
Iași, Romania*

Doctoral committee:

Professor Dr. DOREL LUCANU, Supervisor, UAIC

Associate Prof. Dr. ADRIAN IFTENE, Chair, UAIC

Associate Prof. Dr. MARIUS MINEA, Politehnica University of Timișoara

Researcher Dr. VLAD RUSU, INRIA Lille-Nord Europe

Professor Dr. GHEORGHE ȘTEFĂNESCU, University of Bucharest

A Generic Framework for Symbolic Execution: Theory and Applications

Andrei Arusoaie

Abstract

The modern world is shifting from the traditional workmanship to a more automated work environment, where software systems are increasingly used for automating, controlling and monitoring human activities. In many cases, software systems appear in critical places which may immediately affect our lives or the environment. Therefore, the software that runs on such systems has to be safe. This requirement has led to the development of various techniques to ensure software safety.

In this dissertation we present a language-independent framework for *symbolic execution*, which is a particular technique for testing, debugging, and verifying programs. The main feature of this framework is that it is parametric in the formal definition of a programming language. We formally define programming languages and symbolic execution, and then we prove that the feasible symbolic executions of a program and the concrete executions of the same program mutually simulate each other. This relationship between concrete and symbolic executions allow us to perform analyses on symbolic programs, and to transfer the results of those analyses to concrete instances of the symbolic programs in question. We use our symbolic execution framework to perform program verification using Hoare Logic and Reachability Logic. For the latter, we propose an alternative proof system, and we show that under reasonable conditions, a certain strategy executing our proof system is sound and weakly complete.

A prototype implementation of our symbolic execution framework has been developed in \mathbb{K} . We illustrate it on the symbolic execution, model checking, and deductive verification of nontrivial programs.

Contents

1	Introduction	5
1.1	Motivation	5
1.2	Contributions	7
1.3	Related work	9
1.4	Acknowledgements	21
1.4.1	Grants	22
1.5	List of Publications	22
1.6	Participation to Summer Schools, Workshops, Internships	23
2	Background	27
2.1	Algebraic Specifications	27
2.2	Many-Sorted First Order Logic	35
2.3	Matching Logic	38
2.4	Reachability Logic	43
2.5	The \mathbb{K} framework	48
3	Language Independent Symbolic Execution Framework	54
3.1	Language Definitions	55
3.2	Unification	58
3.3	Symbolic transition relation and properties	64
3.3.1	Coverage	65
3.3.2	Precision	67
3.4	Symbolic Execution via Language Transformation	70
4	Applications	74
4.1	Hoare-Logic Verification	75

4.1.1	IMP	77
4.1.2	From Hoare Logic to Reachability Logic	80
4.1.3	Hoare Logic by symbolic execution	81
4.2	Reachability-Logic Verification	88
4.2.1	Default strategy for automation	101
4.2.2	Weak Completeness: disproving RL formulas	102
5	Symbolic execution within \mathbb{K} Framework	105
5.1	Symbolic Execution within the \mathbb{K} Framework	105
5.1.1	Compiling the symbolic semantics	106
5.1.2	Running programs with symbolic input	108
5.2	Symbolic execution: use cases	112
5.2.1	LTL model checking	112
5.2.2	SIMPLE, symbolic arrays, and bounded model checking	113
5.2.3	KOOL: testing virtual method calls on lists	115
5.3	A prototype for Reachability-Logic Verification	117
5.3.1	Verifying a parallel program: FIND	117
5.3.2	Verifying the Knuth-Morris-Pratt string matching algorithm: KMP	121
6	Conclusions	126
6.1	Future Work	127
	Bibliography	129

*“Every beginning is a consequence
– every beginning ends something.”*
Paul Valery

Chapter 1

Introduction

This dissertation proposes a language independent framework for symbolic execution, based on formal definitions of programming languages and used for program analysis and verification.

1.1 Motivation

Programmable electronic devices have made our lives more efficient, pleasant, and convenient. Most of them are designed for personal usage (e.g. laptops, smartphones, tablets, etc.) and significantly improve communications and the information flow. Other devices have been adapted for monitoring and controlling industry automated systems, avionic flight management units, nuclear reactor control systems, spaceflight control, and so on. Many of them can be found in places which immediately affect our daily activities, like automotive braking systems, robotic surgery machines, or traffic control devices. No matter what the destination of an electronic device is, the most relevant component is the software that runs on it. For such critical systems, software has to be *safe* in order to avoid harmful situations. The growing demand for this type of systems has led to the development of various techniques to ensure software safety. A particular technique is *symbolic execution*, which has been used for testing, debugging, and verifying programs written in various programming languages. The concept mainly consists in sending *symbolic values* as inputs for programs instead of concrete ones, and the execution is done by

manipulating expressions involving symbolic values. A symbolic program execution typically memorises symbolic values of program variables and a *path condition*, which accumulates constraints on the symbolic values on the path leading to the current instruction. When the next instruction to be executed is a conditional statement, which depends on symbolic values, the execution is separated into distinct branches. The path condition is then updated to distinguish between the different branches. All possible executions can be merged together into a *symbolic execution tree*, where each node is associated with a *symbolic state* and each edge with a transition between these states. Symbolic states typically include the statements to be executed, the program counter, the set of variables and their values, and the current path condition.

The main advantage of symbolic execution is that one can reason about multiple concrete executions of a program at once. On the other hand, its main disadvantage is the state space explosion caused by decision statements and loops. Recently, the technique has found renewed interest in the research community due to the progress in decision procedures, new algorithmic developments, and availability of powerful computers.

Currently, a large number of tools use symbolic execution for different purposes: automated test case generation, program verification, static detection of runtime errors, predictive testing, etc. Despite the fact that symbolic execution is a language independent technique, the existing symbolic execution engines target specific programming languages. As a consequence, it is quite hard to extend or reuse them for a new language. Often, a new implementation is preferred rather than using an existing one. An alternative is to translate programs from one language to another and then run symbolically the transformed programs. However, this solution may have implications on different analyses because the obtained results depend on the translation. Following this approach, some tools implement symbolic execution for assembly languages. Then, compilers are used to translate programs into assembly code. Symbolic execution of a given program is actually done by executing the corresponding assembly program, not the initial one. Thus, one has to be sure that the semantics of the original program is preserved during compilation, otherwise symbolic execution may not explore the desired execution tree.

A non-negligible aspect, that the majority of the symbolic execution

tools are less concerned about, are formal definitions of programming languages. Current implementations are mainly based on existing compilers or interpreters. Due to the fact that standard manuals for programming languages usually contain informal descriptions, compiler developers may have a different understanding of the semantics of a language. Therefore, different compilers can exhibit different behaviours of the same program. This has a negative impact on symbolic execution too, since any compiler can possibly explore different execution paths of the same program. As a consequence, any results obtained by analysing programs using symbolic execution are not transferrable across platforms that use a different compiler. In contrast, formal semantics provides an exact (mathematical) meaning for each language construct, and symbolic execution based on formal semantics captures all possible behaviours of a program.

The aim of this dissertation is to create a language independent formal framework and a tool for symbolic execution which is based on the formal definition of a programming language. There are several advantages of this framework over the existing approaches. First, a generic framework for symbolic execution can be reused for any programming language, without additional knowledge about its implementation. Second, the framework is based on the trusted formal semantics of a language, and thus, misinterpretations of language constructs is avoided. Finally, in a formally defined framework for symbolic execution, the relationship between concrete and symbolic execution can be also formalised. This relationship essentially states if symbolic execution can be used as a sound program analysis technique. Additionally, a practical impact of this framework is modularity, since the symbolic execution engine is a standalone tool, which can be easily integrated by other tools.

1.2 Contributions

Here is a summary of the main contributions of the research presented in this dissertation:

1. The main contribution of this dissertation is a formal, language-independent theory and tool for symbolic execution, based on a language's operational semantics defined by term-rewriting.

2. On the theoretical side, we define symbolic execution as the application of rewrite rules in the semantics by *symbolic rewriting*. We prove that the symbolic execution thus defined has the following properties, which ensure that it is related to concrete program execution in a natural way:

Coverage: for every concrete execution there is a feasible symbolic execution on the same program path, and

Precision: for every feasible symbolic execution there is a concrete execution on the same program path,

where a symbolic execution is feasible if the path conditions along it are satisfiable. In terms of simulations, we say that the feasible symbolic executions and the concrete executions of any given program mutually simulate each other.

3. On the practical side, we present a prototype implementation of our approach in \mathbb{K} [97], a framework dedicated to defining formal operational semantics of languages. Since the current version of \mathbb{K} is based on standard rewriting, we show how symbolic rewriting can be achieved by applying certain modified rewrite rules (obtained by automatically transforming the original ones) in the standard manner over a symbolic domain.
4. We show that our symbolic execution framework can be used for verifying programs using Hoare-like Logic. The approach consists in translating Hoare triples into particular Reachability Logic [100] formulas, which can be automatically proved using symbolic execution and an SMT solver.
5. We present an automatic, language-independent program verification approach and prototype tool based on symbolic execution. The program-specification formalism we consider is Reachability Logic, a sound and relatively complete deduction system, which offers a lot of freedom (but very few guidelines) for constructing proofs. Hence, we propose an alternative proof system, in which symbolic execution becomes a rule for systematic proof construction. We show that, under reasonable conditions on the semantics of programming

languages and of the Reachability-Logic formulas, a certain strategy executing our proof system is sound and weakly complete.

1.3 Related work

In this section we present some of the approaches/tools that use/implement symbolic execution for various applications. We identify the key ingredients of symbolic execution (exhibited by previous research on this topic) which are incorporated into our framework, and we emphasise the need for a language independent framework, which should be based on the formal semantics of programming languages. Note that we focus more on the symbolic execution basic features exhibited by various tools and less on applications.

In 1976, James C. King presents one of the first tools for automated test generation (EFFIGY) based on symbolic execution. In his paper [62], King intuitively describes symbolic execution as follows:

“Instead of supplying the normal inputs to a program (e.g. numbers) one supplies symbols representing arbitrary values”.

The key concepts of symbolic execution introduced in [62] are: *symbolic values*, *path conditions*, *symbolic states* and *symbolic execution trees*. Briefly, symbolic values are nothing else than variables of a given type (e.g. integers), path conditions are logical formulas involving symbolic values, and symbolic execution trees are program executions merged together in a tree, where each branch is induced by a conditional statement and nodes are *symbolic states*. Let us consider that α is an integer symbolic value. The value of α can be any integer, but it is not known which one. For this reason, α can be considered and it will be used as an integer constant. When such values are provided as inputs to programs, the execution semantics of the program is changed into *symbolic execution semantics*. The symbolic execution of regular assignments is quite natural: the expression on the right hand side is evaluated and the result is assigned to the variable. For instance, if \mathbf{a} and \mathbf{b} are two variables, where $\mathbf{b} = \alpha$ and $\mathbf{a} = \mathbf{b} + 1$, then the evaluation of the second assignment will attach to \mathbf{a} the sym-

bolic value $\alpha + 1$. Thus, program variables are assigned with *symbolic expressions*, i.e. expressions involving symbolic values.

A “state” of a symbolic program execution contains the current statement, symbolic values of program variables, the program counter, and the current path condition. The states are connected through edges which represent the transition from one state to another. When the control flow of a program depends on symbolic values, say an IF-THEN-ELSE statement, then there are more possible choices, one for each branch of the statement. For instance, let pc be the current path condition and q be the condition of IF-THEN-ELSE. If q is a symbolic expression then we are in one of the following cases (we use \rightarrow for the logical implication):

- either $pc \rightarrow q$ holds and then the execution continues on the THEN branch;
- either $pc \rightarrow \neg q$ holds and the execution follows on the ELSE branch;
- neither expression $pc \rightarrow q$ nor expression $pc \rightarrow \neg q$ holds; in this case there is at least one set of inputs which satisfies the pc and would take the THEN branch, **and** there is at least one set of inputs which satisfies the pc and would take the ELSE branch.

In the last case the execution *forks* into two “parallel” executions, and thus, from the current state there are two transitions in the symbolic execution tree. A path condition has the property that it determines a unique control flow path from the program. Initially, the path condition has the initial value set to “true”, and it is constructed during the execution by collecting branch conditions (q): $pc = pc \wedge q$ and $pc = pc \wedge \neg q$. Note that a path condition constructed as above is always a satisfiable formula, since for every execution there are some input values which led the execution to that specific state. Thus, there are two important properties that path conditions can have:

- pc can never become identically “false” because there does exist a concrete input to the program which, when executed normally, will trace the same path.
- If pc_1 and pc_2 are two path conditions then $\neg(pc_1 \wedge pc_2)$ holds. If pc_1 and pc_2 have a common part then, there exists, by construction,

a state which caused the two paths to be split by adding q and $\neg q$ to the common part of the two paths.

Another property of symbolic execution exhibited in [62] is that if all symbolic values are instantiated by concrete values then the result is exactly the same as the one obtained when executing the program with the same concrete values. This property (also called *commutativity* in [62]) states that symbolic executions capture precisely concrete executions, and this is why symbolic execution is of interest also in program verification. The author presents a methodology for proving programs correct using the Floyd's method [42]. Three new ancillary statements (ASSUME, ASSERT, and PROVE) are introduced in the language in order to help with the proving process as follows: the user introduces assertions (statements) of the form $\text{ASSERT}(b)$, where b is the formula expected to be valid at (carefully chosen) specific points in the program, such that a fixed set of paths, which start and end with an ASSERT, is obtained. Now, one must show that using *any* set of variable values which satisfy the assertion at the beginning of the path, the execution resulting along the path must satisfy the predicate at the end. This can be easily achieved by symbolic execution as follows:

1. change the ASSERT at the beginning of a path with an ASSUME, and the ASSERT at the end with a PROVE statement
2. initialize pc to *true* and all program variables with symbolic values
3. perform symbolic execution as usual;
4. in the end, execute the PROVE statement and if it displays *true* then the program is verified successfully.

Using this technique, the author claims that for any program, which has a finite symbolic execution tree and the correctness criteria is made explicit with assertions, the exhaustive symbolic execution and the proof of correctness (done following the steps above) are exactly the same process.

EFFIGY, the system described in [62], is among one of the first tools which implemented symbolic execution (there were also other systems developed in parallel: SELECT [19], DISSECT [56], and Clarke's system [28]). It is able to execute programs of a simple language whose

statements are assignments, conditional statements, GOTOs, and statements for dealing with input/output operations. The only supported symbolic values are integers. From the beginning, EFFIGY was designed as an interactive tool which allows step-by-step execution, tracing and state manipulation (load/save/edit). To handle state space explosion, EFFIGY allows users to either limit (bound) the state exploration giving the maximum number of statements to be executed or, in the interactive mode, to decide when to stop the execution. Moreover they have the freedom to explore the execution tree by navigating in any direction (up/down, left/right). Due to its practical limitations, EFFIGY has been used more like an experimental tool. However, considerable insight into the general notion of symbolic execution has been gained during its development.

Symbolic execution presented in [62] served as the main specification for the development of the symbolic execution framework presented in this dissertation. As in [62], our approach deals with symbolic data, computing path conditions for each program state, and generating the symbolic execution tree using the same principles. On the other hand, we generalize symbolic execution in a few directions. First, we do not restrict the data domain only to integers, users having the freedom to decide the symbolic values domain. Second, we implement a language independent symbolic execution framework, so we do not stick to a particular language. Finally, we formally define symbolic execution and prove the Coverage and Precision properties, which formalize the *commutativity* property intuitively presented in [62]. Based on these properties, we have shown in Chapter 4 how this framework can be used for program verification. Particularly, we present in Section 4.1 a similar (but more generic) approach to the one in [62]. In terms of implementation, our tool for symbolic execution covers all the basic features of EFFIGY (visualisation of symbolic execution paths, displaying the whole execution tree, and interactive exploration of the execution tree).

A more complex symbolic execution tool is Symbolic (Java) PathFinder [81] (or SPF). SPF is an extension of Java PathFinder [82] (JPF), whose development started in 1991 as a front end for the SPIN model-checker [55]. Initially, JPF was designed to model-check Java programs. The latest de-

velopments transformed it in more than a model checker; it became an extensible JVM which can be used as a workbench for efficient implementations of analysis and verification tools. Unlike the experimental EFFIGY tool, JPF is more practical, being successfully used for test case generation [61], invariant detection (using iterative invariant strengthening) [83], predicate abstraction [33], model-checking for Java byte code [84], etc.

SPF has been successfully used to generate tests for a component of a NASA flight software system, achieving full test coverage, for a special coverage metric required by the developer. The basic idea behind SPF is to use the JPF's model-checker to generate the state space of a Java (bytecode) program and to collect the path condition for every state. This is the equivalent of generating the symbolic execution tree for the program in question. For a state s , the symbolic execution tree contains a satisfiable path condition. The satisfiability ensures the existence of a *model* (i.e. a set of concrete values for the variables occurring in the path condition for which the path condition holds). SPF obtains such a model using either an internal constraint solver or an external SMT solver. The model itself is actually a test, in the sense that its concrete execution follows a path in the control flow graph until it reaches state s . To avoid state space explosion, JPF performs bounded model-checking, up to a specific depth.

Being developed for a real life language, SPF revealed new practical challenges raised by symbolic execution, like support for complex data structures (lists, trees, arrays), inheritance, polymorphism, the need for powerful SAT, SMT or constraint solvers, etc. During the development of our language independent framework, we have considered all these challenges and we reused most the solutions proposed for SPF. However, there are some features which cannot be covered by our framework by default, since they purely depend on the language or the implementation platform. For instance, SPF permits symbolic execution to be started at any point in the program and at any time during concrete execution. This allows symbolic execution for specific parts of the program. Moreover, parts of the program can be annotated with pre-conditions that help improving the precision of the unit-level symbolic analysis (e.g. for avoiding generation of tests that violate the precondition). In our framework such features are not present by default (since we do not have a generic way to instrument when to switch from symbolic to concrete execution) but we offer means

to implement them. On the other hand, SPF is hard to extend or to reuse for implementing symbolic execution for a different language. From the theoretical point of view, SPF does not concern about defining symbolic execution or relating it with the concrete execution.

There are some modern challenges of symbolic execution which are related mostly to library classes and native code. Such code needs an explicit modelling to make symbolic execution possible. One approach, which is meant to trick the explicit modelling by avoiding it, is combining concrete and symbolic execution, also known as *concolic* execution. This type of hybrid analysis, which performs both concrete and symbolic execution has been implemented by several tools in order to get dynamic test generation. The main point of concolic execution is to perform concrete and symbolic execution together which outputs inputs to cover the “new” behaviour. For instance, consider that we first run a program p with concrete values for the input variables x and y . The execution follows naturally until, at some point, a conditional statement has to be executed. Here, the statement’s condition, say $x \leq y$, can be evaluated as usual, since we know the concrete values of x and y ; thus, the concrete execution may continue on one branch or another. In parallel, we also collect the condition which has to be satisfied by the variables such that the execution will follow the other path, i.e. $\neg(x \leq y)$. For the next run, we choose concrete values for x and y such that $\neg(x \leq y)$, and this will allow us to explore a different program path. This a *directed* symbolic execution, the new program paths and corresponding path conditions being discovered using concrete executions. The technique is also known as *dynamic symbolic execution* and it is implemented by many tools for test case generation, e.g. DART [45], CUTE [107], jCUTE [106], EXE [24], PEX [38] SAGE [46], CREST [22]. Among these, only PEX and SAGE are used at large scale. PEX is a test case generation tool implemented for Microsoft’s .NET platform and is distributed with Visual Studio. It performs a systematic program analysis (using dynamic symbolic execution) to determine test inputs for parameterised Unit Tests. PEX learns the program behaviour by monitoring execution traces. It uses a constraint solver to produce new test inputs which exercise different program behaviour. The result is an automatically generated small test suite which often achieves high code coverage.

A similar tool, also developed at Microsoft, is SAGE which uses whitebox fuzz testing that generates tests for x86 assembly code. The tool executes the program under test with a set of concrete valid inputs and records the explored path. Then, it evaluates symbolically the recorded path and gathers for constraints on inputs capturing how the program use these. The collected constraints are negated one by one in order to obtain new input values using an SMT solver. SAGE found many security bugs in Windows applications.

Other approaches [103, 20] use dynamic symbolic execution to detect errors in programs written in interpreted languages, e.g. Python. There are many challenges when performing symbolic execution on interpreted languages. For instance, statements in programs can wrap complex operations, so expressing path conditions as precise as possible requires a complex theory. Additionally, due to the fact that most of these languages are dynamically typed, symbolic execution has to handle type information also. Some variations of concolic execution consist in performing first symbolic execution, and then concrete execution. The reason is that the symbolic execution will first detect imprecisely inputs which produce errors and the concrete execution will filter the initial results.

Our framework supports both concrete and symbolic execution. For concrete inputs the tool acts as expected in a concrete run of the program. Thus, one can develop different variations of concolic execution.

Another body of related work is symbolic execution in term-rewriting systems. The technique called *narrowing*, initially used for solving equation systems in abstract datatypes, has been extended for solving reachability problems in term-rewriting systems and has successfully been applied to the analysis of security protocols [78]. Such analyses rely on powerful unification-modulo-theories algorithms [37], which work well for security protocols since there are unification algorithms modulo the theories involved there (exclusive-or, ...). This is not always the case for programming languages with arbitrarily complex datatypes.

All the aforementioned approaches are based on symbolic execution engines developed for specific programming languages. However, all of them follow the same principles of symbolic execution stated in [62]. A question which naturally arises in here is: why not having a tool which is able to run symbolically programs of any programming language? An approach which

might answer to this question is KLEE [23] (which generalises EXE [24]), a symbolic virtual machine build on top of the LLVM [65] compiler infrastructure. Since there is a symbolic virtual machine for LLVM and compilers from different languages (C, C++, Objective-C, Ruby, Python, Haskell, Java, D, PHP, Pure, Lua) to this virtual machine then, as a consequence, symbolic execution is available for any of these languages.

In terms of usability, KLEE proved its usefulness, especially in tools for automated test case generation and bug detection. It has been used to detect errors in GNU COREUTILS (version 6.10). The results were quite impressive, the tool achieving the highest code coverage (84.5% of the total number of lines of C code), and discovered ten unique bugs. One of them was introduced back in 1992, and has not been discovered for 16 years, even if COREUTILS was analysed by lots of tools, before KLEE.

However, there is a critical issue which is not addressed by KLEE. What is supposed to be its main advantage, namely the symbolic virtual machine built on top the LLVM compiler, is also one of its main disadvantages. When performing symbolic execution for a particular language, say \mathcal{L} , one expects the execution to happen as specified in the semantics of \mathcal{L} . However, in KLEE, \mathcal{L} programs are compiled into LLVM, and then, their execution depends on the semantics of LLVM. This could be a problem, since the semantics of \mathcal{L} and LLVM may differ for some language constructs, and the KLEE does not treat such cases. A very simple example is the semantics of the $+$ operator in C, which does not specify an evaluation order for the operands of $+$. If those operands are functions with side effects, then the evaluation from left to right might produce a different result than the evaluation from right to left. The LLVM compiler for C assumes a default evaluation order for the arguments of $+$, and thus, when performing symbolic execution the tool will explore a single path instead of two paths (corresponding to evaluation from left to right and evaluation from right to left). When performing some analysis based on symbolic execution of a program, then the results of the analysis might not be valid anymore when using a different compiler. Therefore, there is no guarantee that the translation from \mathcal{L} to LLVM preserves the semantics of \mathcal{L} 's constructs. This approach makes symbolic execution compiler dependent. On the other hand, the symbolic execution framework proposed in this dissertation is parametric in the formal semantics of a programming

language. We perform symbolic execution using *directly* the semantics of the language, no translations being involved in the process.

Recently, rewriting modulo SMT was proposed as a new (language independent) technique which combines the power of SMT solving, rewriting modulo SMT, and model checking, for analysing infinite-state open systems [89, 90]. An open system is modelled as a triple (Σ, E, R) , where (Σ, E) is an equational theory describing the system states and R is a set of rewrite rules describing the system's transitions. The state of an open system must include the state changes due to the environment. These changes are captured by new fresh symbolic variables introduced in the right hand side of the rewrite rules. Thus, the system states are represented not as concrete states, but as symbolic ones, i.e. terms with variables (ranging in the domains handled by the SMT solver) which are constrained by an SMT-solvable formula. Rewriting modulo SMT can symbolically rewrite such states (which may describe possibly an infinite number of concrete states). The approach from [90] is very similar with the one presented in this dissertation, since both of them deal with symbolic states, symbolic rewriting, and use SMT solvers to decide the satisfiability of path conditions. Moreover, the soundness and completeness results from [90] are very similar with our coverage and precision properties. On the other hand, the approach allows symbolic execution modulo axioms, which is not supported by our framework. However, this feature comes with hard conditions that the model used for interpretation has to meet.

A major application of symbolic execution that we are interested in this dissertation is program verification. Almost all of the tools performing verification are specialised for a particular verification target. For instance, BitBlaze [105] is a powerful binary analysis platform designated to detect malicious code and/or prevent it. This platform uses symbolic execution directly on binary code for verifying if particular program paths have security vulnerabilities (e.g. buffer overflows).

S2E [25] is a platform for analysing properties and behaviour of software, which is based on two ideas: *selective symbolic execution* and relaxed execution consistency models. Selective symbolic execution is a way to automatically minimise the amount of code to be executed symbolically, given a target analysis. The tool [26] has been used to verify different

properties of systems, e.g. check if memory safety holds along all critical paths, verify if the code that handles license keys in proprietary program is safe when used with different keys, etc. S2E allows users to specify what code to explore paths for, and treats the other code as an “environment” which runs concretely. It includes a set of heuristics to deal with the transitions between concrete and symbolic executions. S2E also supports two interesting features of symbolic execution that we have not addressed in this dissertation: use a cluster for speeding up symbolic execution and state/path merging [64]. CHEF [21] is a platform for obtaining symbolic execution engines for interpreted languages which relies on S2E. CHEF reuses the interpreter as an executable language specification, which reduces the effort of creating a symbolic execution engine. The authors of CHEF have also implemented engines for Python and Lua, and then used them to generate tests for popular code libraries. An interesting feature of this tool is that it handles very well native implementations for functions that are present in interpreted languages. For instance, in Python strings library there are lots of functions (e.g. `find`) which are natively implemented by the interpreter. When performing symbolic execution, CHEF also explores the internal code for these functions, and thus, it increases the tests coverage. In contrast, the symbolic execution framework that we propose in this dissertation handles only those language constructs that have formal semantics. Therefore, for native functions, one has to provide the code and the semantics of the language that code has been written in.

The systems code must obey many rules, e.g. assertions must succeed, allocated memory must be freed, opened files must be closed, etc. WOODPECKER [31] is a tool which verifies such rules in C programs. It explores the symbolic execution tree in order to find bugs, but in an optimised manner: instead of exploring the whole tree, the tool *prunes* it in a *sound* manner depending on the rules to be verified. This is a form of path slicing. Unlike S2E, WOODPECKER is only designed to check rules and it can probably make use of some features of S2E, e.g. state merging or speeding up the execution using a cluster. On the other hand, S2E may use the WOODPECKER’s capabilities to guide symbolic execution or to prune execution paths.

Some of the existent approaches use a combination of symbolic execution and static analysis to check whether a particular line of code is

reached [72, 112]. Symbolic execution is sometimes used to only validate feasible paths discovered a priori using static analysis of the program’s data and control dependencies [86]. This is the opposite of WOODPECKER’s approach, which uses first symbolic execution and then performs static analysis. Other known tools for program verification are [29], [60], and [87]. Several techniques are implemented to improve the performance of these tools, such as *bounded verification* [27] and *pruning* the execution tree by eliminating redundant paths [31]. The major advantage of these tools is that they perform very well, being able to verify substantial pieces of C or assembly code, which are parts of actual safety-critical systems. On the other hand, these verifiers hardcode the logic they use for reasoning, and verify only specific programs (e.g. written using subsets of C) for specific properties (e.g. allocated memory is eventually freed). In contrast, our goal is to create a generic platform for symbolic execution where such applications can be developed with less effort and in a reasonable amount of time. In this dissertation we present a general approach for program verification based on Reachability Logic and symbolic execution (Section 4.2).

The closest approach to the one presented in this dissertation is [98]. Reachability Logic (presented in detail in Section 2.4) was introduced as an alternative to Floyd/Hoare logics and it is based on the operational semantics of a programming language. It consists in a language independent proof system (Figure 2.1) designed for program verification. It includes a special rule called *Circularity* which is primarily used to handle circular behaviours in programs, due to loops, recursivity, jumps, etc. In Reachability Logic, program properties are expressed using *reachability rules* of the form $\varphi \Rightarrow \varphi'$, where φ and φ' are called *patterns*, i.e. program configurations (with variables) constrained by a first-order formula.

In [92] the authors show that the proof system of Reachability Logic without the *Circularity* rule is *sound* and *complete*, i.e. any program property derived using the proof rules corresponds to an operational behaviour, and any operational behaviour can be derived using the rules from the proof system. In fact, the soundness result is the equivalent of our *Coverage* property, while the completeness result is a weak form of *Precision*. The proof system of Reachability Logic without *Circularity* captures symbolic execution, but at a high level of abstraction. In this dissertation,

we formalise symbolic execution using unification and then we show how we implement it using standard rewriting. Here, we make a clear distinction between code and data (integers, booleans, . . .) in programs, and we only allow data to be symbolic. Our intent was to bring symbolic execution closer to the implementation, by formalising all its details. By doing so, we filled the gap between the (abstract) symbolic execution captured by the Reachability Logic proof system and its implementation based on rewriting. As a consequence, we were able to develop a (language independent) prototype for verifying programs using Reachability Logic. Our tool proves (or disproves) program properties given directly as reachability formulas, and we used it to verify non-trivial programs (Section 5.3).

An implementation closely related to ours is MatchC [93], which has been used for verifying several challenging C programs such as the Schorr-Waite garbage collector. MatchC uses the Reachability Logic formalism for program specifications; it is, however, dedicated to a specific programming language, and uses a particular implementation of the RL proof system. By contrast, we focus on genericity, i.e., on language-independence: given a programming language defined in an algebraic/rewriting setting, we automatically generate the semantics for performing symbolic execution on that language. We propose our own proof system and a default program-verification strategy on the resulting symbolic execution engine. The soundness of our approach has also been proved. It relies on a Circularity Principle adapted to Reachability Logic, which has been formulated in a different setting in [95].

Other approaches offer support for verification of code contracts over programs. Spec# [15] is a tool developed at Microsoft that extends C# with constructs for non-null types, preconditions, postconditions, and object invariants. Spec# comes with a sound programming methodology that permits specification and reasoning about object invariants even in the presence of callbacks and multi-threading. A similar approach, which provides functionality for checking the correctness of a JAVA implementation with respect to a given UML/OCL specification, is the KeY [1] tool. In particular, KeY allows to prove that after running a method, its postcondition and the class invariant hold, using Dynamic Logic [52] and symbolic execution. The VeriFast tool [59] supports verification of single and multi-threaded C and Java programs annotated with preconditions and

postconditions written in Separation Logic [88]. An advanced tool used for verifying C programs is FramaC [32], an extensible and collaborative platform which allows to verify that the source code is compliant with a given specification. For functional specifications FramaC uses a dedicated language, ACSL (ANSI/ISO C Specification Language) [16]. The specifications can be partial, concentrating on one aspect of the analysed program at a time.

All these tools are designed to verify programs that belong to a specific programming language. Regarding performance, our generic tool for verification is (understandably) not in the same league as tools targeting specific languages and/or specific program properties. We believe, however, that the building of fast language-specific verification tools can benefit from the general principles presented in this dissertation, in particular, regarding the building of program-verification tools on top of symbolic execution engines.

1.4 Acknowledgements

Since before I started my PhD, when I was a master student, my research activity has been guided by Professor Dorel Lucanu. Thanks to him I was able to develop myself as a researcher, rather than being only a blind programmer. I would like to especially thank him for his efforts and countless hours he spent for every single issue that I had. He also deserves the credits for creating and managing a research group which was essential for my formation as a member. His contributions to the \mathbb{K} framework and his inspiring ideas on symbolic execution played a major role in the development of this dissertation.

An important contribution to the development of this dissertation was made by Vlad Rusu. I would also like to especially thank him for the fruitful collaboration that we have.

During my master studies I published my first two papers under the supervision of prof. Adrian Iftene. I want to thank him for introducing me to the wonderful world of research and further support during my thesis.

Since before I started my PhD I am a member of the Formal Methods on Software Engineering group, which has a strong collaboration with the Formal Systems Laboratory group from UIUC. In this period, my former

and current colleagues have always been there to give me advices and ideas: Traian Florin Șerbănuță, Radu Mereuță, Ștefan Ciobăcă, Grigore Roșu, Chucky Ellison, Andrei Ștefănescu, Emilian Necula, Raluca Necula, Elena Naum, Mihai Asăvoae, Irina Măriuca Asăvoae. Thank you all!

Probably the key factor which helped me go through this period was my family. I would like to thank my parents, Gheorghită and Anișoara, and my sister, Andreea, for the fact that they have always been there for me. They have always trusted me and encouraged me with every choice that I have made. I would also want to especially thank my girlfriend, Andreea, for her patience, understanding, and support in the most difficult period of this PhD.

Last, but not least, I would like to thank my research committee, made up of Adrian Iftene, Marius Minea, Vlad Rusu, and Gheorghe Ștefănescu, for their patience and time spent to read and review this dissertation.

1.4.1 Grants

- In 2014, this work was supported by the strategic grant POSDRU /159/1.5/S/137750, “Project Doctoral and Postdoctoral programs support for increased competitiveness in Exact Sciences research” co-financed by the European Social Fund within the Sectorial Operational Program Human Resources Development 2007-2013
- Between 2011-2013, this work was supported by strategic grant POSCCE/161/15.06.2010, SMIS-CSNR 602-12516, “An Executable Semantic Framework for Rigorous Design, Analysis and Testing of Systems”.

1.5 List of Publications

Below is the list of publications directly related to this dissertation:

1. *Andrei Arusoaic*, Dorel Lucanu and Vlad Rusu, “A Generic Framework for Symbolic Execution”, Proceedings of 6th International Con-

- ference on Software Language Engineering (SLE'13), p. 260-281, LNCS, Indianapolis, USA, October, 2013
2. *Andrei Arusoaie*, “*Engineering Hoare Logic-based Program Verification in \mathbb{K} Framework*”, Proceedings of SYNASC'13, p. 179-186, IEEE, Timișoara, Romania, September 23-26, 2013.
 3. *Andrei Arusoaie*, Dorel Lucanu, Vlad Rusu, Traian Florin Șerbănuță, Grigore Roșu and Andrei Ștefănescu, “*Language Definitions as Rewrite Theories*”, Proceedings of WRLA'14 , LNCS, Grenoble, France, april 5-6, 2014.
 4. David Lazăr, *Andrei Arusoaie*, Traian Șerbănuță, Chucky Ellison, Radu Mereuță, Dorel Lucanu and Grigore Roșu, “*Executing Formal Semantics with the \mathbb{K} Tool*”, Proceedings of FM'12, p. 267-271, LNCS, Paris, France, August 27-31, 2012.
 5. *Andrei Arusoaie* and Daniel Ionuț Vicol, “*Automating Abstract Syntax Tree construction for Context Free Grammars*”, Proceedings of SYNASC'12, p. 152-160, IEEE, Timișoara, Romania, September 26-29, 2012
 6. *Andrei Arusoaie*, Traian Florin Șerbănuță, Chucky Ellison and Grigore Roșu, “*Making Maude Definitions more Interactive*”, Proceedings of WRLA'12, p. 83-99, LNCS, Tallinn, Estonia, March 24-25, 2012
 7. *Andrei Arusoaie*, Dorel Lucanu and Vlad Rusu, “*Towards a K semantics of OCL*”, ENTCS, 2014, p. 81-96
 8. Traian Florin Șerbănuță, *Andrei Arusoaie*, David Lazăr, Chucky Ellison, Dorel Lucanu and Grigore Roșu, *The \mathbb{K} Primer (version 3.3)*, ENTCS, 2014, p. 57-80

1.6 Participation to Summer Schools, Workshops, Internships

Summer Schools:

1. *Summer School on Coq*, NII Shonan Meeting, 2014, Shonan, Kanagawa, Japan
2. International Summer School on *Software Systems Safety*, MOD 2013, Marktobendorf, Germany
3. Summer School on Language Frameworks, SSLF 2012, Sinaia, Romania.

Workshops:

1. The 22nd *International Workshop on Algebraic Development Techniques*, WADT, 2014, Sinaia, Romania
2. The 2nd *International \mathbb{K} workshop*, K'11, Cheile Grădiștei, Braşov, Romania
3. The 1st *International Workshop on the \mathbb{K} Framework and its Applications*, K'10, Nags Head, North Carolina, USA

Internships:

1. INRIA Lille Nord Europe, October 1st - November 15th 2011, France
2. INRIA Lille Nord Europe, July 1st - July 30th 2013, BQR grant, France

List of Figures

2.1	The language-independent proof system of Reachability Logic.	47
2.2	Fragment of CinK syntax	50
2.3	CinK configuration	50
2.4	Subset of rules from the \mathbb{K} semantics of CinK	52
4.1	\mathbb{K} Syntax of IMP	76
4.2	\mathcal{S}_{IMP} : the \mathbb{K} Semantics of IMP	77
4.3	[HL-IMP]: Floyd-Hoare proof system for IMP	79
4.4	[HL'-IMP]: the equivalent Floyd-Hoare proof system given using reachability rules. Each Hoare triple from the proof system in Figure 4.3 has been transformed into an equivalent reachability rule using (4.1).	82
4.5	CinK program: gcd	89
4.6	Proof System for $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(G)$	97
5.1	sum.cink	109
5.2	log.cink	113
5.3	SIMPLE program: init-arrays	114
5.4	lists.kool: implementation of lists in KOOL	116
5.5	FIND program.	118

5.6	RL formulas necessary to verify FIND. We use <code>a</code> , <code>i</code> , <code>j</code> , <code>oddtop</code> , <code>eventop</code> , <code>N</code> , <code>k</code> to denote program variables, <code>a</code> , <code>i</code> , <code>j</code> , <code>o</code> , <code>e</code> , <code>N</code> , <code>k</code> to denote locations, and <code>a</code> , <code>i</code> , <code>j</code> , <code>o</code> , <code>e</code> , <code>N</code> , <code>k</code> for variables values. We also use $s(i)$ to denote a common sequence in the proofs of (\clubsuit) and (\diamond). <code>CaseAnalysis</code> splits the proof in two goals separated by \vee , while <code>CircularHypothesis(i)</code> represents the application of the formula (i) as a circularity. <code>[SymbolicStep] \times n</code> is the equivalent of applying <code>[SymbolicStep]</code> n times.	120
5.7	The KMP algorithm annotated with pre-/post-conditions and invariants: failure function (left) and the main function (right). Note that we used <code>Pi</code> , <code>Theta</code> , and <code>allOcc</code> to denote functions π and θ , and predicate <code>allOcc</code> , respectively.	123
5.8	Properties needed to prove KMP.	125

Chapter 2

Background

This chapter revisits some theoretical material used in this dissertation. The basics of algebraic specifications, together with the definitions and the notational conventions that we are going to use through the dissertation, are presented in Section 2.1. Section 2.2 contains a brief description of many-sorted First Order Logic. In Section 2.3 we present Matching Logic, a logic designed to state and reason about structural properties of arbitrary program configurations. Matching Logic formulas are the basic ingredients of Reachability Logic (Section 2.4), a logic used for defining the operational semantics of languages, and for stating program properties. At the end of the chapter, in Section 2.5, we present \mathbb{K} , a framework for defining formal semantics of programming languages. All these will be assumed known for the subsequent development of the dissertation.

2.1 Algebraic Specifications

In this section we briefly introduce, inspired from [47], the basic knowledge about algebraic specifications (definitions and notations) used through this dissertation. We assume that the reader is already familiar with set theory as it is used in mathematics today.

The use of algebraic specifications to model computer programs is motivated by the fact that they can manipulate several kinds of *sorts* of data, in the same way as programs do. Programs can be represented in terms of sets of data values and mathematical functions over these sets.

When using an algebra to model a program, there is a natural one-to-one correspondence between the sets of values in the algebra and the sorts of data in the program.

Let S be a set of sorts. We use S^* to denote the set of all lists of elements from S . For example, if $S = \{a, b, c, d\}$ then $a, ab, abd, [] \in S^*$, where $[]$ is the empty list. We often use w to denote non-empty lists of elements from S^* (e.g. $w = abd$).

A *map* is a function $f : A \rightarrow B$ with given source (domain) and target (codomain). In the dissertation, we make use of the following definitions and notations:

- The composition of two functions $f : A \rightarrow B$ and $g : B \rightarrow C$ is denoted by $g \circ f : A \rightarrow C$, where $(g \circ f)(x) = g(f(x))$ for all $x \in A$.
- Given two functions $f : A \rightarrow B$ and $g : A \rightarrow B$ we say that f and g are (extensionally) equal, denoted $f = g$, if for all $x \in A$ we have $f(x) = g(x)$.
- If $f : A \rightarrow B$ is a function and A' is a subset of A (i.e. $A' \subseteq A$) then the *restriction of f to A'* is the function $f|_{A'} : A' \rightarrow B$ defined as $f|_{A'}(x) = f(x)$, for all $x \in A'$.
- We use $1_A : A \rightarrow A$ to denote the identity function on set A : $1_A(x) = x$ for each $x \in A$.

Next, we define sets indexed by a set of sorts:

Definition 2.1.1 (S-indexed set) *Given S a set of sorts, an S-indexed set A is a set-valued map with source S and target $\{A_s \mid s \in S\}$, where the value at $s \in S$ is denoted by A_s .*

By $|A|$ we denote $\bigcup_{s \in S} A_s$ and we let $a \in A$ mean that $a \in |A|$. The *empty* S-indexed set has $\emptyset_s = \emptyset$ for each sort s in S , and we (ambiguously) denoted it by \emptyset . If A and B are two S-indexed sets then:

- $A \subseteq B$ iff $A_s \subseteq B_s$, for each $s \in S$, and
- $A \cup B \triangleq \{C_s \mid C_s = A_s \cup B_s, s \in S\}$.

Thus, symbols \subseteq and \cup , which correspond to inclusion and union operations are overloaded for S -indexed sets. The other known operations on sets, namely intersection (\cap), set difference (\setminus), and cartesian product (\times) are similarly extended to S -indexed sets.

In general, S -indexed sets allow us to manipulate a family of sets as a unit, in such a way that operations on this unit respect the “typing” of data values. We provide examples of S -indexed sets in Example 2.1.2.

The syntax of a programming language consists in a set of rules which specify how different symbols can be combined to obtain a (syntactically) correct program. It includes the set of data entities (sorts), the operations on them together with the sorts of their arguments, and the return sort. In an algebraic setting the syntax is given using *signatures*.

Definition 2.1.2 (Many-sorted signature) *Given S a set of sorts, an S -sorted signature Σ is an $S^* \times S$ -indexed family of sets $\{\Sigma_{w,s} \mid w \in S^*, s \in S\}$ of sets whose elements are called operation symbols. An operation symbol $f \in \Sigma_{w,s}$ is said to have arity w and sort s . A symbol $c \in \Sigma_{\square,s}$ is called a constant symbol.*

In general, the syntax of a programming language is specified using a Backus-Naur Form [75] (abbreviated BNF) grammar, which consists of specification rules of the form

$$Symbol ::= seq_1 \mid seq_2 \mid \dots,$$

where *Symbol* is a *nonterminal* and seq_1, seq_2, \dots are sequences of *terminals* (symbols that never appear in the left side of $::=$) or nonterminals; these sequences are separated by the \mid operator which indicates a *choice* between sequences, i.e. each such sequence can be a possible substitution for *Symbol*. We often refer to these sequences as *syntax productions*.

In Example 2.1.1 we show the BNF syntax of a simple language which deals only with expressions. *Int*, *Bool*, and *Exp* are nonterminals, while $\underline{0}$, *true* are terminal symbols, and *Int*, *Exp* + *Exp* are syntax production corresponding to nonterminal *Exp*. The choice operator “ \mid ” inside *true* \mid *false* specifies the fact that *Bool* can be either the *true* or *false*.

Example 2.1.1 Here we show the syntax of a very simple language which only deals with integer and boolean expressions:

$$\begin{aligned}
\text{Bool} & ::= \underline{\text{true}} \mid \underline{\text{false}} \\
\text{Int} & ::= \dots \mid \underline{-1} \mid \underline{0} \mid \underline{1} \mid \dots \\
\text{Exp} & ::= \text{Int} \mid \text{Bool} \\
& \quad \mid \text{Exp} * \text{Exp} \\
& \quad \mid \text{Exp} / \text{Exp} \\
& \quad \mid \text{Exp} + \text{Exp} \\
& \quad \mid \text{Exp} - \text{Exp} \\
& \quad \mid \neg \text{Exp} \\
& \quad \mid \text{Exp} \wedge \text{Exp} \\
& \quad \mid (\text{Exp})
\end{aligned}$$

The grammar above includes the non-terminals *Bool*, *Int*, and *Exp*, and the terminals $(*, /, +, -, (,), \neg, \wedge)$, $\underline{\text{true}}$, $\underline{\text{false}}$, $\underline{-1}$, $\underline{0}$, $\underline{1}$, \dots .

The BNF grammar of a language has a corresponding S -sorted signature Σ . For each syntax production of the form $\text{Symbol} ::= \text{seq} \mid \dots$, the nonterminal *Symbol* is a sort in S , and Σ includes the indexed sets obtained as below:

- If $\text{seq} \triangleq e$, where e is a terminal symbol, then $t \in \Sigma_{[], \text{Symbol}}$.
- If $\text{seq} \triangleq e$, where e is a nonterminal symbol, then $\Sigma_{e, \text{Symbol}}$ includes an operation symbol $\iota_{e, \text{Symbol}}$.
- If $\text{seq} \triangleq e_1 e_2 \dots e_n$ then the concatenation $e \triangleq t(e_1) t(e_2) \dots t(e_n) \in \Sigma_{w, \text{Symbol}}$; t is the function which maps e_i to symbol $_$ (if e_i is a nonterminal) or to e_i (if e_i is a terminal), for all $i \in \{1, 2, \dots, n\}$; w is obtained by concatenating the sorts of e_i .

In Example 2.1.2 we show the corresponding many sorted signature of the BNF grammar from Example 2.1.1.

Example 2.1.2 The set S includes the nonterminals *Bool*, *Int*, and *Exp* as sorts, that is, $S = \{\text{Bool}, \text{Int}, \text{Exp}\}$. For each syntax production we generate the operation names by concatenating the symbol $_$ (corresponding to nonterminal symbols) and the terminal symbols in the same order they

appear in the production (e.g. for production $Exp + Exp$ we generate the operation name $_{+}$).

The following indexed sets:

- $\Sigma_{[], Bool} = \{\underline{true}, \underline{false}\}$,
- $\Sigma_{Bool, Bool} = \{\neg_{Bool}\}$
- $\Sigma_{Bool Bool, Bool} = \{- \wedge_{Bool} -, - \vee_{Bool} -, \dots\}$
- $\Sigma_{[], Int} = \{\dots, \underline{-1}, \underline{0}, \underline{1}, \dots\}$,
- $\Sigma_{Int Int, Int} = \{- +_{Int} -, - -_{Int} -, - *_{Int} -, - /_{Int} -, \dots\}$
- $\Sigma_{Exp, Exp} = \{\neg, (-)\}$, and
- $\Sigma_{Exp Exp, Exp} = \{- + -, - - -, - * -, - / -, - \wedge -\}$
- $\Sigma_{Int, Exp} = \{\iota_{Int, Exp}\}$.
- $\Sigma_{Bool, Exp} = \{\iota_{Bool, Exp}\}$.

are included in the S -sorted signature Σ , which corresponds to the syntax of the language shown in Example 2.1.1.

There are cases when we may want to refer to a specific part of the signature corresponding to a language, for instance, the one which includes only the *data* sorts in the language (e.g. integers, booleans, ...). We call that part a *subsignature*:

Definition 2.1.3 (Subsignature) Given S and S' two sets of sorts such that $S' \subseteq S$, an S' -sorted signature Σ' is a subsignature of an S -sorted signature Σ if $\Sigma' \subseteq \Sigma$ as $S^* \times S$ -indexed sets.

Remark 2.1.1 The subsignature Σ' can be regarded as an S -sorted signature where $\Sigma'_{w,s} = \emptyset$ for any $(w, s) \in (S^* \times S) \setminus (S'^* \times S')$.

Example 2.1.3 Recall the simple language from Example 2.1.2 and its corresponding signature Σ . The signature Σ' containing only the indexed sets:

- $\Sigma_{[], Bool} = \{\underline{true}, \underline{false}\}$ and
- $\Sigma_{Bool, Bool} = \{\neg_{Bool}\}$
- $\Sigma_{Bool Bool, Bool} = \{- \wedge_{Bool} -, - \vee_{Bool} -, \dots\}$

is a subsignature of Σ .

Signatures specify the syntax of programming languages but we are also interested in their semantics, that is, the entities of different sorts and particular functions that interpret the function symbols from signatures. For this we define Σ -algebras:

Definition 2.1.4 (Σ -algebra) *Given Σ an S -sorted signature, a Σ -algebra M consists of an S -indexed set (also denoted M), i.e., a carrier set M_s for each sort $s \in S$, plus:*

- an element $M_c \in M_s$ for each $c \in \Sigma_{[], s}$ interpreting the constant symbol c as an actual element, and
- a function $M_f : M_{s_1} \times \dots \times M_{s_n} \rightarrow M_s$ for each $f \in \Sigma_{w, s}$, where $w = s_1 \dots s_n$ ($n > 0$), interpreting each operation symbol as a function.

A Σ -algebra interprets every language construct (symbol in signature Σ) either as a constant or as a function. In this dissertation we assume that the carrier sets M_s are disjoint.

Example 2.1.4 *Consider the many sorted signature Σ shown in Example 2.1.1. A Σ -algebra M includes the carrier sets M_{Int} (which is in fact the set of integers), M_{Bool} (the set of boolean values, i.e. true and false), and M_{Exp} . For every constant symbol we have a corresponding element in one these carrier sets, e.g. $M_{\underline{1}} = 1 \in M_{Int}$. We always interpret operation symbols $\iota_{Sort', Sort}$ as injections $M_{Sort'} \rightarrow M_{Sort}$. For instance, we have the injection $M_{\iota_{Int, Exp}} : M_{Int} \rightarrow M_{Exp}$. We often write I for $M_{\iota_{Int, Exp}}(I)$ whenever the injection is deduced from the context. The algebra M also contains functions interpreting all operation symbols, i.e. the symbol $-+_-$ is interpreted by a function $- \dagger_{Int} - : Int \times Int \rightarrow Int$ which is the addition over integers.*

The terms over a signature Σ form a special algebra, called term algebra. In the following, we define the set of Σ -terms with variables, and the corresponding algebra. The set of variables, denoted Var , is an infinite S -indexed set of symbols disjoint from Σ .

Definition 2.1.5 (Σ -terms) *Let Σ be an S -sorted signature and Var an S -indexed set of variables disjoint from Σ . The S -indexed set $T_\Sigma(Var) = \{T_{\Sigma,s}(Var) \mid s \in S\}$ of Σ -terms t is defined by:*

$$t ::= c \mid X \mid f_{\underline{}}(t, \dots, t),$$

where c ranges over constant symbols from $\Sigma_{[],s}$, X ranges over variables from Var_s , and f over operation symbols from $\Sigma_{s_1 \dots s_n, s}$.

Definition 2.1.6 (Term algebra) *Given an S -sorted signature Σ and an S -indexed set of variables Var , the term Σ -algebra $T_\Sigma(Var)$ has Σ -terms $T_\Sigma(Var)$ as carrier sets and interprets each constant symbol c by itself and each operation symbol $f : s_1 \dots s_n \rightarrow s$ by the function $T_{s_1} \times \dots \times T_{s_n} \rightarrow T_s$ that maps (t_1, \dots, t_n) into the term $f_{\underline{}}(t_1, \dots, t_n)$.*

Example 2.1.5 *Let Σ be the signature shown in Example 2.1.2. Then $true$, $_{\underline{}} +_{\underline{}}(2, 3)$, and $_{\underline{}} *_{\underline{}}(_{\underline{}} +_{\underline{}}(2, 3), x)$, where $X \in Var_{Exp}$, are all Σ -terms interpreted as themselves, i.e. $_{\underline{}} true$, $_{\underline{}} +_{\underline{}}(2, 3)$, and $_{\underline{}} *_{\underline{}}(_{\underline{}} +_{\underline{}}(2, 3), X)$ respectively, in the corresponding term algebra.*

We often use only $($ and $)$ instead of $_{\underline{}}$ and $_{\underline{}}$, respectively. The Σ -algebra T_Σ of *ground* Σ -terms is the Σ -algebra $T_\Sigma(\emptyset)$ of terms without variables.

The *subterms* of a term can be referred using *positions*, which are lists of integers. The *subterm of term t at position p* is denoted by $t|_p$, and is defined inductively as follows:

- $t|_{[]} = t$, where $_{\underline{}}$ is the empty list, and
- $f(t_1, \dots, t_n)|_{iq} = t_i|_q$, if $1 \leq i \leq n$.

Note that constants and variables do not have proper subterms. We denote by $t[u]_p$ the term obtained from t by *replacing the subterm at position p by u* , i.e.:

- $t[u]_{\emptyset} = u$, and
- $f(t_1, \dots, t_n)[u]_{iq} = f(t_1, \dots, t_i[u]_q, \dots, t_n)$, if $1 \leq i \leq n$.

Example 2.1.6 Let $t \triangleq _ * _ (_ + _ (Y, 3), X)$. Then, $t|_{\emptyset} = t$, $t|_{12} = 3$, $t[Z]_{\emptyset} = Z$, and $t[Z]_2 = _ * _ (_ + _ (Y, 3), Z)$.

By $vars(t)$ we denote the set of variables occurring in the term t , $vars : T_{\Sigma}(Var) \rightarrow 2^{Var}$ being defined as follows:

- $vars(X) = \{X\}$, $X \in Var$,
- $vars(c) = \emptyset$, c is a constant, and
- $vars(f(t_1, \dots, t_n)) = vars(t_1) \cup \dots \cup vars(t_n)$.

For terms t_1, \dots, t_n we let $vars(t_1, \dots, t_n) \triangleq vars(t_1) \cup \dots \cup vars(t_n)$.

Example 2.1.7 If $t \triangleq _ * _ (_ + _ (Y, 3), X)$ then $vars(t) = \{X, Y\}$.

Let M be a Σ -algebra. A *valuation* ρ is a function $\rho : Var \rightarrow M$ which maps variables to values from M . Any valuation ρ can be extended to $\rho^{\#} : T_{\Sigma}(Var) \rightarrow M$ as follows:

- $\rho^{\#}(X) = \rho(X)$, $X \in Var$,
- $\rho^{\#}(c) = M_c$, and
- $\rho^{\#}(f(t_1, \dots, t_n)) = M_f(\rho^{\#}(t_1), \dots, \rho^{\#}(t_n))$.

For simplicity, in the rest of the dissertation, we use ρ instead of its extension $\rho^{\#}$.

Example 2.1.8 Let $t \triangleq _ * _ (+ (Y, Z), X)$ and $\rho : Var \rightarrow M$ a valuation such that $\rho(X) = 1$, $\rho(Y) = 2$, and $\rho(Z) = 3$, where $\{1, 2, 3\} \subseteq M_{Int}$. Then, $\rho(t) \triangleq \rho^{\#}(t) = M_*(M_+(2, 3), 1)$.

Given two valuations $\rho : V \rightarrow M$ and $\rho' : V' \rightarrow M$, where $V, V' \subseteq \text{Var}$, such that for any $X \in V \cap V'$ we have $\rho(X) = \rho'(X)$, we define the valuation $\rho \uplus \rho' : V \cup V' \rightarrow M$ as

$$(\rho \uplus \rho')(X) = \begin{cases} \rho(X) & , X \in V, \\ \rho'(X) & , X \in V'. \end{cases}$$

Note that the valuation $\rho \uplus \rho'$ is well defined since $\rho(X) = \rho'(X)$ for all $X \in V \cap V'$.

A *substitution* is a valuation $\sigma : \text{Var} \rightarrow T_\Sigma(\text{Var})$, i.e., a valuation where M is the algebra of Σ -terms with variables. The identity substitution is denoted $1_V : V \rightarrow V$, and $1_V(X) = X$, for all $X \in V (\subseteq \text{Var})$.

2.2 Many-Sorted First Order Logic

This section contains an overview of many-sorted First Order Logic, abbreviated FOL in this dissertation. FOL is a formal system which provides a mechanism to express properties about objects together with their logical relationships and dependences.

The basic ingredients of the FOL language are terms, which can be constants, variables or complex terms built by applying functions to other terms. In addition to having terms, FOL contains a set of predicate symbols and quantifiers. Many-sorted FOL is obtained by adding to the FOL formalism the notion of *sort*. Thus, in contrast to usual (unsorted) FOL, the arguments of functional and predicate symbols may have different sorts, and constant and functional symbols also have particular sorts.

In the following we formally define the syntax of many-sorted FOL:

Definition 2.2.1 (Many-Sorted First-Order Signature) *Given S a set of sorts, an S -sorted first order signature Φ is a pair (Σ, Π) , where Σ is an S -sorted signature and Π is an indexed set of the form $\{\Pi_w \mid w \in S^*\}$ whose elements are called predicate symbols, where $\pi \in \Pi_w$ is said to have arity w .*

Next, we define the syntax of FOL formulas over a first order signature $\Phi = (\Sigma, \Pi)$. We let S denote the set of sorts in Φ and Var an infinite S -indexed set of variable symbols (disjoint from Σ).

Definition 2.2.2 (FOL formula) *The set of Φ -formulas is defined by*

$$\phi ::= \top \mid p(t_1, \dots, t_n) \mid \neg\phi \mid \phi \wedge \phi \mid (\exists V)\phi$$

where p ranges over predicate symbols Π , each t_i ranges over $T_\Sigma(\text{Var})$ of appropriate sort, and V over finite subsets of Var .

The other known connectives, \vee , \rightarrow , and the quantifier \forall can be expressed using the ones above:

- $\phi \vee \phi' \triangleq \neg(\neg\phi \wedge \neg\phi')$,
- $\phi \rightarrow \phi' \triangleq \neg\phi \vee \phi'$, and
- $(\forall V)\phi \triangleq \neg(\exists V)\neg\phi$

For convenience, we often use parenthesis for disambiguation or scoping. The syntactical constructs have the following priorities, in decreasing order, starting with the ones which bind tighter:

1. \neg, \exists, \forall ,
2. \wedge, \vee ,
3. \rightarrow .

In a FOL formula a variable may occur *free* or *bound*.

Definition 2.2.3 (Free variables) *A variable X is free in a FOL formula ϕ if:*

- $\phi \triangleq \top$;
- $\phi \triangleq p(t_1, \dots, t_n)$;
- $\phi \triangleq \neg\psi$ and X is free in the FOL formula ψ ;
- $\phi \triangleq \psi_1 \wedge \psi_2$ and X is free in both FOL formulas ψ_1 and ψ_2 ;
- $\phi \triangleq (\exists V)\psi$ and $X \notin V$ and X is free in the FOL formula ψ .

An occurrence of a variable is bound in a formula if it is in the scope of a quantifier.

Definition 2.2.4 (Bound variables) A variable X is bound in a FOL formula ϕ if $\phi \triangleq (\exists V)\psi$ and $X \in V$.

Example 2.2.1 In the formula $(\forall X)(X \wedge Y \vee \pi(Z))$ the variable X is bound, while Y and Z are free.

Definition 2.2.5 A substitution $\sigma : Var \rightarrow T_\Sigma(Var)$ can be extended to FOL formulas:

- $\sigma^\#(\top) = \top$;
- $\sigma^\#(p(t_1, \dots, t_n)) = p(\sigma^\#(t_1), \dots, \sigma^\#(t_n))$;
- $\sigma^\#(\neg\phi) = \neg\sigma^\#(\phi)$;
- $\sigma^\#(\phi_1 \wedge \phi_2) = \sigma^\#(\phi_1) \wedge \sigma^\#(\phi_2)$;
- $\sigma^\#((\exists V)\phi) = (\exists V)\sigma'^\#(\phi)$,

where $\sigma' : Var \rightarrow T_\Sigma(Var)$, $\sigma'(Y) = Y$ when $Y \in V$, and $\sigma'(Y) = \sigma(Y)$ when $Y \notin V$.

Notice that a substitution σ automatically avoids substituting for bound variables.

Example 2.2.2 If σ is a substitution, where $\sigma(X) = Y$, $\sigma(Y) = Z$ then $\sigma((\forall X)(X \wedge \pi(Y))) = (\forall X)(X \wedge \pi(Z))$.

The set of free variables occurring in a FOL formula ϕ , denoted $vars(\phi)$, is defined below:

- $vars(\top) = \emptyset$;
- $vars(p(t_1, \dots, t_n)) = vars(t_1) \cup \dots \cup vars(t_n)$;
- $vars(\neg\phi) = vars(\phi)$;
- $vars(\phi_1 \wedge \phi_2) = vars(\phi_1) \cup vars(\phi_2)$;
- $vars((\exists V)\phi) = vars(\phi) \setminus V$.

Example 2.2.3 Given the FOL formula $F \triangleq (\forall X)(X \wedge Y \vee \pi(Z))$ we have $\text{vars}(F) = \{Y, Z\}$.

The truth value of a FOL formula depends on the choice of values for variables and the meaning of the operation and predicate symbols involved. Thus, we require a *model* of all operations and predicates.

Definition 2.2.6 (FOL Model) Given a FOL signature $\Phi = (\Sigma, \Pi)$, a Φ -model consists of a Σ -algebra M together with a subset $M_p \subseteq M_{s_1} \times \dots \times M_{s_n}$ for each predicate $p \in \Pi_w$, where $w = s_1 \dots s_n$.

Definition 2.2.7 (FOL Satisfaction relation) Given a first order Φ -model M , a Φ -formula ϕ , and a valuation $\rho : \text{Var} \rightarrow M$, the satisfaction relation $\rho \models_{\text{FOL}} \phi$ is defined as follows:

1. $\rho \models_{\text{FOL}} \top$;
2. $\rho \models_{\text{FOL}} p(t_1, \dots, t_n)$ iff $(\rho(t_1), \dots, \rho(t_n)) \in M_p$;
3. $\rho \models_{\text{FOL}} \neg\phi$ iff $\rho \models_{\text{FOL}} \phi$ does not hold;
4. $\rho \models_{\text{FOL}} \phi_1 \wedge \phi_2$ iff $\rho \models_{\text{FOL}} \phi_1$ and $\rho \models_{\text{FOL}} \phi_2$;
5. $\rho \models_{\text{FOL}} (\exists V)\phi$ iff there is $\rho' : \text{Var} \rightarrow M$ with $\rho'(X) = \rho(X)$, for all $X \notin V$, such that $\rho' \models_{\text{FOL}} \phi$.

A formula ϕ is valid in M , denoted by $M \models_{\text{FOL}} \phi$, if it is satisfied by all valuations ρ .

When the model M is clearly understood from the context we only write $\models_{\text{FOL}} \phi$ instead of $M \models_{\text{FOL}} \phi$, to denote that ϕ is valid in M . As a remark, in this dissertation we do not use $\models_{\text{FOL}} \phi$ as notation for universally valid formulas since we work with a fixed model M .

2.3 Matching Logic

Matching Logic (ML) was first introduced in [94] as a novel framework for defining axiomatic semantics for programming languages. In general, logics designed for program reasoning do not concern about the low-level

details of program configurations. In contrast, ML has the program configuration at its core, capturing information about the structure of the various data in a program's state. ML specifications are symbolic program configurations constrained by a FOL formula, called *patterns*. A concrete configuration satisfies a pattern if it *matches* the symbolic configuration of the pattern and satisfies the FOL constraint. We recall here the syntax and the semantics of Matching Logic as presented in [98].

Definition 2.3.1 (ML Signature) *An ML signature $\Phi = (\Sigma, \Pi, Cfg)$ is a first-order signature (Σ, Π) together with a distinguished sort Cfg for configurations.*

The sort Cfg is intended to model program configurations. The configuration terms may contain informations about the program state (the heap, the stack, the input, the output, etc.) and appear in ML formulas.

Example 2.3.1 *Let us consider a simple program configuration which is a pair $\langle\langle\text{code}\rangle\langle m\rangle\rangle$ with `code` a fragment of program and m a mapping from program variables to integers. Then, $\langle\langle\mathbf{x} = 2;\rangle\langle\mathbf{x} \mapsto 5\rangle\rangle$ and $\langle\langle\text{if } \mathbf{x} \leq 2 \text{ then } \{\mathbf{x} = 0;\}\rangle\langle\mathbf{x} \mapsto 1\rangle\rangle$ are terms of sort Cfg .*

In the rest of the section, we consider an ML-signature $\Phi = (\Sigma, \Pi, Cfg)$, and we let Var be an S -indexed set of variables (disjoint from Σ). We now define the syntax of ML formulas over the ML-signature Φ .

Definition 2.3.2 (ML formula) *The set of ML-formulas is defined by*

$$\varphi ::= \pi \mid \top \mid p(t_1, \dots, t_n) \mid \neg\varphi \mid \varphi \wedge \varphi \mid (\exists V)\varphi$$

where π ranges over $T_{\Sigma, Cfg}(Var)$, p ranges over predicate symbols Π , each t_i ranges over $T_{\Sigma}(Var)$ of appropriate sorts, and V over finite subsets of Var .

The syntax of ML formulas is quite similar with the syntax of FOL formulas, except that the former allows configuration terms $\pi \in T_{\Sigma, Cfg}(Var)$ to be also formulas. The ML formulas can be encoded in FOL [98] (Definition 2.3.4), but in the dissertation we prefer to give a direct semantics. In addition to the above syntax we also allow \vee , \rightarrow , and \forall in ML formulas,

which can be expressed in terms of the existent ones in the same way as we did for FOL (see Section 2.2).

We distinguish the two particular types of ML formulas: *basic patterns* and (elementary) *patterns*. A *basic pattern* is a term $\pi \in T_{\Sigma, Cfg}(Var)$. A *pattern* is an ML formula of the form $\pi \wedge \phi$, where π is a basic pattern and ϕ is a FOL formula. The basic pattern π is a configuration term with variables, and thus, it defines a set of (concrete) configurations. The condition ϕ gives additional constraints these configurations must satisfy.

Example 2.3.2 Consider the program configuration $\langle\langle\text{code}\rangle\langle m\rangle\rangle$ shown in Example 2.3.1. Then, $\pi \triangleq \langle\langle\text{if } x > 0 \{ x = x + 1; \}\rangle\langle x \mapsto X\rangle\rangle$ is a basic pattern, where x is a program variable and X is a logical variable, and $\pi \wedge \phi \triangleq \langle\langle\text{if } x > 0 \{ x = x + 1; \}\rangle\langle x \mapsto X\rangle\rangle \wedge X > 0$ is a pattern, which consists of a basic pattern plus an additional constraint over X .

We now define the ML satisfaction relation.

Definition 2.3.3 (ML satisfaction relation) Given an ML signature $\Phi = (\Sigma, \Pi, Cfg)$, M a (Σ, Π) -model, ϕ a ML formula, $\gamma \in M_{Cfg}$ a concrete configuration, and $\rho : Var \rightarrow M$ a valuation, then the satisfaction relation $(\gamma, \rho) \models_{ML} \varphi$ is defined as follows:

1. $(\gamma, \rho) \models_{ML} \pi$ iff $\rho(\pi) = \gamma$;
2. $(\gamma, \rho) \models_{ML} \top$;
3. $(\gamma, \rho) \models_{ML} p(t_1, \dots, t_n)$ iff $(\rho(t_1), \dots, \rho(t_n)) \in M_p$;
4. $(\gamma, \rho) \models_{ML} \neg\varphi$ iff $(\gamma, \rho) \models_{ML} \varphi$ does not hold;
5. $(\gamma, \rho) \models_{ML} \varphi_1 \wedge \varphi_2$ iff $(\gamma, \rho) \models_{ML} \varphi_1$ and $(\gamma, \rho) \models_{ML} \varphi_2$; and
6. $(\gamma, \rho) \models_{ML} (\exists V)\varphi$ iff there is $\rho' : Var \rightarrow M$ with $\rho'(X) = \rho(X)$, for all $X \notin V$, such that $(\gamma, \rho') \models_{ML} \varphi$.

An ML formula ϕ is valid, written $\models_{ML} \varphi$, when $(\gamma, \rho) \models \varphi$ for all $\gamma \in T_{\Sigma, Cfg}$ and all $\rho : Var \rightarrow M$.

In ML, a basic pattern π is satisfied by (γ, ρ) , where γ is a configuration term and ρ a valuation, if $\rho(\pi) = \gamma$. The valuation ρ actually gives a possible *match* of configuration γ against basic pattern π , wherefrom the name of the logic: Matching Logic.

Example 2.3.3 Consider the program configuration $\langle\langle\text{code}\rangle\langle m\rangle\rangle$ shown in Example 2.3.1. If $\pi \triangleq \langle\langle C\rangle\langle \mathbf{x} \mapsto X\rangle\rangle$ is a basic pattern (C and X are variables), then the configuration $\gamma \triangleq \langle\langle\text{skip}\rangle\langle \mathbf{x} \mapsto 7\rangle\rangle$ matches against π since there exists $\rho : \text{Var} \rightarrow M$, where $\rho(X) = 7$ and $\rho(C) = \text{skip}$, such that $\gamma = \rho(\pi)$. Moreover, γ also matches against $\pi \wedge \phi \triangleq \langle\langle C\rangle\langle \mathbf{x} \mapsto X\rangle\rangle \wedge X \geq 0$ using the same valuation ρ , since $\gamma = \rho(\pi)$, and $\rho \models_{\text{FOL}} \phi$, since $7 \geq 0$.

Next, we recall from [98] how pattern reasoning in ML reduces to FOL reasoning in the (Σ, Π) -model M :

Definition 2.3.4 (FOL encoding of ML) Let \square be a special fresh Cfg variable, such that $\square \notin \text{Var}$, and $\text{Var}^\square = \text{Var} \cup \{\square\}$. For a pattern φ , let φ^\square be the FOL formula obtained by replacing basic patterns $\pi \in T_{\Sigma, \text{Cfg}}(\text{Var})$ with equalities $\pi = \square$. If $\gamma \in M_{\text{Cfg}}$ and $\rho : \text{Var} \rightarrow M$, then let $\rho^\gamma : \text{Var}^\square \rightarrow M$ be a valuation, such that $\rho^\gamma(X) = \rho(X)$ for all $X \in \text{Var}$ and $\rho^\gamma(\square) = \gamma$.

Using the notation from Definition 2.3.4, we have $(\gamma, \rho) \models_{\text{ML}} \varphi$ iff $\rho^\gamma \models_{\text{FOL}} \varphi^\square$, which is proved by structural induction on φ (the complete proof can be found in [99]). This result allows us to extend all the constructions over FOL formulas to ML formulas.

The set of concrete configurations that match against a specific pattern is defined below.

Definition 2.3.5 If φ is a ML formula, then $\llbracket \varphi \rrbracket$ denotes the set of concrete configurations $\{\gamma \mid \text{there exists } \rho \text{ s.t. } (\gamma, \rho) \models_{\text{ML}} \varphi\}$. For any set of ML formulas F we let $\llbracket F \rrbracket$ denote the set $\bigcup_{\varphi \in F} \llbracket \varphi \rrbracket$.

Remark 2.3.1 Let us consider $\varphi \triangleq \langle\langle \text{if } B \{ S \} \rangle\langle E \rangle\rangle \wedge (B = \text{true})$ and $\varphi' \triangleq (\exists B) \langle\langle \text{if } B \{ S \} \rangle\langle E \rangle\rangle \wedge (B = \text{true})$ two ML formulas, with B , S , and E from Var . We can observe that $\llbracket \varphi \rrbracket = \llbracket \varphi' \rrbracket$, since $\gamma \in \llbracket \varphi \rrbracket$ obviously implies $\gamma \in \llbracket \varphi' \rrbracket$, and $\gamma \in \llbracket \varphi' \rrbracket$ implies $\gamma \in \llbracket \varphi \rrbracket$ (by taking $\rho(B) = \text{true}$). Thus, by the semantics of $\llbracket _ \rrbracket$, φ and φ' are equivalent. However, φ^\square and φ'^\square are not equivalent as FOL formulas, because of the existential quantifier in φ'^\square (i.e. consider a valuation $\rho : \text{Var} \rightarrow M$, such that $\rho(B) = \text{false}$ and $\rho \models_{\text{FOL}} \varphi'^\square$; then $\rho \not\models_{\text{FOL}} \varphi^\square$).

The next lemma shows that a formula $(\exists V)\varphi$, $V \subseteq \text{vars}(\varphi)$, matches the same set of concrete configurations as φ .

Lemma 2.3.1 *Let φ be an ML formula. Then $\llbracket (\exists V)\varphi \rrbracket = \llbracket \varphi \rrbracket$.*

Proof We prove $\llbracket (\exists V)\varphi \rrbracket \subseteq \llbracket \varphi \rrbracket$ and $\llbracket (\exists V)\varphi \rrbracket \supseteq \llbracket \varphi \rrbracket$:

(\subseteq) Let $\gamma \in \llbracket (\exists V)\varphi \rrbracket$. By Definition 2.3.5, there exists $\rho : \text{Var} \rightarrow M$, such that $(\gamma, \rho) \models_{\text{ML}} (\exists V)\varphi$. Then, there exists $\rho' : \text{Var} \rightarrow M$, with $\rho'(Y) = \rho(Y)$, for all $Y \notin V$ such that $(\gamma, \rho') \models_{\text{ML}} \varphi$ (cf. Definition 2.3.3). Therefore, there exists a valuation ρ' such that $\gamma \in \llbracket \varphi \rrbracket$.

(\supseteq) Let $\gamma \in \llbracket \varphi \rrbracket$. By Definition 2.3.5, there exists $\rho : \text{Var} \rightarrow M$, such that $(\gamma, \rho) \models_{\text{ML}} \varphi$. We have to show that there is a valuation $\rho' : \text{Var} \rightarrow M$ such that $(\gamma, \rho') \models_{\text{ML}} (\exists V)\varphi$. Let $\rho' = \rho$. Then, we have $\rho(Y) = \rho'(Y)$ for all $Y \notin V$ and $(\gamma, \rho) \models_{\text{ML}} \varphi$. Thus, cf. Definition 2.3.3, $(\gamma, \rho) \models_{\text{ML}} (\exists V)\varphi$ and $\gamma \in \llbracket (\exists V)\varphi \rrbracket$.

■

The set of free variables occurring in an ML formula $\varphi \triangleq \pi \wedge \phi$ is $\text{vars}(\varphi) \triangleq \text{vars}(\pi) \cup \text{vars}(\phi)$, where $\text{vars}(\pi)$ and $\text{vars}(\phi)$ are defined in Section 2.1 and Section 2.2, respectively. For some of our subsequent results we are also interested well-definedness of patterns:

Definition 2.3.6 ((Weakly) well-defined patterns) *Let φ be a pattern. Then φ is weakly well-defined iff for any valuation $\rho : \text{Var} \rightarrow \mathcal{T}$ there is some configuration $\gamma \in \mathcal{T}_{\text{Cfg}}$ such that $(\gamma, \rho) \models_{\text{ML}} \varphi$, and φ is well-defined if γ is unique.*

Example 2.3.4 *Let $\varphi \triangleq \langle \langle \text{if } B \{ S \} \rangle \langle E \rangle \rangle \wedge (B = \text{true})$. Then, the pattern φ is weakly well-defined because there are a configuration $\gamma \triangleq \langle \langle \text{if true } \{ \text{skip}; \} \rangle \langle \mathbf{x} \mapsto 0 \rangle \rangle$ and a valuation $\rho : \text{Var} \rightarrow M$, where $\rho(B) = \text{true}$, $\rho(S) = \text{skip}$, and $\rho(E) = \mathbf{x} \mapsto 0$, such that $(\gamma, \rho) \models_{\text{ML}} \varphi$. The pattern $\varphi' \triangleq \langle \langle \text{if } B \{ \text{skip}; \} \rangle \langle \mathbf{x} \mapsto 0 \rangle \rangle \wedge (B = \text{true})$ is well defined, since there is an unique configuration $\gamma' = \gamma$ such that $(\gamma', \rho) \models_{\text{ML}} \varphi'$. On the other hand, the pattern $\varphi'' \triangleq \langle \langle \text{if } B \{ S \} \rangle \langle E \rangle \rangle \wedge \text{false}$ is not weakly well-defined because it is impossible to find a configuration γ'' and valuation ρ'' such that $(\gamma'', \rho'') \models_{\text{ML}} \varphi''$.*

2.4 Reachability Logic

Reachability Logic [98, 101, 100, 93] is a logic designed to use operational semantics for program verification and reasoning about programs. It includes a *language-independent* proof system for deriving reachability properties, which takes any operational semantics of a programming language, and derives any program properties that can be derived with *language-specific* proof systems (e.g. Hoare logics or dynamic logics). The main ingredients of RL are *reachability rules*, also called RL formulas, which generalise both term-rewrite rules and Hoare triples. The proof system derives reachability rules using the trusted operational semantics of the target language. Thus, the verification of a program reduces to checking if its specification (given as a reachability rule) is derivable.

In contrast to ML, which is a static logic of program configurations designed to state and reason about their structural properties, RL can be seen as a dynamic logic of configurations, expressing their evolution over time. In Section 2.3 we showed how one ML pattern φ specifies all the configurations γ that match against it. Here, specifications (i.e. reachability rules) are pairs of patterns, written $\varphi \Rightarrow \varphi'$.

In the following, we let \mathcal{S} be a set of reachability rules and an underlying ML signature (Σ, Π, Cfg) . \mathcal{S} induces a transition system on any Σ -algebra/model M . Let us fix an arbitrary model M , which we call *configuration model*, and a sort-wise set of variables Var . We now formally define reachability rules and the transition system induced by \mathcal{S} .

Definition 2.4.1 (RL formula) *A reachability rule or a RL formula is an expression of the form $\varphi \Rightarrow \varphi'$ where φ and φ' are ML formulas. A reachability system is a set of reachability rules.*

A reachability rule $\varphi \Rightarrow \varphi'$ specifies that configurations that match against φ eventually transit to configurations that match against φ' . In other words, the transitions over program configurations are instances of reachability rules.

Example 2.4.1 *Let us consider the program configuration $\langle\langle\text{code}\rangle\langle m\rangle\rangle$ shown in Example 2.3.1. Then, $\langle\langle\text{x} = I; \text{skip}\rangle\langle\text{x} \mapsto J\rangle\rangle \Rightarrow \langle\langle\text{skip}\rangle\langle\text{x} \mapsto I\rangle\rangle$ is a reachability rule specifying that configurations that match against the*

pattern $\langle\langle \mathbf{x} = I; \mathbf{skip} \rangle\langle \mathbf{x} \mapsto J \rangle\rangle$, (e.g. $\langle\langle \mathbf{x} = 5; \mathbf{skip} \rangle\langle \mathbf{x} \mapsto 10 \rangle\rangle$) transit to configurations that match against $\langle\langle \mathbf{skip} \rangle\langle \mathbf{x} \mapsto I \rangle\rangle$ (e.g. $\langle\langle \mathbf{skip} \rangle\langle \mathbf{x} \mapsto 5 \rangle\rangle$).

Any reachability system \mathcal{S} induces a transition system on the configuration model:

Definition 2.4.2 (RL transition system) \mathcal{S} induces a transition system $(M_{Cfg}, \Rightarrow_{\mathcal{S}})$, where $\Rightarrow_{\mathcal{S}} \subseteq M_{Cfg} \times M_{Cfg}$ is defined by $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ iff there is $\varphi \Rightarrow \varphi'$ in \mathcal{S} and $\rho : Var \rightarrow M$ with $(\gamma, \rho) \models_{ML} \varphi$ and $(\gamma', \rho) \models_{ML} \varphi'$.

The transition system $(M_{Cfg}, \Rightarrow_{\mathcal{S}})$ contains the set of all transitions $\gamma \Rightarrow_{\mathcal{S}} \gamma'$, and thus, it captures precisely how the language defined by \mathcal{S} operates. We extend the notation $\llbracket _ \rrbracket$ to rules, $\llbracket \alpha \triangleq \varphi \Rightarrow \varphi' \rrbracket = \{\gamma \Rightarrow_{\{\alpha\}} \gamma' \mid (\exists \rho : Var \rightarrow M)((\gamma, \rho) \models_{ML} \varphi \wedge (\gamma', \rho) \models_{ML} \varphi')\}$. The next lemma shows that a RL formula of the form $(\exists X)\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$ generates the same transitions as $\pi \wedge \phi \Rightarrow \pi' \wedge \phi'$.

Lemma 2.4.1 *If $X \cap Y = \emptyset$ then $\llbracket (\exists X)\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi' \rrbracket = \llbracket \pi \wedge \phi \Rightarrow \pi' \wedge \phi' \rrbracket$.*

Proof We may assume without loss of generality that $X \cap vars(\pi', \phi') = \emptyset$ and $Y \cap vars(\pi, \phi) = \emptyset$. Note that the condition $X \cap Y = \emptyset$ from the hypothesis of the lemma, can be easily obtained by variable renaming.

(\subseteq) Assume $\gamma \Rightarrow_{\mathcal{S}} \gamma' \in \llbracket (\exists X)\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi' \rrbracket$. Then, there is ρ such that $(\gamma, \rho) \models_{ML} (\exists X)\pi \wedge \phi$ and $(\gamma', \rho) \models_{ML} (\exists Y)\pi' \wedge \phi'$. By Definition 2.3.3, there are ρ_X and ρ_Y such that $(\gamma, \rho_X) \models_{ML} \pi \wedge \phi$ and $\rho_X(z) = \rho(z)$ for all $z \in Var \setminus X$, and $(\gamma', \rho_Y) \models_{ML} \pi' \wedge \phi'$ and $\rho_Y(z) = \rho(z)$ for all $z \in Var \setminus Y$. Let $\rho' = \rho_X \uplus \rho_Y$. Note that ρ' is well-defined since ρ_X and ρ_Y coincide on $Var \setminus (X \cup Y)$. Obviously, we have $(\gamma, \rho') \models_{ML} \pi \wedge \phi$ (ρ' and ρ_X coincide on $vars(\pi, \phi)$) and $(\gamma', \rho') \models_{ML} \pi' \wedge \phi'$ (ρ' and ρ_Y coincide on $vars(\pi', \phi')$). So, $\gamma \Rightarrow_{\mathcal{S}} \gamma' \in \llbracket \pi \wedge \phi \Rightarrow \pi' \wedge \phi' \rrbracket$.

(\supseteq) Assume $\gamma \Rightarrow_{\mathcal{S}} \gamma' \in \llbracket \pi \wedge \phi \Rightarrow \pi' \wedge \phi' \rrbracket$. There is ρ such that $(\gamma, \rho) \models_{ML} \pi \wedge \phi$ and $(\gamma', \rho) \models_{ML} \pi' \wedge \phi'$. Indeed, we have $(\gamma, \rho) \models_{ML} (\exists X)\pi \wedge \phi$ and $(\gamma', \rho) \models_{ML} (\exists Y)\pi' \wedge \phi'$.

■

Lemma 2.4.1 allows us to omit explicitly the writing of the existential quantifiers in RL formulas.

In the following lemma we show that the transition system does not change if, in a rule $\pi \wedge \phi \Rightarrow \pi' \wedge \phi'$, we replace a subterm in π with a variable, and we add the equality between the variable and the replaced term to ϕ . If $\alpha \triangleq \varphi \Rightarrow \varphi'$, we denote by $\text{vars}(\alpha)$ the set $\text{vars}(\varphi) \cup \text{vars}(\varphi')$.

Lemma 2.4.2 *Let $\alpha \triangleq \pi \wedge \phi \Rightarrow \pi' \wedge \phi'$ be a rule in \mathcal{S} , p a position in π , and a rule $\alpha' \triangleq \pi[X]_p \wedge (\phi \wedge X = \pi|_p) \Rightarrow \pi' \wedge \phi'$, where $X \notin \text{vars}(\alpha)$ is a fresh variable which has the same sort as $\pi|_p$. Then $\llbracket \alpha \rrbracket = \llbracket \alpha' \rrbracket$.*

Proof First, we prove that for all valuations $\rho : \text{Var} \rightarrow M$ and configurations γ , such that $(\gamma, \rho) \models_{\text{ML}} \pi[X]_p \wedge (\phi \wedge X = \pi|_p)$ we have $\rho(X) = \pi|_p$ (\spadesuit). Proof by contradiction: let ρ be a valuation and γ a configuration such that $(\gamma, \rho) \models_{\text{ML}} \pi[X]_p \wedge (\phi \wedge X = \pi|_p)$ and $\rho(X) \neq \pi|_p$. From $(\gamma, \rho) \models_{\text{ML}} \pi[X]_p \wedge (\phi \wedge X = \pi|_p)$ we obtain $\rho \models_{\text{ML}} (\phi \wedge X = \pi|_p)$. Since $X \notin \text{vars}(\alpha)$ (i.e. $X \notin \text{vars}(\phi)$) then $\rho \models_{\text{ML}} \phi$ and $\rho \models_{\text{ML}} X = \pi|_p$. But $X \notin \text{vars}(\pi)$, so the only possible valuation of X is $\rho(X) = \pi|_p$, which is a contradiction.

(\subseteq) Let $\gamma \Rightarrow \gamma' \in \llbracket \alpha \rrbracket$. Then, there exists $\rho : \text{Var} \rightarrow M$ such that $(\gamma, \rho) \models_{\text{ML}} \pi \wedge \phi$ and $(\gamma', \rho) \models_{\text{ML}} \pi' \wedge \phi'$. Let $\rho' : \text{Var} \rightarrow M$ be a valuation defined as $\rho'(Y) = \rho(Y)$ for $Y \in \text{Var} \setminus \{X\}$, and $\rho'(X) = \pi|_p$. Since $X \notin \text{vars}(\alpha)$ and $(\gamma', \rho) \models_{\text{ML}} \pi' \wedge \phi'$ we obtain $(\gamma', \rho') \models_{\text{ML}} \pi' \wedge \phi'$. Moreover, we have $\rho'(\pi[X]_p) = \rho'(\pi[\pi|_p]_p) = \rho'(\pi) = \rho(\pi) = \gamma$, and $\rho'(\phi \wedge X = \pi|_p) = \rho'(\phi \wedge \pi|_p = \pi|_p) = \rho'(\phi) = \rho(\phi) = \top$. Therefore, there exists $\rho' : \text{Var} \rightarrow M$ such that $(\gamma, \rho') \models_{\text{ML}} \pi[X]_p \wedge (\phi \wedge X = \pi|_p)$ and $(\gamma', \rho') \models_{\text{ML}} \pi' \wedge \phi'$, i.e. $\gamma \Rightarrow \gamma' \in \llbracket \alpha' \rrbracket$.

(\supseteq) Let $\gamma \Rightarrow \gamma' \in \llbracket \alpha' \rrbracket$. Then, there exists $\rho' : \text{Var} \rightarrow M$ such that $(\gamma, \rho') \models_{\text{ML}} \pi[X]_p \wedge (\phi \wedge X = \pi|_p)$ and $(\gamma', \rho') \models_{\text{ML}} \pi' \wedge \phi'$. Let $\rho : \text{Var} \rightarrow M$ be a valuation such that $\rho(Y) = \rho'(Y)$ for $Y \in \text{Var} \setminus \{X\}$. Then $\gamma = \rho'(\pi[X]_p) = \rho'(\pi[\pi|_p]_p) = \rho'(\pi) = \rho(\pi)$ (we used (\spadesuit) and the fact that $X \notin \text{vars}(\pi)$) and $\top = \rho'(\phi \wedge X = \pi|_p) = \rho'(\phi) = \rho(\phi)$ (we used (\spadesuit) and the fact that $X \notin \text{vars}(\pi, \phi)$). Therefore, there exists $\rho : \text{Var} \rightarrow M$, such that $(\gamma, \rho) \models_{\text{ML}} \pi \wedge \phi$ and $(\gamma', \rho) \models_{\text{ML}} \pi' \wedge \phi'$, i.e. $\gamma \Rightarrow \gamma' \in \llbracket \alpha \rrbracket$. \blacksquare

The following notions are used in the definition of the satisfaction relation over transition systems and RL formulas.

Definition 2.4.3 A configuration $\gamma \in M_{Cf_g}$ terminates in $(M_{Cf_g}, \Rightarrow_S)$ iff there is no infinite sequence $\gamma \Rightarrow_S \gamma_1 \Rightarrow_S \gamma_2 \Rightarrow_S \dots$.

Definition 2.4.4 (Weakly well-defined) A reachability rule $\varphi \Rightarrow \varphi'$ is weakly well-defined, respectively well-defined, if φ' is weakly well-defined, respectively well-defined (cf. Definition 2.3.6).

RL formulas are used to express program properties. Semantic validity in RL follows the same line of partial correctness in Hoare Logic, but in more general terms of reachability:

Definition 2.4.5 (RL semantics) The satisfaction relation \models_{RL} between the transition systems $(M_{Cf_g}, \Rightarrow_S)$ and RL formulas $\varphi \Rightarrow \varphi'$ is defined as follows: $(M_{Cf_g}, \Rightarrow_S) \models_{\text{RL}} \varphi \Rightarrow \varphi'$ iff for all $\gamma \in M_{Cf_g}$ such that γ terminates in $(M_{Cf_g}, \Rightarrow_S)$ and for all $\rho : \text{Var} \rightarrow M$ such that $(\gamma, \rho) \models_{\text{ML}} \varphi$, there exists some $\gamma' \in M_{Cf_g}$ such that $(\gamma', \rho) \models_{\text{ML}} \varphi'$ and $\gamma \Rightarrow_S^* \gamma'$.

By \Rightarrow_S^* we denote the transitive closure of the relation \Rightarrow_S . Intuitively, $(M_{Cf_g}, \Rightarrow_S) \models_{\text{RL}} \varphi \Rightarrow \varphi'$ specifies reachability: any configuration γ that terminates and is matched against φ transits, on some execution path, to a configuration γ' that is matched against φ' .

Since the model M and the reachability system S are fixed, we often write $\mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$ instead of $(M_{Cf_g}, \Rightarrow_S) \models_{\text{RL}} \varphi \Rightarrow \varphi'$. We let $\models_{\text{FOL}} \varphi \rightarrow \varphi'$ denote the fact that the FOL implication between the FOL encoding of ML patterns φ and φ' is valid (see [98] for details).

The proof system of RL is shown in Figure 2.1. It derives sequents of the form $\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow \varphi'$, where \mathcal{S} and \mathcal{C} are sets of RL formulas. The rules in \mathcal{S} are called *axioms*, while the ones in \mathcal{C} are called *circularities*. When \mathcal{C} is empty (\emptyset) we write $\mathcal{S} \vdash \varphi \Rightarrow \varphi'$ instead of $\mathcal{S} \vdash_{\emptyset} \varphi \Rightarrow \varphi'$. During the proof, the set \mathcal{C} is populated with circularities via [Circularity]. The [Axiom] proof rule states that the rules in \mathcal{S} can be used as initial axioms. Both [Reflexivity] and [Transitivity] stand for the corresponding properties of the reachability relation, while [Consequence] and [CaseAnalysis] are adapted from Hoare Logic. [Abstraction] allows hiding irrelevant details of φ behind an existential quantifier. This deduction rule is particularly useful in combination with [Circularity] which has a co-inductive nature. [Circularity] is used, e.g., to prove properties about loops, recursive functions, jumps,

$$\begin{array}{c}
\text{[Axiom]} \quad \frac{\varphi \Rightarrow \varphi' \in \mathcal{S} \quad \phi \text{ is a (patternless) FOL formula}}{\mathcal{S} \vdash_{\mathcal{C}} \varphi \wedge \phi \Rightarrow \varphi' \wedge \phi} \\
\text{[Abstraction]} \quad \frac{\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow \varphi' \quad X \cap \text{vars}(\varphi') = \emptyset}{\mathcal{S} \vdash_{\mathcal{C}} ((\exists X) \varphi \Rightarrow \varphi')} \\
\text{[Reflexivity]} \quad \frac{\cdot}{\mathcal{S} \vdash \varphi \Rightarrow \varphi} \\
\text{[Consequence]} \quad \frac{\models_{\text{FOL}} \varphi_1 \rightarrow \varphi'_1 \quad \mathcal{S} \vdash_{\mathcal{C}} \varphi'_1 \Rightarrow \varphi'_2 \quad \models_{\text{FOL}} \varphi'_2 \rightarrow \varphi_2}{\mathcal{S} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow \varphi_2} \\
\text{[CaseAnalysis]} \quad \frac{\mathcal{S} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow \varphi \quad \mathcal{S} \vdash_{\mathcal{C}} \varphi_2 \Rightarrow \varphi}{\mathcal{S} \vdash_{\mathcal{C}} (\varphi_1 \vee \varphi_2) \Rightarrow \varphi} \\
\text{[Transitivity]} \quad \frac{\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow \varphi'' \quad (\mathcal{S} \cup \mathcal{C}) \vdash \varphi'' \Rightarrow \varphi'}{\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow \varphi'} \\
\text{[Circularity]} \quad \frac{\mathcal{S} \vdash_{\mathcal{C} \cup \{\varphi \Rightarrow \varphi'\}} \varphi \Rightarrow \varphi'}{\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow \varphi'}
\end{array}$$

Figure 2.1: The language-independent proof system of Reachability Logic.

etc. In particular, it allows us to make a claim of circular behaviour at any moment during a proof derivation; the claim holds if it can be proved using itself. This type of reasoning is sound because the circularity claim will be allowed only after at least one application of a (trusted) rule from \mathcal{S} . In addition to the proof rules from Figure 2.1, the proof system of RL also contains the following derived rules (proved in [93]):

- *Substitution*: $\mathcal{S} \vdash_G \theta(\varphi) \Rightarrow \theta(\varphi')$, if $\theta : \text{Var} \rightarrow T_{\Sigma}(\text{Var})$ and $\mathcal{S} \vdash_G \varphi \Rightarrow \varphi'$;
- *Logical Framing*: $\mathcal{S} \vdash_G (\varphi \wedge \phi) \Rightarrow (\varphi' \wedge \phi)$, if ϕ is a patternless FOL formula and $\mathcal{S} \vdash_G \varphi \Rightarrow \varphi'$;
- *Set Circularity*: if $\mathcal{S} \vdash_G \varphi \Rightarrow \varphi'$ for each $\varphi \Rightarrow \varphi' \in G$ and G is finite then $\mathcal{S} \vdash \varphi \Rightarrow \varphi'$ for each $\varphi \Rightarrow \varphi' \in G$;
- *Implication*: if $\models_{\text{FOL}} \varphi \rightarrow \varphi'$ then $\mathcal{S} \vdash \varphi \Rightarrow \varphi'$;
- *Monotony*: if $\mathcal{S} \subseteq \mathcal{S}'$ then $\mathcal{S} \vdash \varphi \Rightarrow \varphi'$ implies $\mathcal{S}' \vdash \varphi \Rightarrow \varphi'$.

In this dissertation we will make use of the following theorem, which states that the proof system shown in Figure 2.1 is sound:

Theorem 2.4.1 (Soundness) *If \mathcal{S} is weakly well-defined and $\mathcal{S} \vdash \varphi \Rightarrow \varphi'$ then $\mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$.*

The proof of Theorem 2.4.1 can be found in [93].

2.5 The \mathbb{K} framework

In this section we present the basics of \mathbb{K} , a framework for defining formal semantics of programming languages. \mathbb{K} has been used in this dissertation as a platform on top of which we have implemented our symbolic execution framework. This brought us a few advantages in terms of implementation due to the fact that our symbolic execution tool inherited all the features of \mathbb{K} . Moreover, we took advantage of the existing language definitions by testing our symbolic execution on the defined languages. We describe the main features of the \mathbb{K} framework by means of an existing language definition called CinK, which is presented entirely in [70]. CinK is an overly simplified core of the C++ language, which includes the basic integer and boolean expressions, local and global variables, side effects, lvalues and rvalues, functions, call by value, (unidimensional) arrays, pointers, and the basic imperative statements (conditionals, loops, etc.). In [71], the language has been also extended with a few concurrency constructs (e.g. thread creation, lock-based synchronisation, thread-join), but here we stick to the “non-concurrent” version of CinK, since we only highlight the basic ingredients of a \mathbb{K} definition.

Introduced by Grigore Roşu in 2003 [91] for teaching a programming languages class, and continuously refined and developed ever since, \mathbb{K} [97, 108] is a framework for defining programming language based on rewriting which combines the strengths of existing frameworks (expressiveness, modularity, concurrency, and simplicity) while avoiding their weaknesses. If one ignores its concurrent semantics, \mathbb{K} can be seen as a notation within rewriting logic [79], the same as most other semantic frameworks, such as natural (or big-step) semantics, (small-step) SOS, Modular SOS, reduction semantics with evaluation contexts, and so on [111]. However, unlike these other semantic frameworks enumerated above, \mathbb{K} cannot be easily

captured step-for-step in rewriting logic, due to its enhanced concurrency which is best described in terms of ideas from graph rewriting [110].

\mathbb{K} has been successfully used to define a large number of programming languages, from simple languages used for teaching (e.g. CinK, SIMPLE [102], KOOL [53], KernelC [113]) to real-life programming languages, such as, C [35], Java [17], Scheme [85], Python [51], PHP [39]. Some other \mathbb{K} definitions were developed to capture various aspects of features of OCL [6], RISC assembly [14, 11], Verilog [76], Javascript [80], Haskell'98 [66], X10 [44], and LLVM [68]. The framework was also used for designing type-checkers [36], model-checking with predicate abstractions [9, 13, 8], path directed symbolic execution [10], program equivalence [69], computing worst-case time analysis [11, 12], and runtime verification [96]. The theoretical foundations of the \mathbb{K} framework can be found in [97], while the \mathbb{K} compiler and runner are described in [114] and [67], respectively.

The main ingredients of a \mathbb{K} semantic definition of a language are *computations*, *configurations* and *rules*. Computations are special tasks which can be seen as *units* of program execution. Configurations are nested structures which hold both syntactic and semantics information about programs. The \mathbb{K} rules are a particular type of reachability rules, where modularity is achieved by specifying only those parts that change in a configuration. All these ingredients are illustrated on the \mathbb{K} definition of CinK.

The *syntax* of a language is given using a BNF-style grammar. Here we only exhibit a fragment of the \mathbb{K} definition of CinK (Figure 2.2). The grammar productions are usually annotated with \mathbb{K} -specific attributes. Some of them are related to the evaluation order of the arguments of a specific syntactical construct. For instance, a major feature of the C++ expressions is the “sequenced before” relation [57], which defines a partial order over the evaluation of subexpressions. This can be easily expressed in \mathbb{K} using the *strict* attribute to specify an evaluation order for an operation’s operands. If the operator is annotated with the *strict* attribute then its operands will be evaluated in a nondeterministic order. For instance, all the binary operations are strict. Hence, they may induce non-determinism in programs because of possible side-effects in their arguments.

Another feature is given by the classification of expressions into *rval-*

<i>Exp</i>	::= <i>Id</i> <i>Int</i> <i>Bool</i> <i>String</i>	
	<i>Exp</i> (<i>Exp</i>)	[<i>strict</i> (1(<i>context</i> (<i>rvalue</i>))), <i>funcall</i>]
	++ <i>Exp</i>	[<i>strict</i> , <i>prefinc</i>]
	-- <i>Exp</i>	[<i>strict</i> , <i>prefdec</i>]
	<i>Exp</i> / <i>Exp</i>	[<i>strict</i> (<i>all</i> (<i>context</i> (<i>rvalue</i>))), <i>divide</i>]
	<i>Exp</i> + <i>Exp</i>	[<i>strict</i> (<i>all</i> (<i>context</i> (<i>rvalue</i>))), <i>plus</i>]
	<i>Exp</i> > <i>Exp</i>	[<i>strict</i> (<i>all</i> (<i>context</i> (<i>rvalue</i>)))]
	<i>Exp</i> && <i>Exp</i>	[<i>strict</i> , <i>and</i>]
<i>Stmt</i>	::= <i>Exps</i> ;	[<i>strict</i>]
	{ <i>Stmts</i> }	
	while (<i>Exp</i>) <i>Stmt</i>	
	return <i>Exp</i> ;	[<i>strict</i> (<i>all</i> (<i>context</i> (<i>rvalue</i>)))]
	<i>Decl</i> (<i>Decls</i>) { <i>Stmts</i> }	[<i>fundec</i>]
	if (<i>Exp</i>) <i>Stmt</i> else <i>Stmt</i>	[<i>strict</i> (1(<i>context</i> (<i>rvalue</i>)))]
<i>Decl</i>	::= <i>Type Exp</i>	
<i>Type</i>	::= int bool void	

Figure 2.2: Fragment of CinK syntax

ues and *lvalues*. The arguments of binary operations are evaluated as rvalues and their results are also rvalues, while, e.g., both the argument of the prefix-increment operation and its result are lvalues. The *strict* attribute for such operations has a sub-attribute *context* for wrapping any subexpression that must be evaluated as an rvalue. Other attributes (*funcall*, *divide*, *plus*, *minus*, ...) are names associated to each syntactic production, which can be used to refer to them.

A \mathbb{K} *configuration* is used to store program states. It consists in a nested structure of cells, which typically includes the program to be executed, input and output streams, values for program variables, and other additional information. The configuration of CinK (Figure 2.3) includes the $\langle \cdot \rangle_k$ cell containing the code that remains to be executed, which is represented as a list of *computation tasks* $C_1 \curvearrowright C_2 \curvearrowright \dots$ to be executed in the given order. Computation tasks are typically statements and expression evaluations. The memory is modelled using two cells $\langle \cdot \rangle_{\text{env}}$ (which

$$\langle \langle \$PGM \rangle_k \langle \cdot \rangle_{\text{env}} \langle \cdot \rangle_{\text{store}} \langle \cdot \rangle_{\text{stack}} \langle \cdot \rangle_{\text{return}} \langle \cdot \rangle_{\text{in}} \langle \cdot \rangle_{\text{out}} \rangle_{\text{cfg}}$$

Figure 2.3: CinK configuration

holds a map from variables to memory locations) and $\langle \rangle_{\text{state}}$ (which holds a map from memory locations to values). Locations are used as values for pointer types. The configuration also includes a cell for the function call stack $\langle \rangle_{\text{stack}}$ and another one $\langle \rangle_{\text{return}}$ for the return values of functions. The $\langle \rangle_{\text{in}}$ and $\langle \rangle_{\text{out}}$ cells hold the input and (respectively) the output of the program. Both cells are connected to the standard input/output stream. When the configuration is initialised at runtime, a CinK program is loaded in the $\langle \rangle_{\text{k}}$ cell, and all the other cells remain empty.

A \mathbb{K} rule is a topmost rewrite rule specifying transitions between configurations. Since usually only a small part of the configuration is changed by a rule, a *configuration abstraction* mechanism is used, allowing one to only specify the parts transformed by the rule. For instance, the (abstract) rule for addition, shown in Figure 2.4, represents the (concrete) rule

$$\begin{aligned} & \langle \langle I_1 + I_2 \curvearrowright C \rangle_{\text{k}} \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle T \rangle_{\text{stack}} \langle V \rangle_{\text{return}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \rangle_{\text{cfg}} \\ & \Rightarrow \\ & \langle \langle I_1 +_{\text{Int}} I_2 \curvearrowright C \rangle_{\text{k}} \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle T \rangle_{\text{stack}} \langle V \rangle_{\text{return}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \rangle_{\text{cfg}} \end{aligned}$$

where $+_{\text{Int}}$ is the mathematical operation for addition, I_1, I_2, V are integers, E and S are maps, and T, I , and O are lists. The three dots “...” in a cell are meant to help the configuration abstraction mechanism, by specifying that there is some other content in the cell, which is not affected by the current rule.

The rule for division has a side condition which restricts its application, i.e. it applies only when the denominator is not zero. The conditional statement **if** has two corresponding rules, one for each possible evaluation of the condition expression. The rule for the **while** loop performs an unrolling into an **if** statement. The increment and update rules have side effects in the $\langle \rangle_{\text{store}}$ cell, modifying the value stored at a specific address. Reading a value from the memory is specified by the lookup rule, which matches a value in the $\langle \rangle_{\text{store}}$ and places it in the $\langle \rangle_{\text{k}}$ cell. The auxiliary construct `$lookup` is used when a program variable is evaluated as an rvalue. The semantics of a variable declaration consists in adding new entries in the $\langle \rangle_{\text{env}}$ and $\langle \rangle_{\text{store}}$ cells, where the variable is mapped to a *fresh* memory location L , which is initialised with zero. A function declaration is handled in the same way as a variable declaration, except that the value stored in the fresh location is the lambda abstraction of the function. The rule defining the evaluation of a function is a bit more complex: it evaluates

$I_1 / I_2 \wedge I_2 \neq_{Int} 0 \Rightarrow I_1 /_{Int} I_2$	$[division]$
$I_1 + I_2 \Rightarrow I_1 +_{Int} I_2$	$[plus]$
$I_1 > I_2 \Rightarrow I_1 >_{Int} I_2$	$[greater]$
$if(true) St \text{ else } _ \Rightarrow St$	$[if-true]$
$if(false) _ \text{ else } St \Rightarrow St$	$[if-false]$
$while(B) St \Rightarrow if(B)\{ St \text{ while}(B) St \text{ else } \{\}\}$	$[while]$
$V ; \Rightarrow \cdot$	$[stmt-exp]$
$\langle ++lvalue(L) \Rightarrow lvalue(L) \dots \rangle_k \langle \dots L \mapsto (V \Rightarrow V +_{Int} 1) \dots \rangle_{store}$	$[inc]$
$\langle --lvalue(L) \Rightarrow lvalue(L) \dots \rangle_k \langle \dots L \mapsto (V \Rightarrow V -_{Int} 1) \dots \rangle_{store}$	$[dec]$
$\langle \langle lvalue(L) = V \Rightarrow V \dots \rangle_k \langle \dots L \mapsto _ \Rightarrow V \dots \rangle_{store} \dots \rangle_{cfg}$	$[update]$
$\langle \langle \$lookup(L) \Rightarrow V \dots \rangle_k \langle \dots L \mapsto V \dots \rangle_{store} \dots \rangle_{cfg}$	$[lookup]$
$\{ Sts \} \Rightarrow Sts$	$[block]$
$\langle TX ; \Rightarrow \cdot \dots \rangle_k \langle \dots Env \Rightarrow Env[L/X] \dots \rangle_{env}$ $\langle \dots \cdot \Rightarrow L \mapsto 0 \dots \rangle_{store} \wedge fresh(L)$	$[vardecl]$
$\langle D : Decl(X : Decls)\{ S : Stmts \} \Rightarrow \cdot \dots \rangle_k$ $\langle \dots \cdot \Rightarrow getName(D) \mapsto L \dots \rangle_{env}$ $\langle \dots \cdot \Rightarrow L \mapsto \lambda X.S \dots \rangle_{store} \wedge fresh(L)$	$[fundecl]$
$\langle \lambda X.S(E) \Rightarrow eval \text{ app}(E) \text{ to } \cdot_{List} \text{ fol } X ; \curvearrowright seqpoint \curvearrowright \lambda X.S(\square) \dots \rangle_k$	$[funcall]$
$\langle eval \text{ app}(E_{Exps}) \text{ to } V \text{ fol } X ; \curvearrowright seqpoint \Rightarrow$ $seqpoint \curvearrowright eval \text{ app}(E_{Exps}) \text{ to } V \text{ fol } X ; \dots \rangle_k$	
$\langle eval \text{ app}(E_{Exps}) \text{ to } V \text{ fol } Y ; \curvearrowright \lambda X.S(\square) \curvearrowright K \rangle_k \langle E \rangle_{stack} \Rightarrow$ $bind V \text{ to } X ; \curvearrowright S \curvearrowright return \text{ noVal} ; \dots \rangle_k \langle ([Env], K)E \rangle_{stack}$	
$\langle bind(V, Vs) \text{ to } (T X, Xl) ; \rangle_k \langle Env \rangle_{env} \langle M \rangle_{store} \Rightarrow$ $\langle bind(Vs) \text{ to } (Xl) ; \rangle_k \langle Env[L/X] \rangle_{env} \langle M \ L \mapsto V \rangle_{store}$	
$\langle bind(E_{Exps}) \text{ to } (\cdot Decls) ; \rangle_k \Rightarrow \langle \cdot K \rangle_k$	
$\langle return(V) ; \curvearrowright \cdot \rangle_k \langle \cdot \rangle_{env} \langle ([Env], K) \rangle_{stack} \langle \cdot \rangle_{return} \Rightarrow$ $\langle V \curvearrowright K \rangle_k \langle Env \rangle_{env} \langle \cdot List \rangle_{stack} \langle V \rangle_{return}$	$[return]$

Figure 2.4: Subset of rules from the \mathbb{K} semantics of CinK

the parameters, then binds values to the formal parameters, and then it executes the body while saving the calling context. We use \square as a special variable, destined to receive the result of an evaluation.

In addition to these rules (written by the \mathbb{K} user), the \mathbb{K} framework automatically generates so-called *heating* and *cooling* rules, which are induced by *strict* attributes. We show the case of division, which is strict in both arguments:

$$\begin{array}{ll} A_1 / A_2 \Rightarrow rvalue(A_1) \curvearrowright \square / A_2 & rvalue(I_1) \curvearrowright \square / A_2 \Rightarrow I_1 / A_2 \\ A_1 / A_2 \Rightarrow rvalue(A_2) \curvearrowright A_1 / \square & rvalue(I_2) \curvearrowright A_1 / \square \Rightarrow A_1 / I_2 \end{array}$$

The \mathbb{K} rules are classified in two categories: *structural* rules and *computational* rules. Intuitively, only the computational rules yield transitions in the transition system associated to a program, while the structural rules are meant to only re-arrange the configuration so that the computational rules can match.

In terms of implementation, \mathbb{K} consists in a set of command line tools which includes a definition compiler and a program runner. The \mathbb{K} compiler generates a rewrite theory, which is then used to run real programs. At runtime, the runner first loads the given program in the initial configuration and then applies repeatedly the \mathbb{K} rules to the configuration, until no rule can be applied. The result consists in one or more *final* configurations. \mathbb{K} uses an external rewrite engine (Maude), which allows I/O interactions [7], and is integrated with the Z3 SMT solver [34]. Thus, \mathbb{K} can be used as a tool for both executing the formal semantics of a programming language and for analysing programs written in that language.

Chapter 3

Language Independent Symbolic Execution Framework

Symbolic execution is a well-known program-analysis technique used in testing, verification, debugging, and other applications. There are many tools implementing symbolic execution engines but most of them are language dependent (e.g. Java PathFinder [82], PEX [38]). Any existent language independent approaches are based on translations to different languages (e.g. WOODPECKER [31], KLEE [23]), where symbolic execution is implemented at the compiler level. However, this approach has a few downsides mostly related to the fact that the translation of a language into another language is not guaranteed to be sound.

The main goal of this dissertation is to design a language independent framework which is based on the language's formal semantics. In this chapter, we introduce formal definitions for programming languages and symbolic execution. We prove that the symbolic execution thus defined has the following properties, which ensures a natural relation between symbolic and concrete program execution:

- *Coverage*: to every concrete execution there corresponds a feasible symbolic one;
- *Precision*: to every feasible symbolic execution there corresponds a concrete one;

Two executions are said to be corresponding if they take the same path in the control flow graph of the program, and a symbolic execution is feasible if the path conditions along it are satisfiable. Or, stated in terms of simulations: the feasible symbolic executions and the concrete executions of any given program mutually simulate each other. These properties are very important when performing analyses on symbolic programs, since the results of those analyses can be soundly transferred to concrete instances of the symbolic programs in question.

This chapter is organised as follows: we define programming languages in terms of algebraic specifications and RL formulas in Section 3.1. Then, we present (what we mean by) unification (Section 3.2) and we prove a technical lemma which states that unification can be resolved by matching under reasonable conditions. Unification is then used to define symbolic execution and the lemma is used to prove the *Coverage* and *Precision* properties in Section 3.3. Finally, in Section 3.4, we show how our symbolic execution framework can be implemented by standard rewriting.

3.1 Language Definitions

A wide range of programming languages are used nowadays for programming billions of hardware devices. Depending on their level of abstraction, there are low level programming languages, whose instructions are very close to the language circuits (e.g. Assembler, Verilog, VHDL), or high level imperative, functional or object-oriented programming languages (e.g. C, OCaml, Java, etc.). Whatever a programming language is, it requires a *specification* which clearly states the meaning of the programs written in that language. This specification includes the syntax of the language, that is, the constructs of which programs are made, and the semantics of the language, which describes the behaviour of each syntactical construct. Over the years, researchers reached a consensus on the idea that language specifications should be given using a formal specification language. However, formal specification languages have been themselves designed to be appropriate for a single analysis area: structural operational semantics for dynamic semantics, reduction semantics with evaluation contexts for proving type soundness, or axiomatic semantics for program verification. Moreover, formal descriptions of languages are

confined to the power of expression of their specification language.

In this dissertation we choose a formal specification language which is intended to keep the balance between simplicity and power of expressivity and analysis. We propose a general notion of language definition based on algebraic specifications and Reachability Logic. The syntax of a programming language is given using algebraic signatures, while its semantics is given using RL formulas. The intent of having this generic notion of language definition is driven by our goal of defining a generic framework for symbolic execution, which is parametric in the language definition.

We now define languages in terms of ML signatures and RL formulas.

Definition 3.1.1 (Language Definition) *A language definition is a tuple $\mathcal{L} = ((\Sigma, \Pi, Cfg), \mathcal{T}, \mathcal{S})$, where:*

- (Σ, Π, Cfg) is a ML signature,
- \mathcal{T} is a model of (Σ, Π, Cfg) , and
- \mathcal{S} is a set of RL formulas.

In the rest of the dissertation, we consider that language definitions $\mathcal{L} = ((\Sigma, \Pi, Cfg), \mathcal{T}, \mathcal{S})$ meet the assumptions presented in the rest of this section.

Assumption 3.1.1 *The first-order signature (Σ, Π) includes a subsignature $(\Sigma^{\mathfrak{d}}, \Pi^{\mathfrak{d}})$ consisting of all data sorts together with their operations and predicates. We implicitly consider a subset $S^{\mathfrak{d}} \subset S$ of data sorts and that $\Sigma_{w,s} = \Sigma_{w,s}^{\mathfrak{d}}$ for all $(w, s) \in S^{\mathfrak{d}*} \times S^{\mathfrak{d}}$. The signature $(\Sigma^{\mathfrak{d}}, \Pi^{\mathfrak{d}})$ depends on the language \mathcal{L} , which has specific data sorts (e.g. integers, booleans, arrays, ...). We always assume that the sort Cfg is not a data sort.*

The distinction between data and non-data sorts from Assumption 3.1.1 is essential, since in our approach we only allow the data sorts from $\Sigma^{\mathfrak{d}}$ to be symbolic. In Example 3.1.1 we show the signatures Σ and $\Sigma^{\mathfrak{d}}$, corresponding to the fragment of the syntax of CinK (Figure 2.2).

Example 3.1.1 *Consider the syntax fragment of CinK, shown in Figure 2.2. The set S includes nonterminals Int , $Bool$, $String$, Id , Exp , $Exps$,*

Example 3.1.3 For *CinK*, the model \mathcal{D} interprets *Int* as the set of integers, the operations like $+$ as the corresponding usual operation on integers, *Bool* as the set of Boolean values, the operations like $\&\&$ as the usual Boolean operations, the sort *Map* as the set of mappings $X \mapsto I$, where X ranges over identifiers (*Id*) or integers (*Int*), and I over integers (*Int*).

The model \mathcal{T} interprets data terms in the same way as \mathcal{D} (e.g. $\mathcal{T}_1 = \mathcal{D}_1$) and non-data terms as ground terms, where data subterms are replaced by their interpretations in \mathcal{D} . For instance, `if 1 > 0 then { skip; } else { skip; }` is interpreted as $\mathcal{T}_{\text{if } \mathcal{D}_{\text{true}} \text{ then } \{ \text{skip}; \} \text{ else } \{ \text{skip}; \}}$, since $\mathcal{D}_{1>0} = \mathcal{D}_{\text{true}}$.

Example 3.1.4 Let *size* be a function symbol in $\Sigma_{\text{Cfg,Int}}$. Then the interpretation of *size* in \mathcal{T} is given by \mathcal{D} , i.e. $\mathcal{T}_{\text{size}} = \mathcal{D}_{\text{size}}$.

Assumption 3.1.3 The last component of a language definition is a set \mathcal{S} of RL formulas of the form $\pi \wedge \phi \Rightarrow \pi' \wedge \phi'$, where $\text{vars}(\phi) \subseteq \text{vars}(\pi)$. We tacitly assume that the variables in $\text{vars}(\pi' \wedge \phi')$ and not in $\text{vars}(\pi \wedge \phi)$ are existentially quantified.

Example 3.1.5 The \mathbb{K} rules shown in Figure 2.4 are included in the set \mathcal{S} of RL formulas which define the semantics of *CinK*. Note that \mathbb{K} rules are in fact particular cases of RL formulas.

3.2 Unification

Our approach uses unification for formally defining symbolic execution. In this section we define what we mean by unification and we prove a technical lemma which states that unification can be implemented using matching, under some assumptions. This result is essential for proving the *Coverage* and *Precision* properties, also formalised using unification.

We assume a given language definition $\mathcal{L} = ((\Sigma, \Pi, \text{Cfg}), \mathcal{T}, \mathcal{S})$ as in Section 3.1, and an infinite S-indexed set of variables Var .

Definition 3.2.1 (Unifiers) A symbolic unifier of two terms t_1, t_2 , with $\text{vars}(t_1) \cap \text{vars}(t_2) = \emptyset$, is any substitution $\sigma : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow T_\Sigma(Z)$ for some set Z of variables such that $\sigma(t_1) = \sigma(t_2)$. A concrete unifier of terms t_1, t_2 is any valuation $\rho : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow \mathcal{T}$ such that $\rho(t_1) =$

$\rho(t_2)$. A *symbolic unifier* σ of two terms t_1, t_2 is a most general unifier of t_1, t_2 with respect to concrete unification whenever, for all concrete unifiers ρ of t_1 and t_2 , there is a concrete unifier $\eta : Z \rightarrow \mathcal{T}$ such that $\eta \circ \sigma = \rho$.

We often call a symbolic unifier satisfying the above a *most general unifier*, even though the standard notion of most general unifier in algebraic specifications/rewriting is slightly different. We say that terms t_1, t_2 are symbolically (resp. concretely) unifiable if they have a symbolic (resp. concrete) unifier.

Example 3.2.1 *The term $t_1 \triangleq \text{if_then_else_}(B, S_1, S_2)$ and the term $t_2 \triangleq \text{if_then_else_}(B', S_1, S'_2)$ are symbolically unifiable, by the substitution $B \mapsto B', S_2 \mapsto S'_2$, extended to the identity for the other variables occurring in the terms. Since $\text{if_then_else_}(B, S_1, S_2)$ is a non data function symbols in Σ (statements are not data), the two terms are also concretely unifiable, e.g., by any valuation that maps B to true and S_1, S_2 to $\{\}$ (empty block).*

In the rest of the section we prove that unification of two given terms, say t_1 and t_2 , can be achieved by matching. However, this is possible considering some assumptions over t_1 and t_2 . One of these assumptions is that t_1 is *linear*, that is, every variable from $\text{vars}(t_1)$ occurs only once in t_1 . Another assumption is that every subterm of sort data of t_1 is a variable and all the elements of $\text{vars}(t_2)$ have data sorts. The terms we deal with are ML patterns of the form $\pi \wedge \phi$. Every pattern of this form can be transformed such that it is linear and all its data subterms are variables. The transformation consists in replacing the variables and the data subterms with fresh variables and adding equalities between the fresh variables and the terms they replaced to ϕ . Moreover, the transformation does not change the transition system generated by the language (which can be proved by applying Lemma 2.4.2 repeatedly to all subterms of sort data and duplicate variables).

The following lemma shows that for two terms t_1 and t_2 , which are concretely unifiable and satisfy the assumptions above, there exists the most general unifier of t_1 and t_2 . Moreover, this unifier is obtained via a substitution which represents the matching of t_2 against t_1 . In other words, the lemma shows that we can implement symbolic execution (defined in terms of unification) using rewriting.

Lemma 3.2.1 (Unification by Matching) *If t_1 and t_2 are terms such that*

$$(h_1) \text{ vars}(t_1) \cap \text{vars}(t_2) = \emptyset,$$

(h_2) t_1 is linear, has a non-data sort, and all its data subterms are variables,

(h_3) all the elements of $\text{vars}(t_2)$ have data sorts, and

(h_4) t_1, t_2 are concretely unifiable,

then there exists a substitution $\sigma : \text{vars}(t_1) \rightarrow T_\Sigma(\text{vars}(t_2))$ such that $\sigma(t_1) = t_2$ and $\sigma_{t_2}^{t_1} \triangleq \sigma \uplus 1_{\text{vars}(t_2)}$ is a most-general unifier of t_1, t_2 .

Proof By assumptions (h_2), (h_3), and (h_4), the term t_2 has the following properties:

(c_1) t_2 has a non-data sort, and

(c_2) t_2 cannot be a variable.

Indeed, because t_1 and t_2 are concretely unifiable (cf. (h_4)), so there is $\rho : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow \mathcal{T}$ of t_1 and t_2 , such that $\rho(t_1) = \rho(t_2)$. Since the carrier sets in \mathcal{T} are disjoint then t_1 and t_2 have the same sort, and thus, t_2 has a non-data sort (cf. (h_2)). The term t_2 cannot be a variable because t_2 should have simultaneously a non-data sort (cf. (c_1)), and a data sort, since $t_2 \in \text{vars}(t_2)$ (cf. (h_3)). In the rest of the proof we assume (c_1) and (c_2).

We proceed by induction on the structure of t_1 . In the base case t_1 is either a variable or a constant:

- t_1 is a variable. In this case we have $\text{vars}(t_1) = \{t_1\}$, where t_1 is a variable of non-data sort (cf. (h_2)). We consider $\sigma : \text{vars}(t_1) \rightarrow T_\Sigma(\text{vars}(t_2))$ such that $\sigma(t_1) = t_2$, and $\sigma_{t_2}^{t_1} : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow T_\Sigma(\text{vars}(t_2))$ where $\sigma_{t_2}^{t_1} \triangleq \sigma \uplus 1_{\text{vars}(t_2)}$. Since $\sigma_{t_2}^{t_1}(t_1) = \sigma(t_1) = t_2 = 1_{\text{vars}(t_2)}(t_2) = \sigma_{t_2}^{t_1}(t_2)$ then $\sigma_{t_2}^{t_1}$ is a symbolic unifier of t_1, t_2 . We have to show that $\sigma_{t_2}^{t_1}$ is also the most general unifier of t_1, t_2 . Let us consider any concrete unifier, say, $\rho : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow \mathcal{T}$ of t_1 and t_2 . Then, $\rho(t_1) = \rho(t_2)$ (cf. Definition 3.2.1) and

$\rho(t_2) = \rho(\sigma_{t_2}^{t_1}(t_1))$ (where $\sigma_{t_2}^{t_1}(t_1) = t_2$ as shown above). Thus, $\rho(t_1) = \rho(t_2) = \rho(\sigma_{t_2}^{t_1}(t_1))$ and for all variables $x \in \text{vars}(t_2)$ we have $\rho(x) = \rho(1_{\text{vars}(t_2)}(x)) = \rho(\sigma_{t_2}^{t_1}(x))$. Therefore, for all $x \in \text{vars}(t_1) \cup \text{vars}(t_2) (= \{t_1\} \cup \text{vars}(t_2))$, $\rho(\sigma_{t_2}^{t_1}(x)) = \rho(x)$, which proves the fact that $\sigma_{t_2}^{t_1}$ is a most general unifier by taking $\eta = \rho$ in Definition 3.2.1.

- t_1 is a constant symbol. Since t_2 cannot be a variable (cf. (c_2)) then t_2 is either a constant of non-data sort, where $t_1 = t_2$, or t_2 is of the form $g(t_2^1, \dots, t_2^m)$, where $m > 0$.

We consider the case when t_2 is a constant. We choose σ the unique function $\emptyset \rightarrow T_\Sigma$ extended to terms and we have $\sigma(t_1) = t_1 = t_2$. Moreover, $\sigma_{t_2}^{t_1} \triangleq \sigma \uplus 1_{\text{vars}(t_2)} = \sigma$ is a symbolic unifier for t_1 and t_2 , and it is also the most general unifier if we choose η the unique function $\emptyset \rightarrow \mathcal{T}$ in Definition 3.2.1.

Now, we consider $t_2 = g(t_2^1, \dots, t_2^m)$, where $m > 0$. From the fact that ρ is a concrete unifier, we have $\rho(t_1) = \rho(t_2)$. Thus, $\rho(t_1) = \mathcal{T}_f =_{\mathcal{T}} \mathcal{T}_g(t_2^1, \dots, t_2^m) = \rho(g(t_2^1, \dots, t_2^m))$, where $=_{\mathcal{T}}$ is the equality in \mathcal{T} . Since, t_1 and t_2 have non-data sorts (cf. (h_2) and (c_1)) and \mathcal{T} is a \mathcal{D} -grounded model, then according to Definition 3.1.2 we have that $f = g$, $m = 0$. This is a contradiction, since we assumed $m > 0$. Therefore, this case is impossible.

For the inductive step, let $t_1 = f(t_1^1, \dots, t_1^n)$ with $f \in \Sigma \setminus \Sigma^0$, $n \geq 1$, and $t_1^1, \dots, t_1^n \in T_\Sigma(\text{Var})$. There are two subcases regarding t_2 :

- t_2 is a constant. This is impossible, since t_1 is not a constant. The proof is similar with the one in the base case, where we considered t_1 as being constant and t_2 not a constant.
- $t_2 = g(t_2^1, \dots, t_2^m)$ with $g \in \Sigma$, $m \geq 1$, and $t_2^1, \dots, t_2^m \in T_\Sigma(\text{Var})$. Since ρ is a concrete unifier of t_1, t_2 we have $\rho(t_1) = \rho(f(t_1^1, \dots, t_1^n)) = \mathcal{T}_f(\rho(t_1^1), \dots, \rho(t_1^n)) =_{\mathcal{T}} \mathcal{T}_g(\rho(t_2^1), \dots, \rho(t_2^m)) = \rho(t_2)$, where we emphasize by subscripting the equality symbol with \mathcal{T} that the equality is that of the model \mathcal{T} .

Since \mathcal{T} is a \mathcal{D} -grounded model we have $f = g$, $m = n$, and $\rho(t_1^i) = \rho(t_2^i)$ for $i \in \{1, \dots, n\}$ (cf. Definition 3.1.2). The respective

subterms t_1^i and t_2^i of t_1 and t_2 satisfy the hypotheses of our lemma, except maybe for the fact that t_1^i may have a data sort. There are again two cases:

- if for some $i \in \{1, \dots, n\}$, t_1^i has a data sort then by (h_2) t_1^i is a variable, with $\text{vars}(t_1^i) = \{t_1^i\}$, and we let $\sigma^i : \text{vars}(t_1^i) \rightarrow T_\Sigma(\text{vars}(t_2^i))$ such that $\sigma^i(t_1^i) = t_2^i$. Then, $\sigma_{t_2^i}^{t_1^i} \triangleq \sigma^i \uplus 1_{\text{vars}(t_2^i)}$ is a most-general unifier of t_1^i and t_2^i , which is proved like in the base case;
- otherwise, t_1^i and t_2^i satisfy all the the hypotheses of our lemma. We can then use the induction hypothesis and obtain substitutions $\sigma^i : \text{vars}(t_1^i) \rightarrow T_\Sigma(\text{vars}(t_2^i))$ such that $\sigma^i(t_1^i) = t_2^i$ for all $i \in \{1, \dots, n\}$, and the corresponding most-general-unifiers $\sigma_{t_2^i}^{t_1^i}$ for t_1^i and t_2^i , of the form $\sigma_{t_2^i}^{t_1^i} = \sigma^i \uplus 1_{\text{vars}(t_2^i)}$.

Let $\sigma \triangleq \biguplus_{i=1}^n \sigma^i : \text{vars}(t_1) \rightarrow T_\Sigma(\text{vars}(t_2))$, which is a well-defined substitution because t_1 is linear. If $\sigma_{t_2}^{t_1} \triangleq \sigma \uplus 1_{\text{vars}(t_2)}$ then $\sigma_{t_2}^{t_1}$ is a symbolic unifier of t_1, t_2 :

$$\begin{aligned}
\sigma_{t_2}^{t_1}(t_1) &= \sigma_{t_2}^{t_1}(f(t_1^1, \dots, t_1^n)) \\
&= f(\sigma_{t_2}^{t_1}(t_1^1), \dots, \sigma_{t_2}^{t_1}(t_1^n)) \\
&= f(\sigma(t_1^1), \dots, \sigma(t_1^n)) \\
&= f(\sigma^1(t_1^1), \dots, \sigma^n(t_1^n)) \\
&= f(t_2^1, \dots, t_2^n) \\
&= f(1_{\text{vars}(t_2)}(t_2^1), \dots, 1_{\text{vars}(t_2)}(t_2^n)) \\
&= f(\sigma_{t_2}^{t_1}(t_2^1), \dots, \sigma_{t_2}^{t_1}(t_2^n)) \\
&= \sigma_{t_2}^{t_1}(f(t_2^1, \dots, t_2^n)) \\
&= \sigma_{t_2}^{t_1}(t_2)
\end{aligned}$$

Next, we prove that $\sigma_{t_2}^{t_1}$ is the most general unifier. We extend $\sigma_{t_2}^{t_1}$ to $\mu_{t_2}^{t_1} : \text{vars}(t_1) \cup \text{vars}(t_2) \rightarrow \text{vars}(t_2)$ where $\mu_{t_2}^{t_1} \triangleq \sigma^i \uplus 1_{\text{vars}(t_2)}$ (which is well-defined cf. (h_1) , (h_2)). Then:

$$\begin{aligned}
\sigma_{t_2}^{t_1} &= \sigma \uplus 1_{\text{vars}(t_2)} \\
&= (\uplus_{i=1}^n \sigma^i) \uplus 1_{\text{vars}(t_2)} \\
&= \uplus_{i=1}^n (\sigma^i \uplus 1_{\text{vars}(t_2)}) \\
&= \uplus_{i=1}^n \mu_{t_2}^{t_1^i}.
\end{aligned}$$

Since ρ is a concrete unifier of t_1 and t_2 , $\rho(t_1^i) = \rho(t_2^i)$ for $i \in \{1, \dots, n\}$. From the inductive hypothesis we have that each $\sigma_{t_2}^{t_1^i}$ is the most-general-unifier of t_1^i and t_2^i for $i = \{1, \dots, n\}$, so we obtain the existence of valuations $\eta^i : \text{vars}(t_2^i) \rightarrow \mathcal{T}$ such that $\eta^i \circ \sigma_{t_2}^{t_1^i} = \rho|_{\text{vars}(t_1^i) \cup \text{vars}(t_2^i)}$, for $i = 1, \dots, n$. We choose $\eta : \text{vars}(t_2) \rightarrow \mathcal{T}$ such that $\eta(x) = \eta^i(x)$ when $x \in \text{vars}(t_2^i)$. Remember that t_2 may not be linear, so for $i \neq j$ we may have $y \in \text{vars}(t_2^i) \cap \text{vars}(t_2^j)$. Thus, we analyse the cases when $y \in \text{vars}(t_2^i)$ and $y \in \text{vars}(t_2^j)$:

- $\eta(y) = \eta^i(y) = (\eta^i \circ \sigma_{t_2}^{t_1^i})(y) = \rho(y)$, or
- $\eta(y) = \eta^j(y) = (\eta^j \circ \sigma_{t_2}^{t_1^j})(y) = \rho(y)$.

We used the fact that $\sigma_{t_2}^{t_1^i}$ and $\sigma_{t_2}^{t_1^j}$ are defined over $\text{vars}(t_1^i)$ and $\text{vars}(t_1^j)$, respectively. Either way $\eta(y) = \eta^i(y) = \eta^j(y)$, where from we conclude that η is well-defined.

Let $x \in \text{vars}(t_1) \cup \text{vars}(t_2)$. Then there is i such that $x \in \text{vars}(t_1^i)$ (since t_1 linear). For η and $\sigma_{t_2}^{t_1}$ defined as above we have $(\eta \circ \sigma_{t_2}^{t_1})(x) = \eta(\sigma_{t_2}^{t_1}(x)) = \eta(\mu_{t_2}^{t_1^i}(x)) = \eta(\sigma^i(x)) = \eta^i(\sigma^i(x)) = (\eta^i \circ \sigma^i)(x) = \rho|_{\text{vars}(t_1^i) \cup \text{vars}(t_2^i)}(x) = \rho(x)$. Thus, η has the property $\eta \circ \sigma_{t_2}^{t_1} = \rho$, which proves that $\sigma_{t_2}^{t_1}$ is a most general unifier of t_1 and t_2 and concludes the proof.

■

Remark 3.2.1 *The most general unifier $\sigma_{t_2}^{t_1}$ whose existence is stated by Lemma 3.2.1 is also unique, because $\sigma_{t_2}^{t_1}$ is defined to be $\sigma \uplus 1_{\text{vars}(t_2)}$ and σ , which is a (syntactical) match of t_1 on t_2 , is unique when it exists.*

3.3 Symbolic transition relation and properties

In this section we present a symbolic execution approach for languages defined using the language-definition framework presented in Section 3.1. We prove that the transition system generated by symbolic execution forward-simulates the one generated by concrete execution, and that the transition system generated by concrete execution backward-simulates the one generated by symbolic execution (restricted to satisfiable patterns). These properties are the naturally expected ones from a symbolic execution framework. They allow to perform analyses on symbolic programs, and to transfer the results of those analyses to concrete instances of the symbolic programs in question.

In the rest of the section we let $\mathcal{L} \triangleq ((\Sigma, \Pi, Cfg), \mathcal{T}, \mathcal{S})$ be a language definition. Moreover, we assume that $vars(\varphi \Rightarrow \varphi') = vars(\varphi) \cup vars(\varphi')$, and for every rule $\varphi_1 \Rightarrow \varphi_2 \in \mathcal{S}$ with $\varphi_1 \triangleq \pi_1 \wedge \phi_1$, π_1 satisfies the hypotheses regarding t_1 of Lemma 3.2.1, that is, π_1 is linear, has non-data sort, and all its data subterms are variables. Recall that for patterns $\varphi \triangleq \pi \wedge \phi$ we have $vars(\varphi) = vars(\pi) \cup vars(\phi)$ (Section 2.3). For simplicity, from now on we will use $vars(\pi, \phi)$ to denote $vars(\pi) \cup vars(\phi)$, and for any patterns φ and φ' we let $vars(\varphi, \varphi') \triangleq vars(\varphi_1) \cup vars(\varphi_2)$.

Symbolic execution essentially consists of applying the semantical rules over patterns using the most general unifiers whose existence and unicity are given by Lemma 3.2.1 and Remark 3.2.1.

We first define the relation \sim between patterns:

Definition 3.3.1 *Two patterns φ and φ' are in relation \sim , i.e. $\varphi \sim \varphi'$, iff $\llbracket \varphi \rrbracket = \llbracket \varphi' \rrbracket$.*

According to Definition 3.3.1, two patterns φ and φ' are in relation \sim if they are matched by the same set of configurations, i.e. $\{\gamma \mid (\exists \rho)(\gamma, \rho) \models_{\text{ML}} \varphi\}$ and $\{\gamma \mid (\exists \rho)(\gamma, \rho) \models_{\text{ML}} \varphi'\}$ are equal. Therefore, \sim is an equivalence relation. We let $[\varphi]_{\sim}$ denote the equivalence class of φ . The symbolic transition relation is a relation between such equivalence classes:

Definition 3.3.2 (Symbolic transition relation) *We define the symbolic transition relation $\Rightarrow_{\mathcal{S}}^s$ by: $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^s [\varphi']_{\sim}$ iff there is $\tilde{\varphi} \triangleq \pi \wedge \phi$, such that $\tilde{\varphi} \sim \varphi$ and all the variables in $vars(\pi)$ have data sorts, and there is a rule $\alpha \triangleq \varphi_1 \Rightarrow \varphi_2 \in \mathcal{S}$ with $\varphi_i \triangleq \pi_i \wedge \phi_i$ for $i \in \{1, 2\}$,*

$\text{vars}(\tilde{\varphi}) \cap \text{vars}(\alpha) = \emptyset$, π_1 and π are concretely unifiable, and $[\varphi']_{\sim} = [\sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)]_{\sim}$, where $\sigma_{\pi}^{\pi_1}$ is the most general symbolic unifier of π, π_1 (cf. Lemma 3.2.1), extended as the identity substitution over the variables in $\text{vars}(\phi_1, \phi_2) \setminus \text{vars}(\pi, \pi_1)$.

The condition $\text{vars}(\varphi) \cap \text{vars}(\alpha) = \emptyset$ is not a real restriction because we can always use variable renaming for rules in \mathcal{S} .

We call *symbolic rewriting* the manner in which rules are applied in Definition 3.3.2.

Example 3.3.1 Recall the rule for division from the semantics of *CinK* shown in Example 3.1.5. The left hand-side of the rule is linear and all its subterms of data sorts are variables. Let us consider a pattern $\varphi \triangleq \langle \langle X/Y \curvearrowright \cdot \rangle_{\text{k}} \langle Y' \mapsto L \rangle_{\text{env}} \langle L' \mapsto A' \rangle_{\text{store}} \langle S \rangle_{\text{stack}} \langle R \rangle_{\text{return}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \rangle_{\text{cfg}} \wedge Y' =_{\text{Int}} Y \wedge L' =_{\text{Int}} L \wedge A' =_{\text{Int}} A +_{\text{Int}} 1 \wedge A \neq_{\text{Int}} -1$. Note that φ is linear and all its data subterms are variables. The division rule generates a symbolic transition from $[\varphi]_{\sim}$ to $[\varphi']_{\sim}$, where the pattern φ' is $\langle \langle X/_{\text{Int}}Y \curvearrowright \cdot \rangle_{\text{k}} \langle Y' \mapsto L \rangle_{\text{env}} \langle L' \mapsto A' \rangle_{\text{store}} \langle S \rangle_{\text{stack}} \langle R \rangle_{\text{return}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \rangle_{\text{cfg}} \wedge Y' =_{\text{Int}} Y \wedge L' =_{\text{Int}} L \wedge A' =_{\text{Int}} A +_{\text{Int}} 1 \wedge A \neq_{\text{Int}} -1 \wedge Y \neq_{\text{Int}} 0$.

Assumption 3.3.1 From now on we assume that for all elementary patterns $\varphi \triangleq \pi \wedge \phi$, $\varphi' \triangleq \pi' \wedge \phi'$ such that $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\text{s}} [\varphi']_{\sim}$, π and π' may only have variables of data sorts. This can be obtained by starting with an initial pattern satisfying these assumptions and by ensuring that they are preserved by the rules \mathcal{S} . Thus, our symbolic execution framework allows symbolic data but no symbolic code.

We now show that the symbolic execution thus defined is related with concrete execution via the coverage and precision properties stated in the introduction of this chapter.

3.3.1 Coverage

The coverage property states that the symbolic transition system forward-simulates the concrete transition system.

Lemma 3.3.1 (Coverage) If $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in \llbracket \varphi \rrbracket$ (with all variables in $\text{vars}(\varphi)$ of data sorts) then there exists φ' such that $\gamma' \in \llbracket \varphi' \rrbracket$ and $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\text{s}} [\varphi']_{\sim}$.

Proof Let $\varphi \triangleq \pi \wedge \phi$. From $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ we obtain the rule $\alpha \triangleq \pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}$ and the valuation $\rho : Var \rightarrow \mathcal{T}$ such that $\gamma = \rho(\pi_1)$, $\rho \models_{\text{ML}} \phi_1$, $\gamma' = \rho(\pi_2)$, and $\rho \models_{\text{ML}} \phi_2$. From $\gamma \in \llbracket \varphi \rrbracket$ we obtain the valuation $\mu : Var \rightarrow \mathcal{T}$ such that $\gamma = \mu(\pi)$ and $\mu \models_{\text{ML}} \phi$. We assume, without loss of generality, that $\text{vars}(\varphi) \cap \text{vars}(\alpha) = \emptyset$ (which can always be obtained by variable renaming). The basic patterns π_1 and π are concretely unifiable if we choose $\rho|_{\text{vars}(\pi_1)} \uplus \mu|_{\text{vars}(\pi)}$ as their concrete unifier. Then, using Lemma 3.2.1 we obtain their unique most-general symbolic unifier $\sigma_{\pi}^{\pi_1}$, whose codomain is $T_{\Sigma}(\text{vars}(\pi_1) \cup \text{vars}(\pi))$. Let $\eta : \text{vars}(\pi_1) \cup \text{vars}(\pi) \rightarrow \mathcal{T}$ be the valuation such that $\eta \circ \sigma_{\pi}^{\pi_1} = \rho|_{\text{vars}(\pi_1)} \uplus \mu|_{\text{vars}(\pi)}$ (cf. Definition 3.2.1). We extend $\sigma_{\pi}^{\pi_1}$ to $\text{vars}(\varphi) \cup \text{vars}(\alpha)$ by letting it be the identity on $(\text{vars}(\varphi) \cup \text{vars}(\alpha)) \setminus \text{vars}(\pi_1, \pi)$, and extend η to $\text{vars}(\varphi) \cup \text{vars}(\alpha)$ such that $\eta|_{\text{vars}(\phi_1, \pi_2, \phi_2) \setminus \text{vars}(\pi_1)} = \rho|_{\text{vars}(\phi_1, \pi_2, \phi_2) \setminus \text{vars}(\pi_1)}$ and $\eta|_{\text{vars}(\phi) \setminus \text{vars}(\pi)} = \mu|_{\text{vars}(\phi) \setminus \text{vars}(\pi)}$. With these extensions we have $(\eta \circ \sigma_{\pi}^{\pi_1})(x) = (\rho|_{\text{vars}(\alpha)} \uplus \mu|_{\text{vars}(\varphi)})(x)$ for all $x \in \text{vars}(\varphi) \cup \text{vars}(\alpha)$.

Let $\varphi' \triangleq \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$. Then, according to Definition 3.3.2 we have the transition $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\varphi']_{\sim}$. There remains to prove $\gamma' \in \llbracket \varphi' \rrbracket$.

- On the one hand, $\eta(\sigma_{\pi}^{\pi_1}(\pi_2)) = (\eta \circ \sigma_{\pi}^{\pi_1})(\pi_2) = (\rho \uplus \mu)(\pi_2) = \rho(\pi_2) = \gamma'$; thus, there is a valuation η such that $(\gamma', \eta) \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\pi_2)$.
- On the other hand,

$$\begin{aligned}
\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2) & \quad \text{iff} \\
\eta \circ \sigma_{\pi}^{\pi_1} \models_{\text{ML}} (\phi \wedge \phi_1 \wedge \phi_2) & \quad \text{iff} \\
\rho \uplus \mu \models_{\text{ML}} (\phi \wedge \phi_1 \wedge \phi_2) & \quad \text{iff} \\
\mu \models_{\text{ML}} \phi \text{ and } \rho \models_{\text{ML}} \phi_1 \text{ and } \rho \models_{\text{ML}} \phi_2. &
\end{aligned}$$

Since the last relations hold by the hypotheses, it follows $\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$. Note that we used the property: if $\rho : Var \rightarrow \mathcal{T}$ is a valuation and $\sigma : Var \rightarrow T_{\Sigma}(Var)$ a substitution, then $\rho \models_{\text{ML}} \sigma(\varphi)$ iff $\rho \circ \sigma \models_{\text{ML}} \varphi$.

The two above items imply $(\gamma', \eta) \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$, i.e., $(\gamma', \eta) \models_{\text{ML}} \varphi'$, which concludes the proof.

■

Remark 3.3.1 *Note that in the proof of the Lemma 3.3.1, the symbolic transition $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi']_{\sim}$ is given by a rule which corresponds (cf. Definition 3.3.2) to the rule which gives the concrete transition $\gamma \Rightarrow_{\mathcal{S}} \gamma'$.*

The next corollary follows immediately from Lemma 3.3.1. It can be used to draw conclusions about the absence of concrete program executions on a given path from the absence of feasible symbolic executions on the same path.

Corollary 3.3.1 *For every concrete execution $\gamma_0 \Rightarrow_{\mathcal{S}} \gamma_1 \Rightarrow_{\mathcal{S}} \dots \Rightarrow_{\mathcal{S}} \gamma_n \Rightarrow_{\mathcal{S}} \dots$, and pattern φ_0 such that $\gamma_0 \in \llbracket \varphi_0 \rrbracket$ (with all the variables in $\text{vars}(\varphi_0)$ of data sorts), there exists a symbolic execution $[\varphi_0]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi_1]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \dots \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi_n]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} \dots$ such that $\gamma_i \in \llbracket \varphi_i \rrbracket$ for $i = 0, 1, \dots$*

Proof From the hypothesis we have $\gamma_0 \Rightarrow_{\mathcal{S}} \gamma_1 \Rightarrow_{\mathcal{S}} \dots \Rightarrow_{\mathcal{S}} \gamma_n \Rightarrow_{\mathcal{S}} \dots$ and φ_0 such that $\gamma_0 \in \llbracket \varphi_0 \rrbracket$ (with all the variables in $\text{vars}(\varphi_0)$ of data sorts). Then, for each $\gamma_i \Rightarrow_{\mathcal{S}} \gamma_{i+1}$, where $i \in \{0, 1, \dots\}$ we have $[\varphi_i]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi_{i+1}]_{\sim}$ and $\gamma_{i+1} \in \llbracket \varphi_{i+1} \rrbracket$ (cf. Lemma 3.3.1).
■

Remark 3.3.2 *Corollary 3.3.1 states that for every concrete execution there exists a symbolic one. Moreover, symbolic execution takes the same path in the program as concrete execution, since there is a one to one correspondence between the rules that generate symbolic and concrete transitions (cf. Remark 3.3.1).*

3.3.2 Precision

The precision property states that the symbolic transition system is backwards simulated by the concrete one. Forward simulation could not hold in this case, because the patterns resulting from a symbolic transition may be unsatisfiable. A pattern φ is satisfiable if there is a configuration γ such that $\gamma \in \llbracket \varphi \rrbracket$; otherwise, φ is unsatisfiable.

Lemma 3.3.2 (Precision) *If $\gamma' \in \llbracket \varphi' \rrbracket$ and $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi']_{\sim}$ then there exists $\gamma \in \mathcal{T}_{\text{cfg}}$ such that $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in \llbracket \varphi \rrbracket$.*

Proof Let $\tilde{\varphi} \triangleq \pi \wedge \phi$ such that $\tilde{\varphi} \sim \varphi$ (as shown in Definition 3.3.2). Since $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} [\varphi']_{\sim}$ then (by Definition 3.3.2) there is a rule $\alpha \triangleq \pi_1 \wedge \phi_1 \Rightarrow$

$\pi_2 \wedge \phi_2 \in \mathcal{S}$ such that π_1 and π are concretely unifiable, and $[\varphi']_{\sim} = [\sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)]_{\sim}$, where $\sigma_{\pi}^{\pi_1}$ is the most general unifier (cf. Lemma 3.2.1). Thus, $\varphi' \sim \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$, that is $\llbracket \varphi' \rrbracket = \llbracket \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2) \rrbracket$. From $\gamma' \in \llbracket \varphi' \rrbracket$ and the previous equality we obtain $\eta : \text{Var} \rightarrow \mathcal{T}$ such that $\gamma' = \eta(\sigma_{\pi}^{\pi_1}(\pi_2))$ and $\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$.

We extend $\sigma_{\pi}^{\pi_1}$ to $\text{vars}(\varphi) \cup \text{vars}(\alpha)$ by letting it be the identity on $(\text{vars}(\tilde{\varphi}) \cup \text{vars}(\alpha)) \setminus \text{vars}(\pi_1, \pi)$. Let $\rho : \text{Var} \rightarrow \mathcal{T}$ be defined by $\rho(x) = (\eta \circ \sigma_{\pi}^{\pi_1})(x)$ for all $x \in \text{vars}(\tilde{\varphi}) \cup \text{vars}(\pi_1)$, and $\rho(x) = \eta(x)$ for all $x \in \text{Var} \setminus (\text{vars}(\tilde{\varphi}) \cup \text{vars}(\pi_1))$. Let $\gamma \triangleq \rho(\pi_1)$. We have $\gamma' = \eta(\sigma_{\pi}^{\pi_1}(\pi_2)) = \eta(\pi_2) = \rho(\pi_2)$. From $\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$ we get $\eta \circ \sigma_{\pi}^{\pi_1} \models_{\text{ML}} \phi \wedge \phi_1 \wedge \phi_2$. In particular, $\eta \circ \sigma_{\pi}^{\pi_1} \models_{\text{ML}} \phi_1$ and $\eta \circ \sigma_{\pi}^{\pi_1} \models_{\text{ML}} \phi_2$, i.e., $\rho \models_{\text{ML}} \phi_1$ and $\rho \models_{\text{ML}} \phi_2$. Since $\gamma \triangleq \rho(\pi_1)$, $\gamma' = \rho(\pi_2)$ and $\alpha \in \mathcal{S}$ then there is a transition $\gamma \Rightarrow_{\mathcal{S}} \gamma'$.

There remains to prove $\gamma \in \llbracket \varphi \rrbracket (= \llbracket \tilde{\varphi} \rrbracket)$.

- From $\gamma = \rho(\pi_1)$ using the definition of ρ we get $\gamma = \rho(\pi_1) = (\eta \circ \sigma_{\pi}^{\pi_1})(\pi_1) = \eta(\sigma_{\pi}^{\pi_1}(\pi_1)) = \eta(\sigma_{\pi}^{\pi_1}(\pi)) = (\eta \circ \sigma_{\pi}^{\pi_1})(\pi) = \rho(\pi)$
- From $\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$ and $\eta \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\pi_1)$ iff $\eta \circ \sigma_{\pi}^{\pi_1} \models_{\text{ML}} \phi$ we get $\rho \models_{\text{ML}} \phi$.

Since $\varphi \triangleq \pi \wedge \phi$, the items above imply $(\gamma, \rho) \models \varphi$, i.e., $\gamma \in \llbracket \varphi \rrbracket$, which completes the proof. ■

Remark 3.3.3 *Note that in the proof of Lemma 3.3.2 the rule that generates the concrete transition $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ corresponds (cf. Definition 3.3.2) to the rule that generates the symbolic transition $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^{\S} [\varphi']_{\sim}$.*

We call a symbolic execution *feasible* if all its patterns are satisfiable. The next corollary, which follows from Lemma 3.3.2, can be used to draw conclusions on the existence of concrete program executions on a given path from the existence of feasible symbolic executions on the same path.

Corollary 3.3.2 *For every feasible symbolic execution $[\varphi_0]_{\sim} \Rightarrow_{\mathcal{S}}^{\S} [\varphi_1]_{\sim} \Rightarrow_{\mathcal{S}}^{\S} \dots \Rightarrow_{\mathcal{S}}^{\S} [\varphi_n]_{\sim} \Rightarrow_{\mathcal{S}}^{\S} \dots$ there is a concrete execution $\gamma_0 \Rightarrow_{\mathcal{S}} \gamma_1 \Rightarrow_{\mathcal{S}} \dots \Rightarrow_{\mathcal{S}} \gamma_n \Rightarrow_{\mathcal{S}} \dots$ such that $\gamma_i \in \llbracket \varphi_i \rrbracket$ for $i = 0, 1, \dots$*

Proof Since $[\varphi_0]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\varphi_1]_{\sim} \Rightarrow_{\mathcal{S}}^5 \cdots \Rightarrow_{\mathcal{S}}^5 [\varphi_n]_{\sim} \Rightarrow_{\mathcal{S}}^5 \cdots$ is feasible then for each φ_i , $i \in \{0, 1, \dots\}$, there is $\gamma_i \in \llbracket \varphi_i \rrbracket$. Then, applying Lemma 3.3.2 we choose γ_i such that there is a transition $\gamma_i \Rightarrow_{\mathcal{S}} \gamma_{i+1}$ for $i \in \{0, 1, \dots\}$ and $\gamma_i \in \llbracket \varphi_i \rrbracket$. ■

Remark 3.3.4 *Corollary 3.3.2 states that for every feasible execution there is a corresponding concrete execution. Moreover, the execution follows the same path in the program, since there is a one to one correspondence between the rules that generate transitions (cf. Remark 3.3.3).*

The corollaries in this section say that symbolic execution can be used as a sound program-analysis technique. However, symbolic execution is, in general, not powerful enough for performing program verification, because one can (obviously) only generate bounded-length symbolic executions, whereas program verification, especially in the presence of loops and recursive function calls, would require in general executions of an unbounded length.

Remark 3.3.5 *Let $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2$ be the only rule in \mathcal{S} and patterns $\varphi \triangleq \pi \wedge \phi$, φ' , $\widehat{\varphi} \triangleq \widehat{\pi} \wedge \widehat{\phi}$, $\widehat{\varphi}'$ such that $\varphi \sim \widehat{\varphi}$, $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\varphi']_{\sim}$, and $[\widehat{\varphi}]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\widehat{\varphi}']_{\sim}$, where π_1 and π are concretely unifiable, $[\varphi']_{\sim} \triangleq [\sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)]_{\sim}$, π_1 and $\widehat{\pi}$ are concretely unifiable, and $[\widehat{\varphi}']_{\sim} \triangleq [\sigma_{\widehat{\pi}}^{\pi_1}(\pi_2) \wedge \sigma_{\widehat{\pi}}^{\pi_1}(\widehat{\phi} \wedge \phi_1 \wedge \phi_2)]_{\sim}$.*

Let $\gamma' \in \llbracket \varphi' \rrbracket$. From $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\varphi']_{\sim}$ and Lemma 3.3.2 there is γ such that $\gamma \in \llbracket \varphi \rrbracket$ and $\gamma \Rightarrow_{\mathcal{S}} \gamma'$. Since $\varphi \sim \widehat{\varphi}$ we also obtain $\gamma \in \llbracket \widehat{\varphi} \rrbracket$. By Lemma 3.3.1 with $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in \llbracket \widehat{\varphi} \rrbracket$ as hypotheses, there exists $\widehat{\varphi}''$ such that $[\widehat{\varphi}]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\widehat{\varphi}'']_{\sim}$ and $\gamma' \in \llbracket \widehat{\varphi}'' \rrbracket$. In the proof of Lemma 3.3.1, $\widehat{\varphi}''$ is chosen as $\sigma_{\widehat{\pi}}^{\pi_1}(\pi_2) \wedge \sigma_{\widehat{\pi}}^{\pi_1}(\widehat{\phi} \wedge \phi_1 \wedge \phi_2)$, and thus, $\llbracket \widehat{\varphi}'' \rrbracket = \llbracket \sigma_{\widehat{\pi}}^{\pi_1}(\pi_2) \wedge \sigma_{\widehat{\pi}}^{\pi_1}(\widehat{\phi} \wedge \phi_1 \wedge \phi_2) \rrbracket = \llbracket \widehat{\varphi}' \rrbracket$. Therefore, $\gamma' \in \llbracket \widehat{\varphi}' \rrbracket$.

Conversely, we can show that if $\gamma' \in \llbracket \widehat{\varphi}' \rrbracket$ then $\gamma' \in \llbracket \varphi' \rrbracket$ using a similar approach. Therefore, $\llbracket \varphi' \rrbracket = \llbracket \widehat{\varphi}' \rrbracket$, that is $[\varphi']_{\sim} = [\widehat{\varphi}']_{\sim}$. In conclusion, for two different patterns φ and $\widehat{\varphi}$ such that $\varphi \sim \widehat{\varphi}$, a rule in \mathcal{S} induces at most one symbolic transition, i.e. $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^5 [\varphi']_{\sim} (= [\widehat{\varphi}']_{\sim})$.

Remark 3.3.6 *The symbolic transition relation is finitely branching since the set \mathcal{S} of rules is finite and cf. Remark 3.3.5 each rule determines at most one symbolic transition.*

The Coverage and Precision properties depend on the conditions exhibited in the hypothesis of Lemma 3.2.1. Essentially, these conditions apply to both the rules in the semantics, and the programs that we run symbolically. In practice, we can perform some mechanical transformations over the rules in the semantics such that their are linear and all data subterms are variables. On the other hand, the condition over programs (i.e. all variables are of data sorts) remains a real restriction, and thus, our symbolic execution framework is not able to run programs where pieces of code are data. Moreover, if the conditions of the Lemma are not met then the Coverage and Precision results are not guaranteed to hold, since the symbolic execution relation depends on the existence of the most general unifier (which also depends on these conditions)

3.4 Symbolic Execution via Language Transformation

In this section we show how to achieve symbolic execution in language-definition frameworks, such as \mathbb{K} , where rules are applied by standard rewriting. This forms the basis of our prototype implementation of symbolic execution in the \mathbb{K} framework, which is presented in the Section 5.1.

The symbolic semantics of programming languages (given in Definition 3.3.2) requires rules to be applied in a symbolic manner. A fair question that arises is then: how to implement this symbolic rewriting in a setting where only standard rewriting is available, such as our generic language definition framework, where languages are triples $\mathcal{L} = ((\Sigma, \Pi, Cfg), \mathcal{T}, \mathcal{S})$? The answer is to *transform* a language definition \mathcal{L} into another language definition \mathcal{L}^s , such that standard rewriting in \mathcal{L}^s corresponds to symbolic rewriting in \mathcal{L} .

In the following we define the components of \mathcal{L}^s :

1. (Σ^s, Π^s, Cfg^s) . The signature Σ^s is Σ extended with a new sort *Cond*, with constructors for formulas, and a new sort *Cfg*^s (for symbolic configurations) with the constructor $_{\mathcal{A}^s} : Cfg \times Cond \rightarrow Cfg^s$. This leaves $(\Sigma^s)^\partial$ unchanged, i.e. $(\Sigma^s)^\partial = \Sigma^\partial$. The predicates in Π are

transformed into terms of sort *Cond*, hence, Π is not included in Π^s . Here, we have the freedom to choose Π^s such that $\Rightarrow_{\mathcal{S}^s}$ corresponds to $\Rightarrow_{\mathcal{S}}^s$ (cf. Proposition 3.4.1), or to the subset of $\Rightarrow_{\mathcal{S}}^s$ corresponding to the feasible executions, or to a subset of $\Rightarrow_{\mathcal{S}}^s$ that approximates the feasible executions. We first choose $\Pi^s = \emptyset$. For an elementary pattern $\varphi \triangleq \pi \wedge \phi$, let φ^s be the symbolic configuration $\varphi^s \triangleq \pi \wedge^s \phi$, and for each symbolic configuration $\varphi^s \triangleq \pi \wedge^s \phi$, let $\varphi \triangleq \pi \wedge \phi$ be the corresponding elementary pattern.

2. The data domain \mathcal{D}^s is the set of terms $\bigcup_d T_{\Sigma, d}(Var^{\mathfrak{d}})$, where d is a data sort and $Var^{\mathfrak{d}}$ denotes the subset of variables of data sort. The model \mathcal{T}^s is then the set of ground terms over the signature $(\Sigma^s \setminus (\Sigma^s)^{\mathfrak{d}}) \cup \mathcal{D}^s$.
3. For each rule $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2$ in \mathcal{S} , the set of rules \mathcal{S}^s includes a rule of the form $\pi_1 \wedge^s \psi \Rightarrow \pi_2 \wedge^s (\psi \wedge \phi_1 \wedge \phi_2)$, where ψ is fresh a variable of sort *Cond*.

The first implication in Proposition 3.4.1 says that for each transition in $\Rightarrow_{\mathcal{S}^s}$ there is a transition in $\Rightarrow_{\mathcal{S}}^s$ that is "more general" than the one in $\Rightarrow_{\mathcal{S}^s}$, while the second implication says that for each transition in $\Rightarrow_{\mathcal{S}^s}$ there exists essentially the same transition in $\Rightarrow_{\mathcal{S}}^s$ (up to equivalent pattern conditions).

Proposition 3.4.1 (Relating $\Rightarrow_{\mathcal{S}}^s$ and $\Rightarrow_{\mathcal{S}^s}$)

1. If $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$ then there exist a substitution σ and patterns φ and φ' such that $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^s [\varphi']_{\sim}$ and $\widehat{\varphi} = \sigma(\varphi), \widehat{\varphi}' = \sigma(\varphi')$.
2. If $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^s [\varphi']_{\sim}$ then $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$ for some $\widehat{\varphi} \in [\varphi]_{\sim}, \widehat{\varphi}' \in [\varphi']_{\sim}$.

Proof We recall that, by Assumption 3.3.1, all the free variables occurring in the basic patterns of φ, φ' (and thus also in φ^s, φ'^s and in all patterns \sim -equivalent to φ, φ') are of data sorts. This is used in the proof to apply Lemma 3.2.1.

(1) Assume $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$, where $\widehat{\varphi} \triangleq \pi \wedge \phi$. Then, there is a rule $\pi_1 \wedge^s \psi \Rightarrow \pi_2 \wedge^s (\psi \wedge \phi_1 \wedge \phi_2) \in \mathcal{S}^s$, generated from $\alpha \triangleq \pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}$, and a valuation $\rho^s : Var \cup \{\psi\} \rightarrow \mathcal{T}^s$ such that $\widehat{\varphi}^s = \rho^s(\pi_1 \wedge^s \psi)$ and $\widehat{\varphi}'^s = \rho^s(\pi_2) \wedge^s \rho^s(\psi \wedge \phi_1 \wedge \phi_2)$. From $\rho^s(\pi_1 \wedge^s \psi) = \widehat{\varphi}^s$

we obtain $\pi = \rho^s(\pi_1)$ and $\phi = \rho^s(\psi)$. We assume, without restriction of generality (possibly, after renaming some variables), that $\text{vars}(\widehat{\varphi}^s) \cap \text{vars}(\pi_1 \wedge \phi_1) = \emptyset$. It follows that the restriction of ρ^s to $\text{vars}(\pi_1, \pi)$ is the same with the most-general-unifier $\sigma_\pi^{\pi_1}$ given by Lemma 3.2.1.

Let $\varphi \triangleq \widehat{\varphi}$ and $\varphi' \triangleq \sigma_\pi^{\pi_1}(\pi_2 \wedge \phi \wedge \phi_1 \wedge \phi_2)$. We obviously have $[\varphi]_\sim \Rightarrow_{\mathcal{S}}^s [\varphi']_\sim$ by Definition 3.3.2. We recall that $\sigma_\pi^{\pi_1}$ is extended as identity over $\text{vars}(\alpha, \widehat{\varphi}^s) \setminus \text{vars}(\pi, \pi_1)$ in Definition 3.3.2. Therefore, we consider the substitution σ defined by $\sigma(x) = \rho^s(x)$ for $x \in \text{vars}(\alpha, \widehat{\varphi}^s) \setminus \text{vars}(\pi, \pi_1)$ and as the identity in the rest. We obtain $\sigma(\varphi') = \widehat{\varphi}'$ and $\sigma(\varphi) = \sigma(\widehat{\varphi}) = \widehat{\varphi}$ by the definition of σ , which concludes the first implication of the proposition.

(2) Assume $[\varphi]_\sim \Rightarrow_{\mathcal{S}}^s [\varphi']_\sim$. Then, there is $\widehat{\varphi} \triangleq \pi \wedge \phi \in [\varphi]_\sim$ and a rule $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}$ such that $[\varphi']_\sim = [\sigma_\pi^{\pi_1}(\pi_2) \wedge \sigma_\pi^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)]_\sim$, where $\sigma_\pi^{\pi_1}$ is the most general unifier of π and π_1 , cf. Lemma 3.2.1. We choose $\widehat{\varphi}' \triangleq \sigma_\pi^{\pi_1}(\pi_2) \wedge \sigma_\pi^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)$. Lemma 3.2.1 also says that $\sigma_\pi^{\pi_1}$ can be decomposed into a substitution σ over $\text{vars}(\pi_1)$ such that $\sigma(\pi_1) = \pi$, and the identity substitution over $\text{vars}(\pi)$. Since $\text{vars}(\pi) \cap \text{vars}(\pi_1) = \emptyset$ we can obtain, possibly by renaming variables, that $\text{vars}(\pi \wedge \phi) \cap \text{vars}(\pi_1) = \emptyset$, and thus σ has no effect on ϕ . We obtain $\widehat{\varphi}' = \sigma(\pi_2) \wedge \phi \wedge \sigma(\phi_1 \wedge \phi_2)$. Let then $\rho^s : \text{Var} \cup \{\psi\} \rightarrow \mathcal{T}^s$ be any valuation such that $\rho^s(x) = \sigma(x)$ for all $x \in \text{vars}(\pi_1) \cup \text{vars}(\pi_2)$ and $\rho^s(\psi) = \phi$. We obtain $\rho^s(\pi_1 \wedge \psi) = \widehat{\varphi}^s$ and $\rho^s(\pi_2 \wedge (\psi \wedge \phi_1 \wedge \phi_2)) = \widehat{\varphi}'^s$, which means that the transition $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$ is generated by the rule $\pi_1 \wedge \psi \Rightarrow \pi_2 \wedge (\psi \wedge \phi_1 \wedge \phi_2) \in \mathcal{S}^s$ and the valuation ρ^s .

■

A coverage result (cf. Section 3.3.1) relating $\Rightarrow_{\mathcal{S}}$ and $\Rightarrow_{\mathcal{S}^s}$ follows from the original coverage result and the second implication of Proposition 3.4.1. Let us assume $\gamma \Rightarrow_{\mathcal{S}} \gamma'$ and $\gamma \in \llbracket \widehat{\varphi}^s \rrbracket$, where $\llbracket \widehat{\varphi}^s \rrbracket \triangleq \llbracket \widehat{\varphi} \rrbracket$. Since $\varphi \sim \widehat{\varphi}$ (i.e. $\llbracket \varphi \rrbracket = \llbracket \widehat{\varphi} \rrbracket$) then $\gamma \in \llbracket \varphi \rrbracket$, which, together with $\gamma \Rightarrow_{\mathcal{S}} \gamma'$, implies, thanks to the original coverage result, that there is φ' such that $\gamma' \in \llbracket \varphi' \rrbracket$ and $[\varphi]_\sim \Rightarrow_{\mathcal{S}}^s [\varphi']_\sim$. From the second implication of Proposition 3.4.1 we have $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$ and $\widehat{\varphi}' \in [\varphi']_\sim$ (i.e. $\llbracket \varphi' \rrbracket = \llbracket \widehat{\varphi}' \rrbracket$), and thus we obtain $\gamma' \in \llbracket \widehat{\varphi}'^s \rrbracket$. This gives the coverage of $\Rightarrow_{\mathcal{S}}$ by $\Rightarrow_{\mathcal{S}^s}$.

A precision result relating $\Rightarrow_{\mathcal{S}}$ by $\Rightarrow_{\mathcal{S}^s}$ also follows from the original precision result and the first implication of Proposition 3.4.1, by using the fact (taken from its proof) that $\varphi = \widehat{\varphi}$. Indeed, assume $\widehat{\varphi}^s \Rightarrow_{\mathcal{S}^s} \widehat{\varphi}'^s$

and $\gamma' \in \llbracket \widehat{\varphi}'^s \rrbracket$. Since $\widehat{\varphi}' = \sigma(\varphi')$ it follows that $\gamma' \in \llbracket \varphi' \rrbracket$ (i.e., φ being more general than $\widehat{\varphi}'$, it contains all the instances of $\widehat{\varphi}'$, including the configuration γ'). Using $[\varphi]_{\sim} \Rightarrow_{\mathcal{S}}^s [\varphi']_{\sim}$ and the original precision result, we obtain a configuration $\gamma \in \llbracket \varphi \rrbracket$ such that $\gamma \Rightarrow_{\mathcal{S}} \gamma'$, and since the proof of the first implication gave $\varphi = \widehat{\varphi}$, we have $\gamma \in \llbracket \widehat{\varphi}^s \rrbracket$, proving the precision of $\Rightarrow_{\mathcal{S}^s}$ with respect to $\Rightarrow_{\mathcal{S}}$.

Remark 3.4.1 *The coverage and precision results relating $\Rightarrow_{\mathcal{S}}$ and $\Rightarrow_{\mathcal{S}^s}$ are important because they say that $\Rightarrow_{\mathcal{S}^s}$ has the natural properties expected from symbolic execution.*

The feasible symbolic executions can be obtained as executions of another slightly different definition of \mathcal{L}^s , at least theoretically, by considering $\Pi_{Cond}^s = \{sat\}$ and $\Pi_s^s = \emptyset$ for all sorts s other than $Cond$, with the interpretation $\mathcal{T}_{sat}^s = \{\phi \mid \phi \text{ is satisfiable in } \mathcal{T}\}$, and by taking in \mathcal{S}^s the following conditional rules, for each $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}$:

$$(\pi_1 \wedge^s \psi) \wedge sat(\psi \wedge \phi_1 \wedge \phi_2) \Rightarrow \pi_2 \wedge^s (\psi \wedge \phi_1 \wedge \phi_2)$$

Recall that $\pi_1 \wedge^s \psi$, $\pi_2 \wedge^s (\psi \wedge \phi_1 \wedge \phi_2)$ are terms in \mathcal{L}^s , thus, the expression $(\pi_1 \wedge^s \psi) \wedge sat(\psi \wedge \phi_1 \wedge \phi_2)$ is a pattern and the above rule is well formed.

However, this definition of \mathcal{L}^s cannot be implemented in practice because the satisfiability problem for FOL is undecidable. To deal with this issue we implemented a slightly modified version of \mathcal{L}^s which approximates feasible symbolic executions. This can be done by using a predicate symbol $nsat$ instead of sat , such that its interpretation is sound, i.e., $\mathcal{T}_{nsat}^s \subsetneq \{\phi \mid \phi \text{ not satisfiable in } \mathcal{T}\}$, and is computable. Then, the rules in \mathcal{S}^s have the form $(\pi_1 \wedge^s \psi) \wedge \neg nsat(\psi \wedge \phi_1 \wedge \phi_2) \Rightarrow \pi_2 \wedge^s (\psi \wedge \phi_1 \wedge \phi_2)$. The predicate $nsat$ is an approximation of the theoretical $unsat$ predicate. This version of \mathcal{L}^s is what we actually implemented in our prototype tool in \mathbb{K} Framework, which is presented in Section 5.1.

Chapter 4

Applications

Symbolic execution has been used as a powerful program analysis technique in many applications (Section 1.3). In this dissertation, we focused mainly on using it for program verification. This chapter describes two approaches based on Hoare Logic and Reachability Logic, respectively. In Section 4.1 we present our approach from [4], where we show that, for a given language, we can use our symbolic execution framework to verify Hoare triples. We present a Hoare Logic proof system for IMP, and its equivalent proof system given using reachability rules, obtained by a transformation proposed in [100]. We prove that symbolic execution yields a reachability rule which can be derived with this proof system. By its nature, verification based on Hoare Logic is a bit limited. First, it is not easy to express with Hoare triples properties for programs which involve recursive function calls, blocks with local variables, virtual method calls, etc. Second, Hoare Logic is language dependent, i.e. for every new language construct we may need additional proof rules. On the other hand, our symbolic execution is language independent, and we take advantage of this in Section 4.2, where we propose a language independent proof system, which has symbolic execution as a distinguished proof rule, and a default strategy for applying the proof rules, such that we can prove RL formulas. Formally, we prove that our deductive system together with the default strategy is sound and weakly-complete. An implementation of this proof system and examples are discussed in Section 5.3.

4.1 Hoare-Logic Verification

Hoare Logic [54] (also known as Floyd-Hoare Logic) is probably the most popular formal system for reasoning about the correctness of computer programs. Proposed in 1969 by C. A. R. Hoare, the logic was inspired from the work of R. Floyd on flowcharts [42], and it consists in a set of deduction rules which prove the validity of Hoare triples. Triples are formulas of the form $\langle\psi\rangle St \langle\psi'\rangle$ representing program properties. St is a piece of code (statements), while ψ and ψ' are program assertions, known as *precondition* and *postcondition*, respectively. Intuitively, such formulas have the following semantics: a triple $\langle\psi\rangle St \langle\psi'\rangle$ holds, w.r.t. the partial correctness, if and only if any terminating execution of St , starting in a state satisfying the precondition ψ , ends in a state satisfying the postcondition ψ' . If further, the precondition ensures the termination of St , then $\langle\psi\rangle St \langle\psi'\rangle$ holds w.r.t. total correctness.

In [4] we present a methodology which, applied to the \mathbb{K} definition of a programming language, turns the \mathbb{K} runner into a Hoare-like verifier. The work was inspired from [49], where the authors present methods for proving Hoare formulas based both on the forward computation of postconditions and on the backwards computation of preconditions. According to [49], a Hoare triple can be verified by symbolically executing separate paths and combining the results. Intuitively, we follow the same approach: we first perform symbolic execution starting in a state where the precondition is the initial path condition, and after the execution we verify if the collected path condition implies the postcondition. If the implication holds for all program paths, then the program is proved correct.

To our best knowledge, there are only a few notable Hoare Logic verification approaches based on the formal semantics of a language (e.g. Caduceus [18], Java-ITP [104], KeY-Hoare [43], Spec# [15]). Others (e.g. ITS [50], JAlgo [58]) are not even based on formal definitions of programming languages.

In this section we show that symbolic execution can be used to verify Hoare triples. The approach is based on a mechanical translation of Hoare triples into reachability rules. The obtained reachability rules can be proved by executing symbolically their left hand side and by checking if the resulted patterns imply the right hand side of the rule.

We present the proof system of Hoare Logic [HL-IMP] by means of a simple imperative language, called IMP (Section 4.1.1). This proof system has been shown sound in [100], using a technique which mechanically translates Hoare triples into reachability rules and Hoare Logic proof derivations into Reachability Logic proof derivations. We show this translation in Section 4.1.2, and then we use it to construct the equivalent Hoare Logic proof system [HL'-IMP] given using reachability rules. Then, we extend IMP with syntax and semantics for Hoare-like annotations (preconditions, postconditions, and invariants) and we show that symbolic execution in the annotated language yields a proof derivation in [HL'-IMP] (Section 4.1.3).

```

Id ::= domain of identifiers
Int ::= domain of integers (including operations)
Bool ::= domain of booleans (including operations)
AExp ::= Int
        | Id
        | AExp / AExp [strict]
        | AExp * AExp [strict]
        | AExp + AExp [strict]
        | ( AExp )
BExp ::= Bool
        | AExp <= AExp [strict]
        | not BExp [strict]
        | BExp and BExp [strict(1)]
        | ( BExp )
Stmt ::= skip
        | Id = AExp
        | if BExp Stmt else Stmt [strict(1)]
        | { Stmt }
        | while BExp do Stmt
        | Stmt ; Stmt

```

Figure 4.1: \mathbb{K} Syntax of IMP

4.1.1 IMP

Our running example is IMP, a simple imperative language intensively used in research papers. The \mathbb{K} syntax of this programming language is shown in Figure 4.1. The IMP statements are either assignments, if statements, while loops, skip (i.e., the empty statement), or a list of statements.

The IMP configuration consists only of the program to be executed and an environment mapping variables to values. Precisely, the $\langle \rangle_k$ cell contains the code and the $\langle \rangle_{\text{env}}$ cell the environment:

$$\text{Cfg} ::= \langle \langle \text{Code} \rangle_k \langle \text{Map}_{Id, Int} \rangle_{\text{env}} \rangle_{\text{cfg}}$$

The semantics of IMP is shown in Figure 4.2. As usual, \mathbb{K} rules specify how the configuration evolves during the execution. Most syntactical constructs require only one semantical rule. The exceptions are the **and** and boolean operation and the **if** statement, which have boolean arguments and require two rules each.

$$\begin{aligned}
\langle \langle I_1 + I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle I_1 +_{Int} I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle I_1 * I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle I_1 *_{Int} I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle I_1 / I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle I_1 /_{Int} I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} I_2 \neq_{Int} 0 \\
\langle \langle I_1 \leq I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle I_1 \leq_{Int} I_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{true and } B \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle B \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{false and } B \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle \text{false} \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{not } B \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle \neg B \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{skip} \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle S_1; S_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle S_1 \curvearrowright S_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \{ S \} \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle S \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{if true } S_1 \ \text{else } S_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle S_1 \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{if false } S_1 \ \text{else } S_2 \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \langle \langle S_2 \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle \text{while } B \ \text{do } S \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} &\Rightarrow \\
&\quad \langle \langle \text{if } B \ \{ S ; \text{while } B \ \text{do } S \} \ \text{else skip} \ \dots \rangle_k \ \dots \rangle_{\text{cfg}} \\
\langle \langle X \ \dots \rangle_k \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} &\Rightarrow \langle \langle \text{lookup}(X, M) \ \dots \rangle_k \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \\
\langle \langle X = I \ \dots \rangle_k \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} &\Rightarrow \langle \langle \ \dots \rangle_k \langle \text{update}(X, M, I) \rangle_{\text{env}} \rangle_{\text{cfg}}
\end{aligned}$$

Figure 4.2: \mathcal{S}_{IMP} : the \mathbb{K} Semantics of IMP

In addition to the rules shown in Figure 4.2 the semantics of IMP includes the heating/cooling rules, induced by the *strict* attribute. We show only the case of the **if** statement, which is strict in the first argument. The evaluation is done by the following rules:

$$\begin{array}{ll}
 \text{(heating)} & \text{(cooling)} \\
 A_1 + A_2 \Rightarrow A_1 \curvearrowright \square + A_2 & I_1 \curvearrowright \square + A_2 \Rightarrow I_1 + A_2 \\
 A_1 + A_2 \Rightarrow A_2 \curvearrowright A_1 + \square & I_2 \curvearrowright A_1 + \square \Rightarrow A_1 + I_2
 \end{array}$$

Here, A_1 and A_2 are terms of sort $AExp$, I_1 and I_2 are integers, and \square is destined to receive an integer value (once it is computed, typically, by the other rules in the semantics).

Hoare Logic

Hoare Logic provides a formal system for reasoning about program correctness. Using the axioms and the deduction rules of the proof system one can prove Hoare triples $\langle \psi \rangle St \langle \psi' \rangle$, where St is a program statement, and ψ and ψ' are assertions involving program variables.

Definition 4.1.1 (Hoare triples) *A Hoare triple is a formula of the form $\langle \psi \rangle St \langle \psi' \rangle$, where St is a program statement, and ψ and ψ' are assertions defined as below:*

- every boolean expression is an assertion;
- if ψ is an assertion then $\neg\psi$ is an assertion;
- if ψ and ψ' are assertions then $\psi \wedge \psi'$ is an assertion;
- if ψ is an assertion and V a set of variables then $(\exists V)\psi$ is an assertion,

where the quantifier \exists , and the connectives \neg , \wedge are the same as in FOL.

In IMP, boolean expressions correspond to sort $BExp$. The other known FOL connectives (\rightarrow , \vee) and the quantifier \forall can be also used in assertions, being expressed in terms of the ones above as explained in Section 2.2.

Next, we define the satisfaction relation between a program configuration and an assertion:

$$\begin{array}{c}
\text{[HL-skip]} \frac{}{\langle \psi \rangle \text{skip} \langle \psi \rangle} \qquad \text{[HL-block]} \frac{\langle \psi \rangle St \langle \psi' \rangle}{\langle \psi \rangle \{ St \} \langle \psi' \rangle} \\
\text{[HL-assign]} \frac{}{\langle \psi \rangle \mathbf{a} = E \langle (\exists u) \mathbf{a} = E[u/\mathbf{a}] \wedge \psi[u/\mathbf{a}] \rangle} \\
\text{[HL-seq]} \frac{\langle \psi_1 \rangle St_1 \langle \psi_2 \rangle \quad \langle \psi_2 \rangle St_2 \langle \psi_3 \rangle}{\langle \psi_1 \rangle St_1 ; St_2 \langle \psi_3 \rangle} \\
\text{[HL-while]} \frac{\langle \psi \wedge BE \rangle St \langle \psi \rangle}{\langle \psi \rangle \mathbf{while} (BE) \mathbf{do} St \langle \psi \wedge \neg BE \rangle} \\
\text{[HL-if]} \frac{\langle \psi \wedge BE \rangle St_1 \langle \psi' \rangle \quad \langle \psi \wedge \neg BE \rangle St_2 \langle \psi' \rangle}{\langle \psi \rangle \mathbf{if} (BE) St_1 \mathbf{else} St_2 \langle \psi' \rangle} \\
\text{[HL-cseq]} \frac{\psi \rightarrow \psi_1 \quad \langle \psi_1 \rangle St \langle \psi'_1 \rangle \quad \psi'_1 \rightarrow \psi'}{\langle \psi \rangle St \langle \psi' \rangle}
\end{array}$$

Figure 4.3: [HL-IMP]: Floyd-Hoare proof system for IMP

Definition 4.1.2 Given a configuration $\gamma \triangleq \langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}$ and an assertion ψ the satisfaction relation $\gamma \models_{\text{HL}} \psi$, read γ satisfies ψ , is defined as follows:

- if ψ is a boolean expression then $\gamma \models_{\text{HL}} \psi$ iff $\langle \langle \psi \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow_{\mathcal{S}_{\text{IMP}}}^* \langle \langle \text{true} \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}$;
- $\gamma \models_{\text{HL}} \neg \psi$ iff $\gamma \not\models_{\text{HL}} \psi$ does not hold;
- $\gamma \models_{\text{HL}} \psi \wedge \psi'$ iff $\gamma \models_{\text{HL}} \psi$ and $\gamma \models_{\text{HL}} \psi'$;
- $\gamma \models_{\text{HL}} (\exists V) \psi$ iff there exists a configuration $\tau \triangleq \langle \langle St \rangle_k \langle \sigma' \rangle_{\text{env}} \rangle_{\text{cfg}}$ with $\sigma'(x) = \sigma(x)$ for all $x \notin V$, such that $\tau \models_{\text{HL}} \psi$.

A triple $\langle \psi \rangle St \langle \psi' \rangle$ is called valid (w.r.t. partial correctness) if and only if for any program configuration $\gamma \triangleq \langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}$ such that γ satisfies ψ , if $\langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow_{\mathcal{S}_{\text{IMP}}} \langle \langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \rangle_{\text{cfg}}$ then $\gamma' \triangleq \langle \langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \rangle_{\text{cfg}}$ satisfies ψ' . We denote by $\mathcal{S}_{\text{IMP}} \models_{\text{HL}} \langle \psi \rangle St \langle \psi' \rangle$ the fact that $\langle \psi \rangle St \langle \psi' \rangle$ is valid for \mathcal{S}_{IMP} .

The Hoare Logic deduction system for IMP is shown in Figure 4.3. It contains a set of deduction rules which correspond to each statement of the language. The deduction system can be used only to prove Hoare triples w.r.t partial correctness. In [100], the authors have presented and proved correct a technique to automatically translate Hoare triples into reachability rules and Hoare Logic proof derivations into Reachability Logic proof derivations. Based in this translation, the authors prove the soundness of the proof system shown in Figure 4.3 by relying on the soundness of Reachability Logic.

4.1.2 From Hoare Logic to Reachability Logic

For the IMP programming language, a Hoare triple $\langle\psi\rangle St \langle\psi'\rangle$ can be encoded as the following reachability rule [100]:

$$(\exists X)\langle\langle St \rangle_{\mathbf{k}} \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X)\langle\langle \cdot \rangle_{\mathbf{k}} \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y} \quad (4.1)$$

X and Y are two disjoint sets of variables corresponding to two sets \mathbf{X} , \mathbf{Y} , which contain the program variables which appear in St , and the additional variables that appear in ψ and ψ' but not in \mathbf{X} , respectively \mathbf{Y} . We use X and Y instead of \mathbf{X} and \mathbf{Y} because Reachability Logic distinguishes between program and logical variables. This restriction does not limit Hoare Logic in any way, because we can always choose \mathbf{X} as being the set of all program variables about which we want to reason and \mathbf{Y} the rest of variables occurring anywhere in the state or the specification, but different from \mathbf{X} . The state σ_X maps each $\mathbf{x} \in \mathbf{X}$ to its corresponding $x \in X$. Formulas $\psi_{X,Y}$ and $\psi'_{X,Y}$ are ψ and respectively ψ' with $\mathbf{x} \in \mathbf{X}$ and $\mathbf{y} \in \mathbf{Y}$ replaced by their corresponding $x \in X$ and $y \in Y$, and each expression construct (e.g. $+$) replaced by its mathematical correspondent (e.g. $+_{Int}$). We use an additional function $eval_{\sigma_X}$ to perform the same transformation over program expressions (e.g. $eval_{\sigma_X}(\mathbf{x} + \mathbf{y}) = x +_{Int} y$). If E is a program expression then $E' = eval_{\sigma_X}(E)$ is given by evaluating E in the environment σ_X , that is, $\langle\langle E \rangle_{\mathbf{k}} \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow^* \langle\langle E' \rangle_{\mathbf{k}} \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}$, where E' is a symbolic expression of sort Int (and, therefore, $\langle\langle E' \rangle_{\mathbf{k}} \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}$ is irreducible).

Example 4.1.1 *Let $\langle\mathbf{a} > 0\rangle \mathbf{a} = \mathbf{a} + \mathbf{b} \langle(\exists u)(\mathbf{a} = u + \mathbf{b} \wedge u > 0)\rangle$ be a Hoare triple. The set of program variables is $\mathbf{X} = \{\mathbf{a}, \mathbf{b}\}$, while the set of logical*

variables is $X = \{a_X, b_X\}$. Both Y and Z are empty sets. The state is $\sigma_X = \{\mathbf{a} \mapsto a_X \mathbf{b} \mapsto b_X\}$, and $\psi_{X,Y} = a_X >_{Int} 0$ and $\psi'_{X,Y} = (\exists u)(a_X =_{Int} u +_{Int} b_X \wedge u >_{Int} 0)$. By applying transformation (4.1) we obtain the following reachability rule:

$$\begin{aligned} & (\exists \{a_X, b_X\}) \langle \langle \mathbf{a} = \mathbf{a} + \mathbf{b} \rangle_k \langle \mathbf{a} \mapsto a_X \mathbf{b} \mapsto b_X \rangle_{env} \rangle_{cfg} \wedge a_X >_{Int} 0 \Rightarrow \\ & (\exists \{a_X, b_X\}) \langle \langle \cdot \rangle_k \langle \mathbf{a} \mapsto a_X \mathbf{b} \mapsto b_X \rangle_{env} \rangle_{cfg} \wedge (\exists u)(a_X =_{Int} u +_{Int} b_X \wedge u >_{Int} 0) \end{aligned} \quad (4.2)$$

The assertion $\psi'_{X,Y}$ was obtained by making use the $eval_{\sigma_X}$ function, as follows: $eval_{\sigma_X}((\mathbf{a} + \mathbf{b})[u/\mathbf{a}]) = eval_{\sigma_X}(u + \mathbf{b}) = u +_{Int} b_X$.

We applied the transformation (4.1) to the proof system from Figure 4.3 and we obtained the equivalent proof system [HL'-IMP] given using reachability rules, which shown in Figure 4.4. This proof system is used in the next section to implement Hoare Logic using symbolic execution.

4.1.3 Hoare Logic by symbolic execution

The goal of this section is to show how to prove Hoare triples using symbolic execution. As pointed out previously in this section, triples $(\psi)St(\psi)$ can be proved by proving their corresponding reachability rules $\varphi \Rightarrow \varphi'$. Note that the transformation (4.1) generates a particular type of reachability rules, i.e. φ' contains an empty $\langle \cdot \rangle_k$ cell. In the case of IMP, we can prove such a rule by executing symbolically φ , and then by verifying if every (result) pattern φ'' implies φ' . However, symbolic execution may not finish due to loops in programs, and thus, we need a way to provide invariants for loops.

In the following, we show how one can extend the definition of IMP with support not only for invariants, but also for preconditions and postconditions. This extension is done in two steps:

1. Add syntax and semantics for preconditions, postconditions, and invariants.
2. Define `assume` and `implies` operations for verifying RL formulas.

For IMP, the syntax is enriched with a new sort *Assert* for assertions, which extends sort *BExp* (boolean expressions of the language). This

$$\begin{array}{c}
\text{[HL'-skip]} \frac{}{(\exists X) \langle \langle \text{skip} \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \rangle_{\text{cfg}}} \\
\text{[HL'-assign]} \frac{}{(\exists X) \langle \langle \mathbf{a} = E \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge ((\exists u)(a_X = \text{eval}_{\sigma_X}(E[u/\mathbf{a}]) \wedge \psi_{X,Y}[u/\text{eval}_{\sigma_X}(\mathbf{a})]))} \\
\text{[HL'-seq]} \frac{(\exists X) \langle \langle St_1 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^1 \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^2 \quad (\exists X) \langle \langle St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^2 \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^3}{(\exists X) \langle \langle St_1 St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^1 \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^3} \\
\text{[HL'-block]} \frac{(\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}}{(\exists X) \langle \langle \{ St \} \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}} \\
\text{[HL'-while]} \frac{(\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge BE \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}}{(\exists X) \langle \langle \text{while } (BE) \text{ do } St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge \neg BE} \\
\text{[HL'-if]} \frac{(\exists X) \langle \langle St_1 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge BE \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y} \quad (\exists X) \langle \langle St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge \neg BE \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}}{(\exists X) \langle \langle \text{if } (BE) \text{ } St_1 \text{ else } St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}} \\
\text{[HL'-cseq]} \frac{\psi_{X,Y} \rightarrow \psi_{X,Y}^1 \quad (\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^1 \Rightarrow \psi_{X,Y}^2 \rightarrow \psi'_{X,Y} \quad (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}^2}{(\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle \langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}}
\end{array}$$

Figure 4.4: [HL'-IMP]: the equivalent Floyd-Hoare proof system given using reachability rules. Each Hoare triple from the proof system in Figure 4.3 has been transformed into an equivalent reachability rule using (4.1).

allows boolean and arithmetic expressions (including program variables) in annotations. Alternatively, one can define a custom syntax for *Assert*, which also permits program variables. Once we have a sort for assertions we can give syntax for the preconditions, postconditions, and invariants. For IMP we assume that triples have the following syntax:

pre: *Assert* *post*: *Assert Stmt*

The syntax of the annotated **while** loop is:

while *BExp inv*: *Assert Stmt*

According to transformation (4.1), the precondition and the program become parts of the left hand side of the corresponding reachability rule. The right hand side of the reachability rule contains the postcondition, and an empty $\langle \cdot \rangle_k$ cell. The semantics of a Hoare triple can be given using the additional constructs **assume** and **implies** (defined later), as follows:

$$\left\langle \frac{\text{pre: } \psi \text{ post: } \psi' \text{ } St}{\text{assume}(\langle St \rangle_k \langle \sigma \rangle_{\text{env}} \langle \psi \rangle_{\text{cond}}) \rightsquigarrow \text{implies}(\langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \langle \psi' \rangle_{\text{cond}})} \right\rangle_k$$

The intuition of the rule above is that it assumes the code to be executed and the precondition, and after the execution, it checks if the postcondition holds. The $\langle \cdot \rangle_{\text{cond}}$ cell is used to hold the path condition.

The operational semantics of the **while** loop with invariants, is given by the following rewrite rules:

$$\left\langle \frac{\text{while}(B)\text{inv}:\psi \text{ } St \rightsquigarrow K}{\text{assume}(\langle St \rangle_k \langle \sigma \rangle_{\text{env}} \langle B \wedge \psi \rangle_{\text{cond}}) \rightsquigarrow \text{implies}(\langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \langle \psi \rangle_{\text{cond}})} \right\rangle_k$$

$$\left\langle \frac{\text{while}(B)\text{inv}:\psi \text{ } St \rightsquigarrow K}{\text{implies}(\langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \langle \psi \rangle_{\text{cond}}) \rightsquigarrow \text{assume}(\langle K \rangle_k \langle \sigma \rangle_{\text{env}} \langle \neg B \wedge \psi \rangle_{\text{cond}})} \right\rangle_k$$

The above rules breaks the annotated **while** statement in two executions: one for proving the body of the loop: $\langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (B \wedge \psi) \Rightarrow \langle \langle \cdot \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'$, and another one for proving if the invariant and the negation of the loop condition hold after the loop has been executed: $\langle \langle K \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\neg B \wedge \psi) \Rightarrow \langle \langle \cdot \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'$.

The additional operations `assume` and `implies` are used to execute patterns and verify implications between patterns, respectively. The `assume(φ)` construct, prepares the pattern given as argument, i.e. φ , for symbolic execution by setting up the initial configuration. In the case of IMP, the semantics of `assume` is the following:

$$\frac{\langle \text{assume} (\langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \langle \phi \rangle_{\text{cond}}) \rangle_{\text{cfg}} \dots \rangle_k \langle \frac{-}{\sigma} \rangle_{\text{env}} \langle \frac{-}{\phi} \rangle_{\text{cond}}}{St}$$

The \mathbb{K} tool produces by symbolic execution of φ one or more result configurations (patterns). For any such result, say φ'' , we have to check whether $\varphi'' \rightarrow \varphi'$. This is done by using the `implies(φ')` operation, which checks if the current symbolic configuration φ'' , implies the pattern given as argument φ' . The following proposition allows us to reduce the verification of $\varphi'' \rightarrow \varphi'$ to a formula which can be sent to an SMT solver. We denote by $\sigma(\mathbf{x})$ the value which is mapped to program variable \mathbf{x} by σ , and by $\text{range}(\sigma)$ the set $\{\sigma(\mathbf{x}) \mid \mathbf{x} \in \mathbf{X}\}$, where \mathbf{X} is the set of program variables.

Proposition 4.1.1 *Any pattern $\varphi \triangleq \pi \wedge \phi$, where $\pi \triangleq \langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}$ is in relation \sim (cf. Definition 3.3.1) with the pattern $(\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma})$, where $\bigwedge_{\sigma} \triangleq \bigwedge_{x \in X} x = \sigma(\mathbf{x})$.*

Proof From Lemma 2.3.1 we have $\llbracket (\exists X) \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma}) \rrbracket = \llbracket \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma}) \rrbracket$. Therefore, we have it is sufficient to show that $\llbracket \varphi \rrbracket = \llbracket \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma}) \rrbracket$.

(\subseteq) Let $\gamma \in \llbracket \varphi \rrbracket$. Then, there is a valuation $\rho : \text{Var} \rightarrow \mathcal{T}$ such that $(\gamma, \rho) \models_{\text{ML}} \pi \wedge \phi$, that is, $\rho(\pi) = \gamma$ and $\rho \models_{\text{ML}} \phi$. We consider a substitution $f_{\sigma} : X \rightarrow \text{range}(\sigma)$ which is defined as $f_{\sigma}(x) = \sigma(\mathbf{x})$, for all $x \in X$. Next, we consider a valuation $\rho' : \text{Var} \rightarrow \mathcal{T}$ such that $\rho' = \rho \circ f_{\sigma}$. Then, we have $\rho'(\langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}) = (\rho \circ f_{\sigma})(\langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}) = \rho(f_{\sigma}(\langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}})) = \rho(\langle \langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}) = \rho(\pi) = \gamma$, and $\rho'(\phi) = (\rho \circ f_{\sigma})(\phi) = \rho(f_{\sigma}(\phi)) = \rho(\phi) = \text{true}$, since ϕ does not contain variables from X . Moreover, we have $\rho'(\bigwedge_{\sigma}) = (\rho \circ f_{\sigma})(\bigwedge_{\sigma}) = \rho(f_{\sigma}(\bigwedge_{x \in X} x = \sigma(\mathbf{x}))) = \rho(\bigwedge_{x \in X} f_{\sigma}(x) = f_{\sigma}(\sigma(\mathbf{x}))) = \rho(\bigwedge_{x \in X} \sigma(\mathbf{x}) = \sigma(\mathbf{x})) = \text{true}$. So, we have ρ' such that $(\gamma, \rho') \models_{\text{ML}} \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma})$.

(\supseteq) Let $\gamma \in \llbracket \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma}) \rrbracket$. Then, there is a valuation $\rho' : \text{Var} \rightarrow \mathcal{T}$, such that $(\gamma, \rho') \models_{\text{ML}} \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi \wedge \bigwedge_{\sigma})$.

Let $\rho : \text{Var} \rightarrow \mathcal{T}$ be a valuation, such that $\rho' = \rho \circ f_\sigma$. Then, we obtain $\gamma = \rho'(\langle\langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}) = (\rho \circ f_\sigma)(\langle\langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}}) = \rho(f_\sigma(\langle\langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}})) = \rho(\langle\langle St \rangle_k \langle \sigma \rangle_{\text{env}} \rangle_{\text{cfg}}) = \rho(\pi)$, and from $\rho'(\phi) = (\rho \circ f_\sigma)(\phi) = \rho(f_\sigma(\phi)) = \rho(\phi)$ we obtain $\rho \models_{\text{ML}} \phi$. Therefore, there is ρ such that $(\gamma, \rho) \models_{\text{ML}} \pi \wedge \phi$, i.e. $\gamma \in \llbracket \varphi \rrbracket$.

■

Example 4.1.2 If $\varphi \triangleq \langle\langle \cdot \rangle_k \langle x \mapsto 4 \ n \mapsto t + 1 \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \text{true}$ is a pattern then, by Proposition 4.1.1, $\varphi \sim (\exists \{x, n\}) \langle\langle \cdot \rangle_k \langle x \mapsto x \ n \mapsto n \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \text{true} \wedge (x = 4 \wedge n = t + 1)$.

The implication $\varphi'' \rightarrow \varphi'$ can be rewritten as $\langle\langle St \rangle_k \langle \sigma'' \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \phi'' \rightarrow \langle\langle St \rangle_k \langle \sigma' \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \phi'$. Next, we use Proposition 4.1.1, and we simplify the implication using the following equivalences:

$$\begin{aligned} & (\exists X'') \langle\langle St \rangle_k \langle \sigma_{X''} \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi'' \wedge \bigwedge_{\sigma''}) \rightarrow (\exists X') \langle\langle St \rangle_k \langle \sigma_{X'} \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi' \wedge \bigwedge_{\sigma'}) \text{ iff} \\ & (\exists X) (\langle\langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi'' \wedge \bigwedge_{\sigma''}) \rightarrow \langle\langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge (\phi' \wedge \bigwedge_{\sigma'})) \text{ iff} \\ & (\exists X) ((\phi'' \wedge \bigwedge_{\sigma''}) \rightarrow (\phi' \wedge \bigwedge_{\sigma'})) \text{ iff } (\exists X) (\neg(\phi'' \wedge \bigwedge_{\sigma''}) \vee (\phi' \wedge \bigwedge_{\sigma'})) \text{ iff} \\ & ((\phi'' \wedge \bigwedge_{\sigma''}) \wedge \neg(\phi' \wedge \bigwedge_{\sigma'})) \text{ is not satisfiable.} \end{aligned}$$

We always assume that φ'' and φ' contain an empty $\langle \cdot \rangle_k$ cell. Now, we give semantics to `implies`:

$$\langle \text{implies} (\langle\langle \cdot \rangle_k \langle \sigma' \rangle_{\text{env}} \langle \psi' \rangle_{\text{cond}} \rangle_{\text{cfg}}) \langle \sigma'' \rangle_{\text{env}} \langle \phi'' \rangle_{\text{cond}} \rangle$$

$$\text{requires } \text{unsat}(\phi'' \wedge \bigwedge_{\sigma''} \wedge \neg(\phi' \wedge \bigwedge_{\sigma'}))$$

In practice, one may need additional constructs for defining \bigwedge_σ . For IMP, we have defined a function which traverses the $\langle \cdot \rangle_{\text{env}}$ cell and computes \bigwedge_σ .

In the rest of the section, we let $\mathcal{S}_{\text{IMP}'}^5$ be $\mathcal{S}_{\text{IMP}}^5$ enriched with the rules shown above. Below we define irreducible patterns and successful executions.

Definition 4.1.3 (irreducible pattern) Let φ and φ' be patterns. φ is irreducible iff there is no φ' such that $\varphi \Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s \varphi'$.

Definition 4.1.4 (complete and successful execution) A complete execution is an execution $\varphi \Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s \dots \Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s \varphi'$ such that φ' is irreducible. A complete execution is successful, and we write $\varphi \Rightarrow^! \varphi'$, if φ' contains an empty $\langle \cdot \rangle_k$ cell. We say that $\varphi \downarrow$ holds iff all complete executions starting from φ are successful.

Example 4.1.3 Let $\varphi \triangleq \langle \langle \mathbf{a} = \mathbf{a} + 1 \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0$ (the left hand side of rule (4.2)). Symbolic execution of φ is shown below:

$$\begin{aligned}
\varphi &\triangleq \langle \langle \mathbf{a} = \mathbf{a} + 1 \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 \Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(heating)} \\
\langle \langle \mathbf{a} + 1 \curvearrowright \mathbf{a} = \square \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(heating)} \\
\langle \langle \mathbf{a} \curvearrowright \square + 1 \curvearrowright \mathbf{a} = \square \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(lookup)} \\
\langle \langle a_X \curvearrowright \square + 1 \curvearrowright \mathbf{a} = \square \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(cooling)} \\
\langle \langle a_X + 1 \curvearrowright \mathbf{a} = \square \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(addition)} \\
\langle \langle a_X +_{\text{Int}} 1 \curvearrowright \mathbf{a} = \square \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(cooling)} \\
\langle \langle \mathbf{a} = a_X +_{\text{Int}} 1 \rangle_k \langle \mathbf{a} \mapsto a_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\Rightarrow_{\mathcal{S}_{\text{IMP}'}}^s && \text{(update)} \\
\langle \langle \cdot \rangle_k \langle \mathbf{a} \mapsto a_X +_{\text{Int}} 1 \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge a_X >_{\text{Int}} 0 &\triangleq \varphi' &&
\end{aligned}$$

Note that this execution is successful, i.e. $\varphi \Rightarrow^! \varphi'$ and $\varphi \downarrow$.

Next, we define the set of reachability rules derived from φ as follows:

Definition 4.1.5 Let φ be a pattern. We define the set $RL(\varphi)$ of reachability rules as follows:

$$RL(\varphi) = \{\varphi \Rightarrow \varphi' \mid \varphi \Rightarrow^! \varphi'\}$$

Example 4.1.4 Recall φ and φ' from Example 4.1.3. Then, $RL(\varphi) = \{\varphi \Rightarrow \varphi'\}$.

$RL(\varphi)$ contains all the reachability rules $\varphi \Rightarrow \varphi'$ given by successful executions $\varphi \Rightarrow^! \varphi'$.

Theorem 4.1.1 Let $\varphi \triangleq \langle \langle St \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$ be a pattern. Then any reachability rule $\varphi \Rightarrow \varphi' \in RL(\varphi)$ can be proved using [HL'-IMP].

Proof We proceed by induction on St :

- $St \triangleq \mathbf{skip}$. In this case, symbolic execution of $\langle\langle \mathbf{skip} \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$ consists in applying directly the rule for \mathbf{skip} from $\mathcal{S}_{\text{IMP}}^5$, and we obtain $\langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$. Thus, symbolic execution derives $(\exists X) \langle\langle \mathbf{skip} \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$, which corresponds to [HL'-skip].
- $St \triangleq \mathbf{a} = E$. After the symbolic execution of $\langle\langle \mathbf{a} = E \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$, the variable \mathbf{a} is mapped in the new environment to the (evaluated) expression $\text{eval}_{\sigma_X}(E)$. Moreover, $\psi_{X,Y}$ holds for the previous value of the program variable \mathbf{a} , that is $\text{eval}_{\sigma_X}(\mathbf{a})$. Therefore, by symbolic execution we derive $(\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge ((\exists u)(a_X = \text{eval}_{\sigma_X}(E[u/\mathbf{a}]) \wedge \psi_{X,Y}[u/\text{eval}_{\sigma_X}(\mathbf{a})]))$, which corresponds to [HL'-assign].
- $St \triangleq \{ S \}$. From the inductive hypothesis, symbolic execution derives $(\exists X) \langle\langle S \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}$. The symbolic execution of $\{ S \}$, consists in applying first the rule for blocks, which reduces the execution of the block to the execution of the body. Using the inductive hypothesis, we obtain $(\exists X) \langle\langle \{ St \} \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}$. In the proof system, this corresponds to [HL'-block].
- $St \triangleq St_1 ; St_2$. From the inductive hypothesis, symbolic execution derives $(\exists X) \langle\langle St_1 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}$, and $(\exists X) \langle\langle St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}$ in [HL'-IMP]. Symbolic execution of a sequence consists in executing St_1 followed by St_2 . Since we have the proofs for St_1 and St_2 , we obtain $(\exists X) \langle\langle St_1 St_2 \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi'_{X,Y}$. The corresponding proof rule is [HL'-seq].
- $St \triangleq \mathbf{while} (BE) S$. Let $\psi_{X,Y}$ be an invariant for this loop. In this case, symbolic execution consists in analysing two possible paths: either BE is *true* and the execution continues by assuming $\langle\langle S \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge BE$, or BE is *false* and the rule for $\mathbf{implies}$ checks whether the path condition implies the invariant and $\neg BE$. First, symbolic execution derives (from the inductive hypothesis) $(\exists X) \langle\langle S \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge BE \Rightarrow (\exists X) \langle\langle \cdot \rangle_k \langle \sigma_X \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge \psi_{X,Y}$.

Second, it derives $(\exists X)\langle\langle\text{while}(BE)\text{ do }St\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X)\langle\langle\cdot\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge \neg BE$ by verifying the implication. This corresponds to [HL'-while] in the proof system.

- $St \triangleq \text{if}(B) St_1 \text{ else } St_2$. Symbolic execution of the `if` statement consists in exploring the two possible branches, St_1 and St_2 , adding to the path condition the corresponding constraint, i.e. $\psi_{X,Y} \wedge B$ and $\psi_{X,Y} \wedge \neg B$, respectively. From the induction hypothesis we have $(\exists X)\langle\langle St_1 \rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge B \Rightarrow (\exists X)\langle\langle\cdot\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi'_{X,Y}$ and $(\exists X)\langle\langle St_2 \rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi_{X,Y} \wedge \neg B \Rightarrow (\exists X)\langle\langle\cdot\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi'_{X,Y}$. By applying the corresponding rules in the semantics and the inductive hypothesis, we obtain $(\exists X)\langle\langle\text{if}(B) St_1 \text{ else } St_2\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi_{X,Y} \Rightarrow (\exists X)\langle\langle\cdot\rangle_k\langle\sigma_X\rangle_{\text{env}}\rangle_{\text{cfg}} \wedge \psi'_{X,Y}$. This corresponds to [HL'-if] proof rule.

■

Essentially, Theorem 4.1.1 states that any reachability rule $\varphi \Rightarrow \varphi' \in RL(\varphi)$ obtained by (successful) symbolic execution of φ can be derived using the [HL'-IMP] proof system. Note that, in this approach, we verify a particular type of reachability rules, i.e. those which are generated from Hoare triples. For such rules, we only need symbolic execution and an SMT solver for verification.

The way we apply the proof rules from [HL'-IMP] proof system can be seen as a strategy for applying the Reachability Logic proof system on a particular language. In Section 4.2 we present a more general approach for proving RL formulas, which is based on symbolic execution and also handles circular behaviours in programs.

4.2 Reachability-Logic Verification

As presented in Section 2.4, Reachability Logic (RL) is a language independent logic for specifying program properties. For instance, on the `gcd` program shown in Fig. 4.5, the RL formula

$$\langle\text{gcd}\rangle_k \langle a \mapsto l_a \ b \mapsto l_b \ x \mapsto l_x \ y \mapsto l_y \ r \mapsto l_r \rangle_{\text{env}}$$

```

x = a;  y = b;
while (y > 0){
  r = x % y;
  x = y;
  y = r;
}

```

Figure 4.5: CinK program:gcd

$$\begin{aligned}
& \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x \ l_{\mathbf{y}} \mapsto y \ l_{\mathbf{r}} \mapsto r \rangle_{\text{store}} \sigma \wedge a \geq 0 \wedge b \geq 0 \Rightarrow \\
& (\exists x')(\exists y')(\exists r') \langle \cdot \rangle_{\mathbf{k}} \langle \mathbf{a} \mapsto l_{\mathbf{a}} \ \mathbf{b} \mapsto l_{\mathbf{b}} \ \mathbf{x} \mapsto l_{\mathbf{x}} \ \mathbf{y} \mapsto l_{\mathbf{y}} \ \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\
& \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x' \ l_{\mathbf{y}} \mapsto y' \ l_{\mathbf{r}} \mapsto r' \rangle_{\text{store}} \sigma \wedge x' = \text{gcd}(a, b),
\end{aligned} \tag{4.3}$$

where $\sigma \triangleq \langle R \rangle_{\text{return}} \langle S \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}}$, specifies that after the complete execution of the gcd program from a configuration where the program variables \mathbf{a} , \mathbf{b} are bound to non-negative values a, b , a configuration where the variable \mathbf{x} is bound to a location which holds the value $\text{gcd}(a, b)$ is reached. Here, gcd is a mathematical definition of the greatest-common-divisor ($\text{gcd}(0, 0) = 0$ by convention), and lookup is function is a standard lookup function in associative maps.

A naive attempt at verifying the RL formula (4.3) consists in symbolically executing the semantics of the CinK language with the gcd program in its left-hand side, i.e., running gcd with symbolic values $a, b \geq 0$ for \mathbf{a}, \mathbf{b} , and searching for a configuration matched by the formula's right-hand side. However, this does not succeed because it gets caught into an infinite symbolic execution, induced by the infinitely many iterations of the loop.

The proof system of RL (Figure 2.1) is compact and elegant but, despite its nice theoretical properties, its use in practice on nontrivial programs is difficult, because it gives the user a lot of freedom regarding the order and manner of rule application, and offers no practical guidelines for constructing proofs. Moreover, it is not designed for disproving formulas: since the system is relatively complete, the only way to disprove a formula is to show that there exists no proof-tree for the formula (in the presence of an oracle able to decide the validity of first-order assertions), which is not practically possible. On the other hand, the proof system of RL is language independent, unlike in Hoare Logic, where the

proof system depends explicitly on the language. By its nature, verification based on Hoare Logic is a bit limited. First, it is not easy to express with Hoare triples properties for programs which involve recursive function calls, blocks with local variables, virtual method calls, etc. Second, since Hoare Logic is language dependent, its soundness has to be proved for every language, which is not the case of Reachability Logic.

In this section we show how symbolic execution can be used for verifying RL formulas. The approach consists in a simpler proof system, where symbolic execution is a main ingredient and is used as a rule during proof construction. We show that a certain strategy for executing the proof system is sound (when it terminates successfully on a given input - set of RL formulas - then the formulas are valid) and is also weakly complete (if it terminates in failure then its input is invalid). The strategy is not relatively complete, since it may not terminate even for valid RL formulas, as illustrated by the naive symbolic execution attempt at proving (4.3). In order to terminate, it requires additional information under the form of RL formulas. Together, these properties say that when it terminates, our approach correctly solves the RL-based program-verification problem. Termination, of course, cannot be guaranteed because the RL verification problem is undecidable, but the soundness and weak completeness results say that any inability to prove/disprove formulas is only due to issues inherent to the RL verification problem, and not to the particular approach we propose for solving it.

The soundness and weak completeness results are based on certain mutual-simulation properties relating symbolic and concrete program execution, and on a so-called *circularity principle* for reachability-logic formulas, which specifies the conditions under which goals can be reused as hypotheses in proofs. This is essential for proving programs with infinite state-spaces induced e.g., by an unbounded (symbolic) number of loop iterations or of recursive calls. Soundness also requires that the semantics of the programming language is *total*; the behaviour of instructions is completely specified, and weak completeness additionally requires that the semantics is *confluent*: any two executions of a program eventually reach the same state. Weak completeness also poses certain additional requirements on the RL goals; none of these requirements is hard to meet.

In order to prove that a semantics \mathcal{S} satisfies a set G of RL formulas called *goals*, i.e., for proving $\mathcal{S} \models_{\text{RL}} G$ (i.e., $\mathcal{S} \models_{\text{RL}} g$ for all $g \in G$) one can reuse the goals G as hypotheses, provided that the set of goals to be proved is replaced by another set of goals $\Delta_{\mathcal{S}}(G)$, obtained by symbolic execution from the set G . This goal-as-hypothesis reuse is essential for proving programs with unbounded execution lengths induced by loops of recursive function calls. Thus, from $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(G)$ one deduces $\mathcal{S} \models_{\text{RL}} G$; we call this implication the *circularity principle* for RL. Here, \Vdash denotes the entailment of a certain proof system, presented below, in which symbolic execution also plays a major role.

The following definition of *derivative* is an essential operation in RL verification approach. It computes (up to the equivalence relation \sim) a disjunction of all the successors of a pattern by the symbolic transition relation. It uses the choice operation ε , which chooses an arbitrary element in a nonempty set.

Definition 4.2.1 (Derivative) *The derivative $\Delta_{\mathcal{S}}(\varphi)$ of a pattern φ for a set \mathcal{S} of rules is $\Delta_{\mathcal{S}}(\varphi) \triangleq \bigvee_{[\varphi] \sim \Rightarrow_{\mathcal{S}}^* [\varphi'] \sim} \varepsilon([\varphi'] \sim)$.*

Remark 4.2.1 *Since the symbolic transition relation is finitely branching (Remark 3.3.6), for finite rule sets \mathcal{S} the derivative is a finite disjunction. Note also that the patterns in the derivative are only defined up to the equivalence relation \sim .*

Definition 4.2.2 (Derivable Pattern) *An elementary pattern $\varphi \triangleq \pi \wedge \phi$ is derivable for \mathcal{S} if $\Delta_{\mathcal{S}}(\varphi)$ is a nonempty disjunction.*

The notion of cover, defined below, is essential for the soundness of RL-formula verification by symbolic execution, in particular, in situations where a proof goal is circularly used as a hypothesis. Such goals can only be used in symbolic execution only when they *cover* the pattern being symbolically executed:

Definition 4.2.3 (Cover) *Consider an elementary pattern $\varphi \triangleq \pi \wedge \phi$. A set of rules \mathcal{S}' satisfying $\models_{\text{FOL}} \phi \rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}'} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$, where $\sigma_{\pi}^{\pi_1}$ is the most general unifier of π and π_1 , is a cover of φ .*

Remark 4.2.2 *The existence of the most-general unifier $\sigma_\pi^{\pi_1}$ in the above definition means the basic patterns in the LHS of rules in \mathcal{S}' are unifiable with the basic pattern π . In particular, $\bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}'} \sigma_\pi^{\pi_1}(\phi_1 \wedge \phi_2)$ is a nonempty disjunction, otherwise the validity in Def. 4.2.3 would not hold (an empty disjunction is equivalent to false).*

The notion of *total semantics* is essential for the soundness of our approach.

Definition 4.2.4 (Total Semantics) *We say that a set \mathcal{S} of semantical rules is total if for all patterns φ derivable for \mathcal{S} , \mathcal{S} is a cover for φ .*

Remark 4.2.3 *The semantics of CinK is not total because the rule for division:*

$$\langle I_1 / I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \wedge I_2 \neq_{\text{Int}} 0 \Rightarrow \langle I_1 / \text{Int} I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle S \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}}$$

does not meet the condition of Definition 4.2.4, since the semantics of CinK does not cover the patterns of the form:

$$\langle I_1 / I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}}$$

which are derivable. This is due to the condition $I_2 \neq 0$, which is not logically valid. The semantics can easily be made total by adding a rule

$$\langle I_1 / I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \wedge I_2 =_{\text{Int}} 0 \Rightarrow \langle \text{error} \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \quad (4.4)$$

that leads divisions by zero into “error” configurations. We assume hereafter that the CinK semantics has been transformed into a total one by adding the above rule.

Using the notion of cover we obtain a derived rule of the RL proof system:

Lemma 4.2.1 *If $\mathcal{S}' \subseteq \mathcal{S}$ is a cover for φ , and G is a (possibly empty) set of RL formulas, then $\mathcal{S} \vdash_G \varphi \Rightarrow \Delta_{\mathcal{S}'}(\varphi)$.*

Proof Let $\varphi \triangleq \pi \wedge \phi$. According to Definition 4.2.1, we have $\Delta_{\mathcal{S}'}(\varphi) \triangleq \bigvee_{[\varphi]_{\sim} \Rightarrow_{\mathcal{S}'}^{\varepsilon} [\varphi']_{\sim}} \varepsilon([\varphi']_{\sim})$. Since \mathcal{S}' is a cover for φ , by Remark 4.2.2, $\Delta_{\mathcal{S}'}(\varphi)$ is a nonempty disjunction. Using Definition 3.3.2 we obtain that $[\varphi']_{\sim} = [\sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2)]_{\sim}$ for some $\alpha \triangleq (\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2) \in \mathcal{S}$ and $\sigma_{\pi}^{\pi_1}$ the most general symbolic unifier of π, π_1 as built in the proof of Lemma 3.2.1. The projection of $\sigma_{\pi}^{\pi_1}$ on $\text{vars}(\pi)$ is the identity, and the projection on $\text{vars}(\pi_1)$ is a substitution of $\text{vars}(\pi_1)$ matching π_1 on π . By a variable renaming we can always assume that $\text{vars}(\phi) \cap \text{vars}(\pi_1) = \emptyset$, which means that the effect of $\sigma_{\pi}^{\pi_1}$ on ϕ is the identity as well, i.e., $\sigma_{\pi}^{\pi_1}(\phi) = \phi$.

Using the above characterisation for the patterns φ' , we obtain that the choices of the ε operation in $\Delta_{\mathcal{S}'} \triangleq \bigvee_{[\varphi]_{\sim} \Rightarrow_{\mathcal{S}'}^{\varepsilon} [\varphi']_{\sim}} \varepsilon([\varphi']_{\sim})$ can be made such that

$$\Delta_{\mathcal{S}'}(\varphi) = \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}'} \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi \wedge \phi_1 \wedge \phi_2) \quad (4.5)$$

On the other hand, by using the RL proof system from Figure 2.1 and the derived rules of RL (shown in Section 2.4):

- first apply *Substitution* with the rule $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2$ and substitution $\sigma_{\pi}^{\pi_1}$:

$$\mathcal{S} \vdash_G \sigma_{\pi}^{\pi_1}(\pi_1 \wedge \phi_1) \Rightarrow \sigma_{\pi}^{\pi_1}(\pi_2 \wedge \phi_2)$$

- apply *LogicalFraming* with the patternless formula $\sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_2)$:

$$\mathcal{S} \vdash_G \sigma_{\pi}^{\pi_1}(\pi_1 \wedge \phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_2) \Rightarrow \sigma_{\pi}^{\pi_1}(\pi_2 \wedge \phi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_2)$$

- using the [Consequence] and the fact that FOL patternless formulas distribute over patterns:

$$\mathcal{S} \vdash_G \sigma_{\pi}^{\pi_1}(\pi) \wedge (\sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2)) \Rightarrow \sigma_{\pi}^{\pi_1}(\pi_2) \wedge (\sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2))$$

Since the effect of $\sigma_{\pi}^{\pi_1}$ on both π and ϕ is the identity, we further obtain:

$$\mathcal{S} \vdash_G \pi \wedge (\phi \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2)) \Rightarrow \sigma_{\pi}^{\pi_1}(\pi_2) \wedge (\sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2))$$

- then, using **CaseAnalysis** and **Consequence** several times we obtain:

$$\mathcal{S} \vdash_G \pi \wedge \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}'} (\phi \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2)) \Rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}'} (\sigma_{\pi}^{\pi_1}(\pi_2) \wedge (\sigma_{\pi}^{\pi_1}(\phi) \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2))) \quad (4.6)$$

We know from (4.5) that the right-hand side of (4.6) is $\Delta_{\mathcal{S}'}(\varphi)$. To prove $\mathcal{S} \vdash_G \varphi \Rightarrow \Delta_{\mathcal{S}'}(\varphi)$ there only remains to prove (\diamond): the condition in the left-hand side: $\bigvee_{(\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2) \in \mathcal{S}'} (\phi \wedge \sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2))$ is logically implied by ϕ in FOL. Since \mathcal{S}' is a cover for φ , we obtain, using Definition 4.2.3, the validity of $\phi \rightarrow \bigvee_{(\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2) \in \mathcal{S}'} (\sigma_{\pi}^{\pi_1}(\phi_1) \wedge \sigma_{\pi}^{\pi_1}(\phi_2))$, which proves (\diamond) and concludes the proof of the lemma. ■

Corollary 4.2.1 *If \mathcal{S} is total and φ is derivable for \mathcal{S} , then $\mathcal{S} \vdash \varphi \Rightarrow \Delta_{\mathcal{S}}(\varphi)$.*

Proof Apply Lemma 4.2.1 with an empty set G . ■

Before we introduce our proof system we need to deal with the issue that operational semantics are not always weakly well-defined as required by the RL original deductive system's soundness. For example, the semantics of CinK is not weakly well-defined due to the rules for division and modulo, which, for valuations ρ mapping divisors to 0, have no instance of their right-hand side. However, due to the introduction of the rule (4.4) in order to make the semantics total (cf. Remark 4.2.3), the semantics of division can now be equivalently rewritten using just one (reachability-logic) disjunctive rule:

$$\begin{aligned} \langle I_1 / I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \Rightarrow \\ \langle \langle I_1 /_{\text{Int}} I_2 \curvearrowright C \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \wedge I_2 \neq_{\text{Int}} 0 \rangle \\ \vee \langle \langle \text{error} \rangle_k \langle E \rangle_{\text{env}} \langle S \rangle_{\text{store}} \langle R \rangle_{\text{return}} \langle St \rangle_{\text{stack}} \langle I \rangle_{\text{in}} \langle O \rangle_{\text{out}} \wedge I_2 =_{\text{Int}} 0 \rangle \end{aligned} \quad (4.7)$$

By using this rule instead of the two original ones, and by applying the same transformation for the rules defining modulo, the semantics becomes weakly well-defined. This transformation is formalised as follows.

Definition 4.2.5 (\mathcal{S}^Δ) Given a set of semantical rules \mathcal{S} , the set of semantical rules \mathcal{S}^Δ is defined by $\mathcal{S}^\Delta \triangleq \{\pi \Rightarrow \Delta_{\mathcal{S}}(\pi) \mid (\pi \wedge \phi \Rightarrow \varphi) \in \mathcal{S}\}$.

The following lemma establishes that \mathcal{S}^Δ has both properties (totality and weak well definedness) required for the soundness of our approach. For this, we need to extend the notion of derivative for rules of the form \mathcal{S}^Δ (containing disjunctions in their right-hand side) by letting $\Delta_{\mathcal{S}^\Delta}(\varphi) \triangleq \Delta(\varphi)$. The derivability of a pattern for \mathcal{S}^Δ is also, by definition, the derivability for \mathcal{S} . The notion of cover is extended for such rules, of the form $\pi_1 \wedge \phi_1 \Rightarrow \bigvee_{j \in J} \pi_2^j \wedge \phi_2^j$ by letting $\varphi \triangleq \pi \wedge \phi$ be covered by a set \mathcal{S}' of such rules if $\models_{\text{ML}} \varphi \rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \bigvee_{j \in J} \pi_2^j \wedge \phi_2^j \in \mathcal{S}'} \bigvee_{j \in J} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2^j)$.

Lemma 4.2.2 If \mathcal{S} is total then \mathcal{S}^Δ is total and weakly well-defined.

Proof We start by proving the *totality* of \mathcal{S}^Δ . If \mathcal{S} is total then, by definition, \mathcal{S} covers all patterns φ derivable for \mathcal{S} . We need to prove that for all φ that is derivable for \mathcal{S}^Δ , \mathcal{S}^Δ is a cover for φ . Assume such a $\varphi \triangleq \pi \wedge \phi$ derivable for \mathcal{S}^Δ . We have defined derivability for \mathcal{S}^Δ as derivability for \mathcal{S} , hence φ is derivable for \mathcal{S} . The totality of \mathcal{S} and the definition of cover then ensures $\models_{\text{FOL}} \phi \rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$ (

Now, the rules in \mathcal{S}^Δ are of the form $\pi \Rightarrow \Delta_{\mathcal{S}}(\pi)$, for some π occurring in the pattern LHS of a rule in \mathcal{S} . Up to FOL conditions we have $\Delta_{\mathcal{S}}(\pi) = \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\pi_2 \wedge \phi_1 \wedge \phi_2)$ (implicitly, the disjunction performed over those rule in \mathcal{S} for which $\sigma_{\pi}^{\pi_1}$ exists). Hence, by the notion of cover applied to \mathcal{S}^Δ , in order to show that \mathcal{S}^Δ is total, we need to show

$$\models_{\text{FOL}} \phi \rightarrow \bigvee_{\left(\pi \Rightarrow \bigvee_{(\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2) \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\pi_2 \wedge \phi_1 \wedge \phi_2) \right) \in \mathcal{S}^\Delta} \left(\bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2) \right) \quad (4.8)$$

where the second disjunction $\bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$ is merely repeated by the first disjunction $\bigvee_{(\dots) \in \mathcal{S}^\Delta}$ as many times as rules in \mathcal{S}^Δ . Hence, (4.8) simplifies to $\models_{\text{FOL}} \phi \rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$, which we have obtained above ().

For *weak well-definedness* of \mathcal{S}^Δ we need to show that for $\pi \Rightarrow \Delta_{\mathcal{S}}(\pi) \in \mathcal{S}^\Delta$ and each valuation ρ , there exists a configuration γ such that $(\gamma, \rho) \models_{\text{ML}}$

$\Delta_{\mathcal{S}}(\pi)$. We have obtained above that $\Delta_{\mathcal{S}}(\pi) = \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\pi_2 \wedge (\phi_1 \wedge \phi_2))$ (where implicitly the disjunction is performed over those rules in \mathcal{S} for which $\sigma_{\pi}^{\pi_1}$ exists). Now, \mathcal{S} is total, and $\pi (= \pi \wedge \text{true})$ is derivable for \mathcal{S} . This implies $\models_{\text{FOL}} \text{true} \rightarrow \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$, i.e., $\models_{\text{FOL}} \bigvee_{\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}} \sigma_{\pi}^{\pi_1}(\phi_1 \wedge \phi_2)$. This means that for any ρ there exists in the above disjunction at least one disjunct, say, $\sigma_{\pi}^{\pi_1}(\phi_1^i \wedge \phi_2^i)$, such that $\rho \models_{\text{FOL}} \sigma_{\pi}^{\pi_1}(\phi_1^i \wedge \phi_2^i)$. By taking the corresponding $\gamma \triangleq \rho(\sigma_{\pi}^{\pi_1}(\pi_2))$ we obtain $(\gamma, \rho) \models_{\text{ML}} \sigma_{\pi}^{\pi_1}(\pi_2) \wedge \sigma_{\pi}^{\pi_1}(\phi_1^i \wedge \phi_2^i)$, i.e., (γ, ρ) models one of the disjuncts of $\Delta_{\mathcal{S}}(\pi)$. This implies $(\gamma, \rho) \models_{\text{ML}} \Delta_{\mathcal{S}}(\pi)$, which completes the proof. \blacksquare

Another property useful in the sequel is stated below.

Lemma 4.2.3 $\mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$ iff $\mathcal{S}^{\Delta} \models_{\text{RL}} \varphi \Rightarrow \varphi'$.

Proof (\Rightarrow) From $\mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$ (cf. Definition 2.4.5) we obtain that for each terminating configuration $\gamma \in \mathcal{T}_{\text{Cfg}}$ and valuation $\rho : \text{Var} \rightarrow \mathcal{T}$ such that $(\gamma, \rho) \models_{\text{ML}} \varphi$, there is $\gamma' \in \mathcal{T}_{\text{Cfg}}$ and a path from γ to γ' , denoted $\gamma \xrightarrow{\alpha^*}_{\mathcal{S}} \gamma'$, with $\alpha^* \in \mathcal{S}^*$, such that $(\gamma', \rho) \models \varphi'$. We prove by induction on the length of the sequence α^* that $\gamma \xrightarrow{\alpha^*}_{\mathcal{S}^{\Delta}} \gamma'$, which implies the (\Rightarrow) direction.

The base case (i.e. the length of the path is 0) is trivial. For the inductive step, we use the fact that each transition $\gamma_1 \xrightarrow{\alpha}_{\mathcal{S}} \gamma_2$ is induced by semantical rule $\alpha \triangleq \pi \wedge \phi \Rightarrow \varphi \in \mathcal{S}$ and a valuation $\rho : \text{Var} \rightarrow \mathcal{T}$. Then, the corresponding rule $\pi \Rightarrow \Delta_{\mathcal{S}}(\pi) \in \mathcal{S}^{\Delta}$ induces a corresponding transition $\gamma_1 \xrightarrow{\alpha}_{\mathcal{S}^{\Delta}} \gamma_2$ with the same valuation ρ .

(\Leftarrow) In this case the proof is almost similar with the previous one, the only difference being in the inductive step where each transition $\gamma_1 \xrightarrow{\alpha}_{\mathcal{S}^{\Delta}} \gamma_2$ is induced by a semantical rule $\alpha^{\Delta} \triangleq \pi \Rightarrow \Delta_{\mathcal{S}}(\pi) \in \mathcal{S}^{\Delta}$ with a valuation $\rho : \text{Var} \rightarrow \mathcal{T}$. Then, the rule $\alpha \triangleq \pi \wedge \phi \Rightarrow \varphi \in \mathcal{S}$ induces a corresponding transition $\gamma_1 \xrightarrow{\alpha}_{\mathcal{S}} \gamma_2$ with the same valuation ρ . \blacksquare

Lemma 4.2.4 A pattern φ is derivable for \mathcal{S} iff φ is derivable for \mathcal{S}^{Δ} .

Proof By Definition 4.2.2, $\varphi \triangleq \pi \wedge \phi$ derivable for \mathcal{S} means that $\Delta_{\mathcal{S}}(\varphi)$ is not the empty disjunction. This holds if and only if there exists a

$$\begin{array}{l}
\text{[SymbolicStep]} \quad \frac{\varphi \text{ derivable for } \mathcal{S}}{\mathcal{S} \cup G \Vdash \varphi \Rightarrow \Delta_{\mathcal{S}}(\varphi)} \\
\text{[CircHypothesis]} \quad \frac{\alpha \in G \quad \alpha \text{ covers } \varphi}{\mathcal{S} \cup G \Vdash \varphi \Rightarrow \Delta_{\{\alpha\}}(\varphi)} \\
\text{[Implication]} \quad \frac{\models_{\text{ML}} \varphi \rightarrow \varphi'}{\mathcal{S} \cup G \Vdash \varphi \Rightarrow \varphi'} \\
\text{[CaseAnalysis]} \quad \frac{\mathcal{S} \cup G \Vdash \varphi_1 \Rightarrow \varphi \quad \mathcal{S} \cup G \Vdash \varphi_2 \Rightarrow \varphi}{\mathcal{S} \cup G \Vdash (\varphi_1 \vee \varphi_2) \Rightarrow \varphi} \\
\text{[Transitivity]} \quad \frac{\mathcal{S} \cup G \Vdash \varphi \Rightarrow \varphi'' \quad \mathcal{S} \cup G \Vdash \varphi'' \Rightarrow \varphi'}{\mathcal{S} \cup G \Vdash \varphi \Rightarrow \varphi'}
\end{array}$$

Figure 4.6: Proof System for $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(G)$.

rule $\pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in \mathcal{S}$ such that π_1 matches π . The proof is finished by noting that π_1 is also the left-hand side of the rule corresponding $\pi_1 \Rightarrow \Delta_{\mathcal{S}}(\pi_1) \in \mathcal{S}^{\Delta}$.

■

We now have almost all the ingredients for proving RL formulas by symbolic execution. In the following we assume a language with a semantics \mathcal{S} , and a finite of RL formulas with elementary patterns in their left-hand sides $G = \{\varphi_i \Rightarrow \varphi'_i \mid i = 1, \dots, n\}$.

We say that a RL formula $\varphi \Rightarrow \varphi'$ is *derivable* for \mathcal{S} if φ is derivable for \mathcal{S} . If G is a set of RL formulas then $\Delta_{\mathcal{S}}(G)$ is the set $\{\Delta_{\mathcal{S}}(\varphi) \Rightarrow \varphi' \mid \varphi \Rightarrow \varphi' \in G\}$, $\mathcal{S} \Vdash G$ denotes the conjunction $\bigwedge_{\varphi \Rightarrow \varphi' \in G} \mathcal{S} \Vdash \varphi \Rightarrow \varphi'$, and $\mathcal{S} \models_{\text{RL}} G$ denotes $\bigwedge_{\varphi \Rightarrow \varphi' \in G} \mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$. The proof system \Vdash is shown in Figure 4.6. The following theorem establishes the soundness of our approach. It is based on a *circularity principle* also encountered in other coinductive frameworks, e.g., [95].

Theorem 4.2.1 (Circularity Principle) *If \mathcal{S} is total and weakly well-defined, and G is derivable for \mathcal{S} , then $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(G)$ implies $\mathcal{S} \models_{\text{RL}} G$.*

Proof For all $i = 1, \dots, n$ we apply the [Transitivity] rule of the original

RL proof system (Figure 2.1), with $\varphi_i'' \triangleq \Delta_{\mathcal{S}}(\varphi_i)$, and we obtain:

$$\frac{\mathcal{S} \vdash_G \varphi_i \Rightarrow \Delta_{\mathcal{S}}(\varphi_i) \quad (\mathcal{S} \cup G) \vdash \Delta_{\mathcal{S}}(\varphi_i) \Rightarrow \varphi_i'}{\mathcal{S} \vdash_G \varphi_i \Rightarrow \varphi_i'}$$

The first hypothesis: $\mathcal{S} \vdash_G \varphi_i \Rightarrow \Delta_{\mathcal{S}}(\varphi_i)$ holds thanks to Lemma 4.2.1. The second one, $(\mathcal{S} \cup G) \vdash \Delta_{\mathcal{S}}(\varphi_i) \Rightarrow \varphi_i'$ holds because all the rules of \Vdash are derived rules of \vdash , thanks to Lemma 4.2.1 again. Hence, we obtain $\mathcal{S} \vdash_G \varphi_i \Rightarrow \varphi_i'$ for all $i = 1, \dots, n$, i.e., $\mathcal{S} \vdash_G G$. Then we obtain $\mathcal{S} \vdash G$ by applying the derived rule *Set Circularity* of RL. Finally, the soundness of \vdash (with the hypothesis that \mathcal{S} is weakly well defined) implies $\mathcal{S} \models_{\text{RL}} G$. ■

If a semantics \mathcal{S} is only total, but not weakly well-defined, one can use Theorem 4.2.1 to prove $\mathcal{S}^\Delta \models_{\text{RL}} G$ instead of $\mathcal{S} \models_{\text{RL}} G$. Indeed, \mathcal{S}^Δ is both total and weakly well-defined (thanks to Lemma 4.2.2) thus it satisfies the hypotheses of Theorem 4.2.1. Lemma 4.2.3 then ensures that by establishing $\mathcal{S}^\Delta \models_{\text{RL}} G$ we also proved $\mathcal{S} \models_{\text{RL}} G$, which is what we wanted in the first place.

Example 4.2.1 *In this example we show how the RL formula (4.3) is proved using the proof system shown in Figure 4.6, which amounts to verifying that the gcd program meets its specification. Instead of (4.3) we write $\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$, where:*

$$\begin{aligned} \varphi^{\text{gcd}} \triangleq & \langle \text{gcd} \rangle_{\text{k}} \langle \mathbf{a} \mapsto l_{\mathbf{a}} \mathbf{b} \mapsto l_{\mathbf{b}} \mathbf{x} \mapsto l_{\mathbf{x}} \mathbf{y} \mapsto l_{\mathbf{y}} \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x \ l_{\mathbf{y}} \mapsto y \ l_{\mathbf{r}} \mapsto r \rangle_{\text{store}} \sigma \wedge a \geq 0 \wedge b \geq 0, \end{aligned}$$

and

$$\begin{aligned} \varphi'^{\text{gcd}} \triangleq & (\exists x') (\exists y') (\exists r') \langle \cdot \rangle_{\text{k}} \langle \mathbf{a} \mapsto l_{\mathbf{a}} \mathbf{b} \mapsto l_{\mathbf{b}} \mathbf{x} \mapsto l_{\mathbf{x}} \mathbf{y} \mapsto l_{\mathbf{y}} \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x' \ l_{\mathbf{y}} \mapsto y' \ l_{\mathbf{r}} \mapsto r' \rangle_{\text{store}} \sigma \wedge x' = \text{gcd}(a, b). \end{aligned}$$

Note that we use σ to denote the rest of the configuration, where each cell contain a variable. The pattern φ^{gcd} contains a loop, and in order for the proof to succeed, we need an additional RL formula as loop invariant. Thus, we consider the RL formula $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$, where:

$$\begin{aligned} \varphi^{\text{while}} \triangleq & \langle \text{while } (\mathbf{y} > 0) \{ \mathbf{r} = \mathbf{x} \% \mathbf{y}; \mathbf{x} = \mathbf{y}; \mathbf{y} = \mathbf{r}; \} \rangle_{\text{k}} \\ & \langle \mathbf{a} \mapsto l_{\mathbf{a}} \mathbf{b} \mapsto l_{\mathbf{b}} \mathbf{x} \mapsto l_{\mathbf{x}} \mathbf{y} \mapsto l_{\mathbf{y}} \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x \ l_{\mathbf{y}} \mapsto y \ l_{\mathbf{r}} \mapsto r \rangle_{\text{store}} \sigma \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x, y) \wedge x \geq 0 \wedge y \geq 0, \end{aligned}$$

and

$$\begin{aligned} \varphi'^{\text{while}} \triangleq & (\exists x')(\exists y')(\exists r') \langle \cdot \rangle_k \langle \mathbf{a} \mapsto l_{\mathbf{a}} \mathbf{b} \mapsto l_{\mathbf{b}} \mathbf{x} \mapsto l_{\mathbf{x}} \mathbf{y} \mapsto l_{\mathbf{y}} \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x' \ l_{\mathbf{y}} \mapsto y' \ l_{\mathbf{r}} \mapsto r' \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x', y') \wedge x' \geq 0 \wedge y' \geq 0. \end{aligned}$$

The rule $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$ says that the **while** loop preserves the invariant: the gcd of the values corresponding to program variables \mathbf{a} , \mathbf{b} equals the gcd of the values of variables \mathbf{x} , \mathbf{y} . At this point, the set G includes the rules $\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$ and $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$.

Now we can prove $\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$ by applying the deductive system \Vdash to the set of goals G as follows:

- We apply [SymbolicStep] a finite number of times until we obtain $\mathcal{S} \cup G \Vdash \varphi^{\text{gcd}} \Rightarrow \varphi_1^{\text{gcd}}$, where

$$\begin{aligned} \varphi_1^{\text{gcd}} \triangleq & \langle \text{while } (y > 0) \{ \mathbf{r} = \mathbf{x} \% \mathbf{y}; \mathbf{x} = \mathbf{y}; \mathbf{y} = \mathbf{r}; \} \rangle_k \\ & \langle \mathbf{a} \mapsto l_{\mathbf{a}} \ \mathbf{b} \mapsto l_{\mathbf{b}} \ \mathbf{x} \mapsto l_{\mathbf{x}} \ \mathbf{y} \mapsto l_{\mathbf{y}} \ \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto a \ l_{\mathbf{y}} \mapsto b \ l_{\mathbf{r}} \mapsto r \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x, y) \end{aligned}$$

- Now we can apply [CircHypothesis] with $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}} \in G$, and we obtain $\mathcal{S} \cup G \Vdash \varphi_1^{\text{gcd}} \Rightarrow \varphi_2^{\text{gcd}}$, where:

$$\begin{aligned} \varphi_2^{\text{gcd}} \triangleq & (\exists x')(\exists y')(\exists r') \langle \cdot \rangle_k \\ & \langle \mathbf{a} \mapsto l_{\mathbf{a}} \ \mathbf{b} \mapsto l_{\mathbf{b}} \ \mathbf{x} \mapsto l_{\mathbf{x}} \ \mathbf{y} \mapsto l_{\mathbf{y}} \ \mathbf{r} \mapsto l_{\mathbf{r}} \rangle_{\text{env}} \\ & \langle l_{\mathbf{a}} \mapsto a \ l_{\mathbf{b}} \mapsto b \ l_{\mathbf{x}} \mapsto x' \ l_{\mathbf{y}} \mapsto y' \ l_{\mathbf{r}} \mapsto r' \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x', y') \wedge x' \geq 0 \wedge y' \geq 0 \end{aligned}$$

- Next, since φ_2^{gcd} and φ'^{while} are essentially the same pattern, then $\Vdash_{\text{ML}} \varphi_2^{\text{gcd}} \rightarrow \varphi'^{\text{while}}$. So, we apply [Implication] and we obtain $\mathcal{S} \cup G \Vdash \varphi_2^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$.

A graphical representation of the proof tree for the formula (4.3) is shown below ($\varphi \xrightarrow{\text{Rule}} \varphi'$ represents $\frac{\dots}{\varphi \Rightarrow \varphi'} [\text{Rule}]$):

$$\varphi^{\text{gcd}} \xrightarrow{\text{SymbolicStep}^+} \varphi_1^{\text{gcd}} \xrightarrow{\text{CircHypothesis}} \varphi_2^{\text{gcd}} \xrightarrow{\text{Implication}} \varphi'^{\text{gcd}}$$

Finally, we apply [Transitivity] twice: first, we obtain $\mathcal{S} \cup G \Vdash \varphi^{\text{gcd}} \Rightarrow \varphi_2^{\text{gcd}}$ from $\mathcal{S} \cup G \Vdash \varphi^{\text{gcd}} \Rightarrow \varphi_1^{\text{gcd}}$ and $\mathcal{S} \cup G \Vdash \varphi_1^{\text{gcd}} \Rightarrow \varphi_2^{\text{gcd}}$, and then we obtain $\mathcal{S} \cup G \Vdash \varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$ from $\mathcal{S} \cup G \Vdash \varphi^{\text{gcd}} \Rightarrow \varphi_2^{\text{gcd}}$ and $\mathcal{S} \cup G \Vdash \varphi_2^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$.

Below we show the proof of the invariant rule $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$:

- By applying [SymbolicStep] several times and [CaseAnalysis] we obtain either $\mathcal{S} \cup G \Vdash \varphi^{\text{while}} \Rightarrow \varphi_1^{\text{while}}$, where

$$\begin{aligned} \varphi_1^{\text{while}} \triangleq & \langle \cdot \rangle_k \langle \mathbf{a} \mapsto l_a \mathbf{b} \mapsto l_b \mathbf{x} \mapsto l_x \mathbf{y} \mapsto l_y \mathbf{r} \mapsto l_r \rangle_{\text{env}} \\ & \langle l_a \mapsto a \ l_b \mapsto b \ l_x \mapsto x \ l_y \mapsto y \ l_r \mapsto r \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x, y) \wedge x \geq_{\text{Int}} 0 \wedge y =_{\text{Int}} 0 \end{aligned}$$

or $\mathcal{S} \cup G \Vdash \varphi^{\text{while}} \Rightarrow \varphi_2^{\text{while}}$, where

$$\begin{aligned} \varphi_2^{\text{while}} \triangleq & \{ \mathbf{r} = \mathbf{x} \% \mathbf{y}; \mathbf{x} = \mathbf{y}; \mathbf{y} = \mathbf{r}; \} \curvearrowright \text{while}(y > 0) \{ \dots \} _k \\ & \langle \mathbf{a} \mapsto l_a \mathbf{b} \mapsto l_b \mathbf{x} \mapsto l_x \mathbf{y} \mapsto l_y \mathbf{r} \mapsto l_r \rangle_{\text{env}} \\ & \langle l_a \mapsto a \ l_b \mapsto b \ l_x \mapsto x \ l_y \mapsto y \ l_r \mapsto r \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x, y) \wedge x \geq_{\text{Int}} 0 \wedge y >_{\text{Int}} 0. \end{aligned}$$

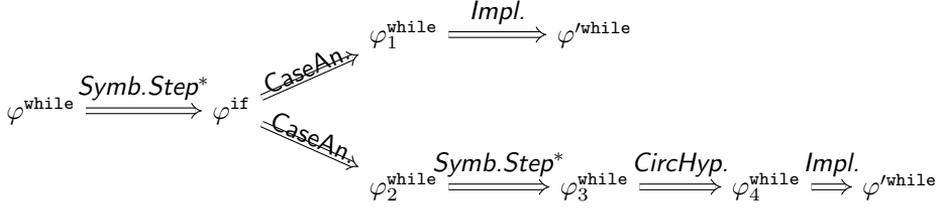
- We observe that φ_1^{while} implies φ'^{while} (because there are $x' = x$, $y' = 0$, and $r' = r$ such that φ'^{while} holds). Thus, we first apply [Implication] to obtain $\mathcal{S} \cup G \Vdash \varphi_1^{\text{while}} \Rightarrow \varphi'^{\text{while}}$, and then, since $\mathcal{S} \cup G \Vdash \varphi^{\text{while}} \Rightarrow \varphi_1^{\text{while}}$, we have $\mathcal{S} \cup G \Vdash \varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$ by [Transitivity].
- On the other hand, we can apply [SymbolicStep] over φ_2^{while} until we obtain $\mathcal{S} \cup G \Vdash \varphi_2^{\text{while}} \Rightarrow \varphi_3^{\text{while}}$, where

$$\begin{aligned} \varphi_3^{\text{while}} \triangleq & \langle \text{while}(y > 0) \{ \mathbf{r} = \mathbf{x} \% \mathbf{y}; \mathbf{x} = \mathbf{y}; \mathbf{y} = \mathbf{r}; \} \rangle_k \\ & \langle \mathbf{a} \mapsto l_a \mathbf{b} \mapsto l_b \mathbf{x} \mapsto l_x \mathbf{y} \mapsto l_y \mathbf{r} \mapsto l_r \rangle_{\text{env}} \\ & \langle l_a \mapsto a \ l_b \mapsto b \ l_x \mapsto y \ l_y \mapsto r \ l_r \mapsto x \% y \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x, y) \wedge x \geq_{\text{Int}} 0 \wedge y >_{\text{Int}} 0 \end{aligned}$$

- Now, since $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$ covers φ_3^{while} we apply [CirCHypothesis] and we get $\mathcal{S} \cup G \Vdash \varphi_3^{\text{while}} \Rightarrow \varphi_4^{\text{while}}$, where the pattern φ_4^{while} is

$$\begin{aligned} \varphi_4^{\text{while}} \triangleq & (\exists x') (\exists y') (\exists r') \langle \cdot \rangle_k \langle \mathbf{a} \mapsto l_a \mathbf{b} \mapsto l_b \mathbf{x} \mapsto l_x \mathbf{y} \mapsto l_y \mathbf{r} \mapsto l_r \rangle_{\text{env}} \\ & \langle l_a \mapsto a \ l_b \mapsto b \ l_x \mapsto x' \ l_y \mapsto y' \ l_r \mapsto r' \rangle_{\text{store } \sigma} \\ & \wedge \text{gcd}(a, b) =_{\text{Int}} \text{gcd}(x', y') \wedge x' \geq_{\text{Int}} 0 \wedge y' >_{\text{Int}} 0 \end{aligned}$$

The proof tree for the formula $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$ is



As in the previous cases, we apply [Implication] to deduce $S \cup G \Vdash \varphi_4^{\text{while}} \Rightarrow \varphi'^{\text{while}}$, and then by a few applications of [Transitivity] we obtain $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$.

This concludes the proof of the set of goals $\{\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}, \varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}\}$, and, in particular, of the fact that `gcd` meets its specification (4.3). Note how the proofs of all goals have used symbolic execution as well other goals as circular hypotheses. Moreover, the proof obligation stating that the loop's body satisfies the invariant, which would be required in Hoare logic, is no longer necessary since it is implicitly proved by the symbolic steps and a circular application of the rule specifying the loop.

4.2.1 Default strategy for automation

Our proof system still leaves some freedom regarding the order of rule applications and still requires some creative user input when, e.g, choosing the patterns $\varphi_1, \varphi_2, \varphi''$ in its rules. We define the following strategy (already applied in the previous example) in order to completely automate proof searches. In the default strategy, the rules are applied with the following priorities (note that all the rules are applied "bottom-up"):

1. Implication has highest priority. It is only applied for closing proof branches;
2. CircularHypothesis, followed by all the applications of CaseAnalysis required to break disjunctive patterns into elementary patterns, has second priority;
3. SymbolicStep, also followed by as many applications of CaseAnalysis as needed to break disjunctive patterns into elementary ones, has third priority.

Note that *Transitivity* is not restricted in any way. In fact, it is just used implicitly for building larger proofs from smaller proofs steps. We call the above strategy the *default strategy*. It transforms proof attempts of $\mathcal{SUG} \Vdash \Delta_{\mathcal{S}}(G)$ into the building of the symbolic transition relation $\Rightarrow^{\mathcal{SUG}}$, which is done by our symbolic execution tool [5]. More details on the implementation are given in the Section 5.1.

4.2.2 Weak Completeness: disproving RL formulas

It is a good idea to try to disprove RL formulas as well as to try to prove them. The default strategy of our proof system can also be used for disproving formulas: if it terminates in by failing to prove its input (a set of RL formulas, with some reasonable restrictions presented below) then the input in question is invalid.

A rule $\varphi \Rightarrow \varphi'$ is *terminal* if φ' is non-derivable for \mathcal{SUG} . For example, the RL formulas $\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$ and $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$ (Example 4.2.1) are terminal, because their right-hand sides contain empty code that cannot be executed further. The specification of a program, like $\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}$ for `gcd`, is typically terminal because it refers to what the program computes when it terminates. Auxiliary formulas, like $\varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}$, may or may not be terminal.

A RL formula $\varphi \Rightarrow \varphi'$ is *terminating* if all configurations $\gamma \in \llbracket \varphi \rrbracket$ are terminating, and a set of formulas is terminating iff every formula in it is terminating. For example, $\{\varphi^{\text{gcd}} \Rightarrow \varphi'^{\text{gcd}}, \varphi^{\text{while}} \Rightarrow \varphi'^{\text{while}}\}$ is terminating. A set \mathcal{S} of RL formulas is *confluent* if the transition relation $\Rightarrow_{\mathcal{S}}^*$ is confluent (i.e. the terms can be rewritten in more than one way, but the result is always the same).

The next theorem is our weak completeness result. It assumes a situation where, on a given proof branch for a terminal goal ($\varphi \Rightarrow \varphi'$), the default strategy of our proof system is "stuck"; thus, its execution terminates in failure.

Theorem 4.2.2 (Weak Completeness) *Consider a confluent set of RL formulas \mathcal{S} , a set of terminating formulas $G = \{\pi_i \wedge \phi_i \Rightarrow \pi'_i \wedge \phi'_i \mid i \in I\}$, a terminal formula $\varphi \Rightarrow \varphi' \in G$, and a proof branch of $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(G)$ generated by the default strategy, starting from $\mathcal{S} \cup G \Vdash \Delta_{\mathcal{S}}(\varphi \Rightarrow \varphi')$ and ending in $\mathcal{S} \cup G \Vdash \varphi'' \Rightarrow \varphi'$, such that φ'' is not derivable for $\mathcal{S} \cup G$ and*

$\not\models_{\text{ML}} \varphi'' \rightarrow \varphi'$. Then $\mathcal{S} \not\models_{\text{RL}} G$.

Proof From the proof branch we extract a symbolic execution $[\varphi]_{\sim} \Rightarrow_{\mathcal{S} \cup G}^{\ast} \varphi''_{\sim}$. Let $\gamma'' \in \llbracket \varphi'' \rrbracket \setminus \llbracket \varphi' \rrbracket$. A (feasible) symbolic execution is simulated by a concrete execution $\gamma \Rightarrow_{\mathcal{S} \cup G}^{\ast} \gamma''$ with $\gamma \in \llbracket \varphi \rrbracket$, thanks to Corollary 3.3.2. We assume (by contradiction) that $\mathcal{S} \models_{\text{RL}} G$.

First, we show (\diamond): $\gamma \Rightarrow_{\mathcal{S}}^{\ast} \gamma''$. For this, we show that for every step, say, $\gamma_1 \Rightarrow_{\{\alpha\}} \gamma_2$ with $\alpha \in G$, there is an execution with rules in \mathcal{S} , i.e., $\gamma_1 \Rightarrow_{\mathcal{S}}^{\ast} \gamma_2$. Then we replace every such step $\gamma_1 \Rightarrow_{\{\alpha\}} \gamma_2$ in $\gamma \Rightarrow_{\mathcal{S} \cup G}^{\ast} \gamma''$ ($\alpha \in G$) by the corresponding execution $\gamma_1 \Rightarrow_{\mathcal{S}}^{\ast} \gamma_2$ and obtain the desired execution $\gamma \Rightarrow_{\mathcal{S}}^{\ast} \gamma''$.

We now prove (\spadesuit) $\gamma_1 \Rightarrow_{\mathcal{S}}^{\ast} \gamma_2$ from $\gamma_1 \Rightarrow_{\{\alpha\}} \gamma_2$ with $\alpha \in G$. From the assumption $\mathcal{S} \models_{\text{RL}} G$ we obtain $\mathcal{S} \models_{\text{RL}} \alpha (= \pi_1 \wedge \phi_1 \Rightarrow \pi_2 \wedge \phi_2 \in G)$. From $\gamma_1 \Rightarrow_{\{\alpha\}} \gamma_2$ we get $(\gamma_1, \rho) \models_{\text{ML}} \pi_1 \wedge \phi_1$ for some valuation ρ and $(\gamma_2, \rho) \models_{\text{ML}} \pi_2 \wedge \phi_2$. Thus, $\gamma_2 = \rho(\pi_2)$. From $\mathcal{S} \models_{\text{ML}} \alpha$ and $(\gamma_1, \rho) \models_{\text{ML}} \pi_1 \wedge \phi_1$ (note that γ_1 is terminating, since G is terminating) we get that there exists γ'_2 such that $\gamma_1 \Rightarrow_{\mathcal{S}}^{\ast} \gamma'_2$ and $(\gamma'_2, \rho) \models_{\text{ML}} \pi_2 \wedge \phi_2$. In particular, $\gamma'_2 = \rho(\pi_2)$. Thus, $\gamma'_2 = \gamma_2$, so $\gamma_1 \Rightarrow_{\mathcal{S}}^{\ast} \gamma_2$. (\spadesuit) is proved, and so is (\diamond).

We thus have $\gamma \Rightarrow_{\mathcal{S}}^{\ast} \gamma''$, where $\gamma'' \in \llbracket \varphi'' \rrbracket \setminus \llbracket \varphi' \rrbracket$, and γ terminating (since $\gamma \in \llbracket \varphi \rrbracket$ and G is terminating). We have assumed $\mathcal{S} \models_{\text{RL}} G$, in particular, $\mathcal{S} \models_{\text{RL}} \varphi \Rightarrow \varphi'$. This means that from the (terminating) γ , there is $\gamma''' \in \llbracket \varphi' \rrbracket$ and $\gamma \Rightarrow_{\mathcal{S}}^{\ast} \gamma'''$. On the other hand, γ'' is terminal (because φ'' is non-derivable) and $\gamma'' \in \llbracket \varphi'' \rrbracket$, so we obtain thanks to the confluence hypothesis that any execution starting in γ ends up in γ'' . Thus, the successor $\gamma''' \in \llbracket \varphi' \rrbracket$ of γ is on such an execution. There are two cases:

- $\gamma''' = \gamma''$: impossible because $\gamma'' \in \llbracket \varphi'' \rrbracket \setminus \llbracket \varphi' \rrbracket$ (hypothesis).
- γ''' strictly precedes γ'' . This is also impossible (since γ''' has successors), but φ' is non-derivable (since the goal $\varphi \Rightarrow \varphi'$ we started with is terminal).

The contradiction was generated by our assumption $\mathcal{S} \models_{\text{RL}} G$, so, we conclude $\mathcal{S} \not\models_{\text{RL}} G$.

■

Note that the default strategy is "stuck" after the last sequent $\mathcal{S} \cup G \Vdash \varphi'' \Rightarrow \varphi'$ because no rule can be applied from there on, in the current proof

branch: Implication cannot be applied to close the proof branch because $\not\models_{\text{ML}} \varphi'' \rightarrow \varphi'$, `CircularHypothesis` and `SymbolicStep` cannot be applied since φ'' is non-derivable, and `CaseAnalysis` cannot be applied because φ'' is an elementary pattern.

Remark 4.2.4 *Theorem 4.2.2 is proved by showing that there are configurations γ, γ'' such that $\gamma \Rightarrow_{\mathcal{S}}^* \gamma''$, $\gamma \in \llbracket \varphi \rrbracket$, and $\gamma'' \in \llbracket \varphi'' \rrbracket \setminus \llbracket \varphi' \rrbracket$. Using confluence, all executions starting in γ end up in the (terminal) γ'' , and using the fact that $\varphi \Rightarrow \varphi'$ is terminal, no configuration on any of these executions may encounter $\llbracket \varphi' \rrbracket$. If the confluence of \mathcal{S} and/or the fact that G is terminating do not hold, one can attempt, as an alternative approach for establishing weak completeness, to use model checking starting from γ , in order to check whether or not there is a reachable configuration in $\llbracket \varphi' \rrbracket$. The terminating nature of G ensures the model checking will always terminate.*

Together with the soundness result, weak completeness says that when our proof system's default strategy terminates (either successfully, by proving all the goals, or unsuccessfully, by getting stuck on a given proof branch), it correctly solves the problem of whether the goals given to it as input are valid or not. That is, if it terminates, our approach correctly solves the program-verification problem.

Chapter 5

Symbolic execution within \mathbb{K} Framework

In this chapter we present a prototype tool implementing our symbolic execution approach. We first briefly describe the tool and its integration within the \mathbb{K} framework (Section 5.1). Then we illustrate the most significant features of the tool by the means of use cases involving nontrivial languages and programs (Section 5.2). In the rest of the chapter we present a verification tool developed on top of the symbolic execution framework, which is used to perform program verification using Reachability Logic (Section 5.3).

5.1 Symbolic Execution within the \mathbb{K} Framework

Our tool is part of \mathbb{K} [97, 109], a rewrite-based executable semantics framework in which programming languages, type systems, and formal analysis tools can be defined. Beside some toy languages used for teaching, there are a few real-life programming languages, supporting different paradigms, that have been successfully defined in \mathbb{K} (e.g. C [35], Java [17] - see Section 2.5 for details).

In \mathbb{K} , the definition of a programming language, say, \mathcal{L} , is compiled into a rewrite theory. Then, the \mathbb{K} runner executes programs in \mathcal{L} by applying the resulting rewrite rules to configurations containing programs. Our implementation follows the same process. The main difference is that our

new \mathbb{K} compiler includes some transformations steps: rule linearisation, replacing the data subterms in a rule with variables and add them into the rule's condition, add a configuration cell for path condition, and modify rules such that the path condition is updated appropriately at runtime. The effect is that the compiled rewrite theory we obtain defines the so-called symbolic semantics of \mathcal{L} instead of its concrete semantics. We note that the symbolic semantics can execute programs with concrete inputs as well, and thus, the initial semantics of \mathcal{L} is not affected. For user convenience we have also improved the \mathbb{K} runner with some specific options which are useful for providing programs with symbolic input and setting up the initial path condition.

By choosing the \mathbb{K} framework as a basis for our tool we take the full advantage of its language-independent capabilities. Given the \mathbb{K} definition of a programming language we can generate fully automatically the symbolic semantics, and then run programs with symbolic values. Therefore, users which already have a \mathbb{K} definition of a language can perform symbolic execution with no additional effort.

The main part of the tool was written in the same language as the \mathbb{K} tool, namely Java. There are also several pieces of code for interaction with the Z3 SMT solver written in \mathbb{K} itself. In the following, we are going to describe in detail the options of the tool. For convenience, we assume that a file named `cink.k` contains the \mathbb{K} definition of CinK (Section 2.5) and CinK programs are stored in files having the `*.cink` extension (e.g. `sum.cink`).

5.1.1 Compiling the symbolic semantics

In order to compile the symbolic semantics of CinK one must type the following command:

```
$ kcompile cink.k --backend symbolic
```

This option enables the following transformations:

- *Generate symbolic values support.* This step is responsible for generating the set *Var* of symbolic values and also the predicates which state their membership to different sorts.

- *Add input support for symbolic values.* This transformation searches the configuration for cells which are connected to the input stream and generates an input variable `$IN` which can be initialised with symbolic values when running programs.
- *Add cell for path condition.* To keep track of the current path condition we use the \mathbb{K} configuration of the language which is enriched with a new cell called `path-condition`. This gives us the freedom to handle the path condition in a language-independent manner and also brings us the advantage of attaching separate path conditions to different result configurations.
- *Rule linearisation.* In practice, the \mathbb{K} rewrite rules are not linear (i.e. they contain more than once the same variable in the left hand side) as we assume in Chapter 3. This step turns them into linear rules by replacing duplicate variables with fresh ones and adding corresponding equalities in the rule's condition.
- *Add rule condition to path condition.* This transformation consists in modifying the rule such that the path condition cell is updated with the rule's condition. If $l \wedge b \Rightarrow r$ is our current rule and ψ the current path condition, then the new path condition will be $\psi \wedge b$ if $\psi \wedge b$ is not unsatisfiable.

The last transformation step is probably the most important. It involves a call to the SMT solver for checking the satisfiability of the future state of the program. If the solver finds it unsatisfiable then the rule is not applied because it makes no sense to investigate an infeasible execution path. However, the rule will apply when the solver is not able to figure out if the formula is satisfiable or not (i.e. returns *unknown*), since otherwise the tool can miss feasible execution paths. Another discussion related to the SMT call is how the tool deals with the rule condition b when it contains expressions which cannot be sent to the solver, like inductive definitions and predicates used by the \mathbb{K} tool for matching. For that, we provide an additional step which splits b into $b_1 \wedge b_2$, where b_1 is the part which is sent to the solver, and b_2 is the part related to matching. The condition b is transformed into its conjunctive normal form and then the split is done using a tagging mechanism which identifies which clauses

are actual matching predicates. After the split, the new path condition becomes $\psi \wedge b_1$ instead of $\psi \wedge b$ and the new rule condition is b_2 , which is relevant for matching purposes. On the other hand, the user must provide by hand the inductive definitions (if any) in \mathbb{K} . Let us consider that gcd , given using an inductive definition in \mathbb{K} is part of b_1 . The tool is able to simplify b_1 before sending it to the solver, by applying the definition of gcd . If the formula reduces to false, then it is unsatisfiable and for sure the new path is unfeasible and will not be explored. If not, b_1 is sent to the solver as it is, gcd being translated as an uninterpreted function. However, a question arises here: can we still trust the response given by the solver when gcd can be interpreted by the solver as *any* function? If the solver says that the formula is unsatisfiable then we know for sure that for any interpretation of gcd (including the real greatest common divisor function) it remains unsatisfiable. On the other hand, if the solver says that the formula is satisfiable then we do not know for sure that it is also satisfiable for the appropriate interpretation of gcd . In this case the tool will consider the formula satisfiable and will continue to explore this execution path. In some sense, because the solver is not able to deal with all types of formulas we are interested in, the tool over-approximates the space of feasible paths.

5.1.2 Running programs with symbolic input

The CinK language also allows I/O interactions. This enables the execution of CinK programs with symbolic input data. We compile the resulting definition by calling the \mathbb{K} compiler with `--backend symbolic` option as discussed in Section 5.1.1. Now, programs such as `sum.cink` (shown in Figure 5.1) can be run with the \mathbb{K} runner in the following ways:

1. with concrete inputs;
2. with symbolic inputs;
3. on one arbitrary execution path, or on all paths up to a given bound;
4. in a step-wise manner, or by letting the program completely execute a given number of paths.

```

void main() {
    int n, s;
    n = read();
    s = 0;
    while (n > 0) {
        s = s + n;
        n = n - 1;
    }
    cout << "Sum = " << s;
}

```

Figure 5.1: `sum.cink`

For example, by running `sum.cink` with a symbolic input n (here and thereafter we use mathematical font for symbolic values) and requiring at most five completed executions, the \mathbb{K} runner outputs the five resulting, final configurations, one of which is shown below, in a syntax slightly simplified for readability:

```

$ krun sum.cink -cIN= $n$  -cPC="true"
<k> . </k>
<path-condition>
     $n > 0 \wedge (n - 1 > 0) \wedge \neg((n - 1) - 1 > 0)$ 
</path-condition>
<env>
     $n \mapsto l_n$ 
     $s \mapsto l_s$ 
</env>
<state>
     $l_n \mapsto (n - 1) - 1$ 
     $l_s \mapsto n + (n - 1)$ 
</state>
...

```

The program is finished since the `k` cell has no code left to execute. The path condition actually means $n = 2$, and in this case the sum s equals $n + (n - 1) = 2 + 1$, as shown by the `state` cell. The `-cIN= n` option sends on the input stream of the program the symbolic value n , while `-cPC="true"` sets the initial path condition.

To get all possible solutions the tool has be run using `--search`:

```
$ krun sum.cink -cIN= $n$  -cPC="true" --search
```

Since `sum.cink` contains a loop that depends on the symbolic value n , the command above will run forever. Still, users can bound the execution to a specific number of solutions:

```
$ krun sum.cink -cIN= $n$  -cPC="true" --search --bound 3
```

```
Solution 1
```

```
...  
<k> . </k>  
<path-condition>  
     $\neg(n > 0)$   
</path-condition>  
...
```

```
Solution 2
```

```
...  
<k> . </k>  
<path-condition>  
     $n > 0 \wedge \neg(n - 1 > 0)$   
</path-condition>  
...
```

```
Solution 3
```

```
...  
<k> . </k>  
<path-condition>  
     $n > 0 \wedge (n - 1 > 0) \wedge \neg((n - 1) - 1 > 0)$   
</path-condition>  
...
```

The command above shows three possible execution paths which compute the sums of numbers up to 0, 1, and 2 respectively. The `--bound 3` option bounds the number of solutions to three. The number of solutions

can be also bounded using the initial path condition:

```
$ krun sum.cink -cIN=n -cPC="0 ≤ n ∧ n < 3" --search
Solution 1
...
<k> . </k>
<path-condition>
    ¬(n > 0)
</path-condition>
...

Solution 2
...
<k> . </k>
<path-condition>
    n > 0 ∧ ¬(n - 1 > 0)
</path-condition>
...

Solution 3
...
<k> . </k>
<path-condition>
    n > 0 ∧ (n - 1 > 0) ∧ ¬((n - 1) - 1 > 0)
</path-condition>
...
```

Users can run the program in a step-wise manner using the \mathbb{K} runner debugger in order to see intermediary configurations in addition to final ones. During this process they can interact with the runner, e.g., by choosing one execution branch of the program among several, feeding the program with inputs, or letting the program run on an arbitrarily chosen path until its completion.

5.2 Symbolic execution: use cases

In this section we illustrate some direct use cases for our symbolic execution tool. Moreover, we run our tool on multiple programming languages, already defined in \mathbb{K} , to emphasise that our tool is language independent. In Section 5.2.1 we show how the symbolic execution framework works together with the Maude LTL (Linear Temporal Logic) model checker [77]. Then, in Section 5.2.2, we show a simple example of bounded model checking over programs which manipulate symbolic arrays. Finally, we test our symbolic execution tool on an object oriented language and we use it to check the well-known virtual method call mechanism (Section 5.2.3).

5.2.1 LTL model checking

The \mathbb{K} runner includes a hook to the Maude LTL (Linear Temporal Logic) model checker [77]. Thus, one can model check LTL formulas on programs having a finite state space (or by restricting the verification to a finite subset of the state space). This requires an (automatic) extension of the syntax and semantics of a language for including labels that are used as atomic propositions in the LTL formulas. Predicates on the program's variables can be used as propositions in the formulas as well, using the approach outlined in [71]. We have enriched the \mathbb{K} definition of CinK with syntax and semantics for atomic propositions. Below is the LTL satisfaction relation for the atomic proposition `logInv(a,x,k)`:

$$B \models_{Ltl} \text{logInv}(A, X, K) \Rightarrow true \text{ requires } \text{val}(B, X) *_{Int} \text{pow}(2, \text{val}(B, K)) \leq_{Int} \text{val}(B, A) \\ \wedge_{Bool} (\text{val}(B, A) <_{Int} (\text{val}(B, X) +_{Int} 1) *_{Int} \text{pow}(2, \text{val}(B, K)))$$

The variable B has sort *Cfg*, while variables A , B , and K are all identifiers, which are mapped in the configuration to integer values. The functions `pow` and `val` are additional functions used for computing the mathematical power function, and for retrieving the value of an identifier from the configuration, respectively. For each such rule, the tool automatically generates support for checking whether the current path condition implies the rule's condition. If this implication does not hold then the rule does not apply, since we are in a state where the LTL property does not hold.

```

void main() {
  int k, a, x;
  a = read();
  x = a;
  while (x > 1) {
    x = x / 2;
    k = k + 1;
    L : {}
  }
}

```

Figure 5.2: `log.cink`

Consider for instance the program `log.cink` in Figure 5.2, which computes the integer binary logarithm of an integer read from the input. We prove that whenever the loop visits the label `L`, the inequalities $x * 2^k \leq a < (x + 1) * 2^k$ hold. The invariant was guessed using several step-wise executions. We let `a` be a symbolic value and restrict it in the interval `(0..10)` to obtain a finite state space. We prove that the above property, denoted by `logInv(a,x,k)` holds whenever the label `L` is visited and `a` is in the given interval, using the following command (again, slightly edited for better readability):

```

$ krun log.cink -cPC="a >Int 0 ^ Bool a <Int 10" -cIN="a"
-tlmc "□Ltl (L →Ltl logInv(a, x, k))"
true

```

The \mathbb{K} runner executes the command by calling the Maude LTL model-checker for the LTL formula $\square_{Ltl} (L \rightarrow_{Ltl} \text{logInv}(a, x, k))$ and the initial configuration having the program `log.cink` in the computation cell `k`, the symbolic value `a` in the input cell `in`, and the constraint $a >_{Int} 0 \wedge_{Bool} a <_{Int} 10$ in the path condition. The result returned by the tool is that the above LTL formula holds.

5.2.2 SIMPLE, symbolic arrays, and bounded model checking

In this section we illustrate symbolic arrays in the SIMPLE [102] language and we show how the \mathbb{K} runner can directly be used for performing bounded model checking. SIMPLE is a non-trivial imperative programming language which includes multidimensional arrays and array refer-

```

void init(int[] a, int x, int j){
    int i = 0, n = sizeof(a);
    a[j] = x;
    while (a[i] != x && i < n) {
        a[i] = 2 * i;
        i = i + 1;
    }
    if (i > j) {
        print("error");
    }
}

void main() {
    int n = read();
    int j = read();
    int x = read();
    int a[n], i = 0;
    while (i < n) {
        a[i] = read();
        i = i + 1;
    }
    init(a, x, j);
}

```

Figure 5.3: SIMPLE program: `init-arrays`

ences, functions and function values, blocks with local variables, exceptions, concurrency via dynamic thread creation/termination and synchronisation. Further details about SIMPLE and its \mathbb{K} definition can be found in [102].

In the program in Figure 5.3, the `init` method assigns the value `x` to the array `a` at an index `j`, then fills the array with ascending even numbers until it encounters `x` in the array; it prints `error` if the index `i` went beyond `j` in that process. The array and the indexes `i`, `j` are parameters to the function, passed to it by the `main` function which reads them from the input. In [3] it has been shown, using model-checking and abstractions on arrays, that this program never prints `error`.

We obtain the same result by running the program with symbolic inputs and using the \mathbb{K} runner as a bounded model checker:

```
$ krun init-arrays.simple -cPC="n >_Int 0" -search -cIN="n j x a b c"
-pattern="<T> <out> error </out> B:Bag </T>"
```

Search results:

No search results

The initial path condition is $n >_{Int} 0$. The symbolic inputs for `n`, `j`, `x` are entered as `n j x`, and the array elements `a b c` are also symbolic. The `--pattern` option specifies a pattern to be searched in the final configuration: the text `error` should be in the configuration's output buffer. The above command thus performs a bounded model-checking with symbolic inputs (the bound is implicitly set by the number of array elements given as inputs - 3). The command above does not return any solution, meaning that the program will never print `error`.

The result was obtained using symbolic execution without any additional tools or techniques. We note that array sizes are symbolic as well, a feature that, to our best knowledge, is not present in other symbolic execution frameworks.

5.2.3 KOOL: testing virtual method calls on lists

Another example we consider for analysis is a program written in the KOOL [53] programming language. KOOL is a concurrent, dynamic, object oriented language inspired from Smalltalk [48], and besides the imperative features (assignments, conditional statements, loops with break and continue) it includes support for many familiar object oriented features (classes, methods, inheritance, dynamic dispatch, exceptions via try/-catch, ...).

In Figure 5.4 we show a KOOL program which implements lists and ordered lists of integers using arrays. We use symbolic execution to check the well-known virtual method call mechanism of object-oriented languages: the same method call, applied to two objects of different classes, may have different outcomes.

The `List` class implements (plain) lists. It has methods for creating, copying, and testing the equality of lists, as well as for inserting and deleting elements in a list. Figure 5.4 shows only insertion and deletion of elements. The class `OrderedList` inherits from `List`. It redefines the `insert` method in order to ensure that the sequences of elements in lists are sorted in increasing order. The `Main` class creates a list `l1`, initializes `l1` and an integer variable `x` with input values, copies `l1` to a list `l2` and then inserts and deletes `x` in `l1`. Finally it compares `l1` to `l2` element by element, and prints *error* if it finds them different.

We use symbolic execution to show that the above sequence of method calls results in different outcomes, depending on whether `l1` is a `List` or an `OrderedList`. We first try the case where `l1` is a `List`, by issuing the following command to the \mathbb{K} runner:

```
$ krun lists.kool -search -cIN="e1 e2 x"
                    -pattern="<T> <out> error </out> B:Bag </T>"
```

Solution 1, State 50:

<path-condition>

$$e_1 = x \wedge_{Bool} \neg_{Bool} (e_1 = e_2)$$

```

class List {
  int a[10];
  int size, capacity;
  ...

  void insert (int x) {
    if (size < capacity) {
      a[size] = x; ++size;
    }
  }

  void delete(int x) {
    int i = 0;
    while(i < size-1 && a[i] != x) {
      i = i + 1;
    }
    if (a[i] == x) {
      while (i < size - 1) {
        a[i] = a[i+1];
        i = i + 1;
      }
      size = size - 1;
    }
  }
  ...
}

class OrderedList extends List {
  ...
  void insert(int x){
    if (size < capacity) {
      int i = 0, k;
      while(i < size && a[i] <= x) {
        i = i + 1;
      }
      ++size; k = size - 1;
      while(k > i) {
        a[k] = a[k-1]; k = k - 1;
      }
      a[i] = x;
    }
  }
}

class Main {
  void Main() {
    List l1 = new List();
    ... // read elements of l1 and x
    List l2 = l1.copy();
    l1.insert(x); l1.delete(x);
    if (l2.eqTo(l1) == false) {
      print("error\n");
    }
  }
}

```

Figure 5.4: lists.kool: implementation of lists in KOOL

</path-condition>

...

The command initializes `l1` with two symbolic values (e_1, e_2) and sets the program variable `x` to the symbolic value x . It searches for configurations that contain *error* in the output. The tool finds one solution, with $e_1 = x$ and $e_1 \neq e_2$ in the path condition. Since `insert` of `List` appends x at the end of the list and `delete` removes the first instance of x from it, `l1` consists of (e_2, x) when the two lists are compared, in contrast to `l2`, which consists of (e_1, e_2) . The path condition implies that the lists are different.

The same command on the same program but where `l1` is an `OrderedList` finds no solution. This is because `insert` in `OrderedList` inserts an element in a unique place (up to the positions of the elements equal to it) in an ordered list, and `delete` removes either the inserted element or one with the same value. Hence, inserting and then deleting an element leaves an ordered list unchanged.

Thus, virtual method call mechanism worked correctly in the tested scenarios. An advantage of using our symbolic execution tool is that the condition on the inputs that differentiated the two scenarios was discovered by the tool. This feature can be exploited in other applications such as test-case generation.

5.3 A prototype for Reachability-Logic Verification

In this section we illustrate on a few examples a prototype tool which implements the default strategy (shown in Section 4.2.1) of the \Vdash deductive system that we have defined (Figure 4.6), in order to verify $\mathcal{S} \models_{\text{RL}} G$, where \mathcal{S} is a given language semantics and G a set of reachability formulas (goals). The implementation is part of the \mathbb{K} tools suite [74] and it has been developed on top of our symbolic execution tool presented in Section 5.1.

In terms of implementation, our prototype reuses components of the \mathbb{K} framework: parsing, compilation steps, support for symbolic execution, and connections to Maude’s [77] state-space explorer and to the Z3 SMT solver [34]. Given a language definition \mathcal{S} and a set of RL formulas G , `kcheck` build a new definition $\mathcal{S} \cup G$ and then it simply applies the default strategy for proving the formulas in G , described in Section 4.2.1. We have used `kcheck` to prove `gcd.cink` (Figure 4.5) as sketched in Example 4.2.1.

Except trivial examples, like `gcd` and `sum`, we tried our tool on more complex ones. Moreover, since our approach is parametric in language definitions, it can be applied to other \mathbb{K} language definitions as well. Thus, in the following sections we describe our experience when verifying two non-trivial programs with `kcheck`, each experiment involving a different language and a different (type of) program. In Section 5.3.1 we have extended `CinK` with threads and we verified that a disjoint parallel program, called `FIND`, successfully finds (in parallel) the index of the first positive integer in an array. In Section 5.3.2, using the \mathbb{K} definition of `CinK` and `kcheck` we verified the Knuth-Morris-Pratt [63] string matching algorithm.

5.3.1 Verifying a parallel program: FIND

The example is inspired from [2]. Given an integer array a and a constant $N \geq 1$, the program in Figure 5.5 finds the smallest index $k \in \{1, \dots, N\}$

```

i = 1;          S1 = while (i < oddtop){
j = 2;          if (a[i] > 0) then{oddtop = i;}
oddtop = N + 1; else {i = i + 2;}
eventop = N + 1; }
S1 || S2;      S2 = while (j < eventop) {
if (oddtop > eventop) if (a[j] > 0) then{eventop = j;}
then { k = eventop; } else{j = j + 2;}
else { k = oddtop; } }

```

Figure 5.5: FIND program.

such that $a[k] > 0$. If such an index k does not exist then $N + 1$ is returned. It is a *disjoint* parallel program, which means that its parallel components only have reading access to the variable a they share.

In order to verify FIND, we have enriched the \mathbb{K} semantics of CinK with dynamic threads and the $||$ operator, which executes in parallel two threads corresponding to S1 and S2. Each thread $\langle \cdot \rangle_{\text{th}}$ has its own computations $\langle \cdot \rangle_k$ and environment $\langle \cdot \rangle_{\text{env}}$ cells, while $\langle \cdot \rangle_{\text{store}}$ is shared among the threads. Threads also have an $\langle \cdot \rangle_{\text{id}}$ (identifier) cell. The configuration is shown below (the $+$ on the $\langle \cdot \rangle_{\text{th}+}$ cell says that the cell contains at least one thread):

$$\langle \langle \langle \cdot \rangle_k \langle \cdot \rangle_{\text{env}} \langle \cdot \rangle_{\text{id}} \rangle_{\text{th}+} \langle \cdot \rangle_{\text{store}} \langle \cdot \rangle_{\text{result}} \langle \cdot \rangle_{\text{stack}} \langle \cdot \rangle_{\text{in}} \langle \cdot \rangle_{\text{out}} \rangle_{\text{cfg}}$$

The $||$ operator yields a non-deterministic behaviour of FIND, i.e. it can generate more than one computation starting in a given initial state. However, in [2] the authors prove that all computations of a disjoint parallel program starting in the same initial state produce the same output. So, all the computations of a disjoint parallel program are equivalent to the sequential execution of its components. For program verification this observation simplifies matters because it allows independent verification of the parallel code, without considering all the interleavings caused by parallelism.

The verification of FIND is given by checking only three rules: one for each of the two loops and one for the main program. This is much simpler than the proof from [2], where more proof obligations must be generated and checked. This is a consequence of the fact that many proof obligations are automatically checked by symbolic execution. Moreover,

when performing mechanised verification, the pre/post conditions and the invariants must be very accurate. Otherwise, the proof will fail even if, intuitively, all the formulas seem to be valid. For example, when using `kcheck` to verify `FIND`, we discovered that the precondition `pre` must be $N \geq 1$ rather than `true` as stated in the (non-mechanised) proof of [2], and in `p2` the value of `j` must be greater-or-equal to 2, a constraint that was also forgotten in [2].

Figure 5.6 shows all the ingredients that we used to prove the correctness of the program `FIND` (Figure 5.5) using our tool. At the figure’s top we show the code macros that we use in our RL formulas. Below the code macros we include the formulas corresponding to the pre/post conditions and invariants used by the authors of [2] in their proof. The program is checked by applying the implementation `kcheck` of our proof system on the consisting of the three RL formula-set $G = \{(\clubsuit), (\diamond), (\spadesuit)\}$. On the bottom lines we show the proofs automatically constructed by `kcheck`.

We believe that the number of three proof obligations, given by G , is minimal for verifying `FIND`. Initially we started we eight rules describing the proof obligations used in [2]. Then, based on the `gcd` examples and others inspired from the same source, we realised that all sequential program fragment specifications can be removed since they can be automatically proved using the `SymbolicStep` rule, which, together with `Transition` and `CaseAnalysis`, amounts to symbolic execution. Since the configuration for the new language is more complex, the syntax for these rules is a bit cumbersome, but it can be generated from the annotations of the program by using symbolic execution to determine the exact structure of the configuration at the point where such a rule should be applied.

The proof trees for the RL formulas (\clubsuit) and (\diamond) are similar to the one for the while-loop in the `gcd` program. However, here the second branch is split by a new use of the `CaseAnalysis` rule, due to the `if` statement from the loop’s body. The proof tree for the RL formula (\spadesuit) , corresponding to the specification of `FIND`, has a single branch because it uses circularities (\clubsuit) and (\diamond) that do not split the proof tree.

The formulas are nontrivial, and it took us several iterations to come up with the exact ones, during which we used the tool in a trial-and-error process. The automatic nature of the tool, as well as the feedback it returned when it failed, were particularly helpful during this process. In

CODE MACROS	
INIT	\triangleq <code>i = 1; j = 2; oddtop = N + 1; eventop = N + 1;</code>
BODY1	\triangleq <code>{if (a[i] > 0) then { oddtop = i; } else { i = i + 2; }}</code>
BODY2	\triangleq <code>{if (a[j] > 0) then { eventop = j; } else { j = j + 2; }}</code>
S1	\triangleq <code>while (i < oddtop) BODY1</code>
S2	\triangleq <code>while (j < eventop) BODY2</code>
MIN	\triangleq <code>if (oddtop > eventop) then { k = eventop; } else { k = oddtop; }</code>
FIND	\triangleq <code>INIT S1 S2; MIN</code>
Formula macros	
<i>pre</i>	\triangleq $N \geq 1$
<i>p</i> ₁	\triangleq $1 \leq o \leq N + 1 \wedge i\%2 = 1 \wedge 1 \leq i \leq o + 1$ $\wedge (\forall_{1 \leq l < i})(l\%2 = 1 \rightarrow a[l] \leq 0) \wedge (o \leq N \rightarrow a[o] > 0)$
<i>p'</i> ₁	\triangleq $1 \leq o' \leq N + 1 \wedge i'\%2 = 1 \wedge 1 \leq i' \leq o' + 1$ $\wedge (\forall_{1 \leq l < i'})(l\%2 = 1 \rightarrow a[l] \leq 0) \wedge (o' \leq N \rightarrow a[o'] > 0)$
<i>q</i> ₁	\triangleq $1 \leq o' \leq N + 1 \wedge (\forall_{1 \leq l < o'})(l\%2 = 1 \rightarrow a[l] \leq 0) \wedge (o' \leq N \rightarrow a[o'] > 0)$
<i>p</i> ₂	\triangleq $2 \leq e \leq N + 1 \wedge j\%2 = 0 \wedge 2 \leq j \leq e + 1$ $\wedge (\forall_{1 \leq l < j})(l\%2 = 0 \rightarrow a[l] \leq 0) \wedge (e \leq N \rightarrow a[e] > 0)$
<i>p'</i> ₂	\triangleq $2 \leq e' \leq N + 1 \wedge j'\%2 = 0 \wedge 2 \leq j' \leq e' + 1$ $\wedge (\forall_{1 \leq l < j'})(l\%2 = 0 \rightarrow a[l] \leq 0) \wedge (e' \leq N \rightarrow a[e'] > 0)$
<i>q</i> ₂	\triangleq $2 \leq e' \leq N + 1 \wedge (\forall_{1 \leq l < e'})(l\%2 = 0 \rightarrow a[l] \leq 0) \wedge (e' \leq N \rightarrow a[e'] > 0)$
<i>post</i>	\triangleq $1 \leq k' \leq N + 1 \wedge (\forall_{1 \leq l < k'})(a[l] \leq 0) \wedge (k' \leq N \rightarrow a[k'] > 0)$
Map macros for environment and store	
<i>Env</i>	\triangleq <code>a → a i → i j → j oddtop → o eventop → e N → N k → k</code>
<i>St</i>	\triangleq <code>a → a i → i j → j o → o e → e N → N k → k</code>
<i>St'</i>	\triangleq <code>a → a i → i' j → j' o → o' e → e' N → N k → k'</code>
RL formulas	
\clubsuit	$\langle \langle S1 \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St \rangle_{st} \wedge i < o \wedge p_1 \Rightarrow \langle \langle \cdot \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St' \rangle_{st} \wedge o' \leq i' \wedge p'_1 \wedge q_1$
\diamond	$\langle \langle S2 \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St \rangle_{st} \wedge j < e \wedge p_2 \Rightarrow \langle \langle \cdot \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St' \rangle_{st} \wedge e' \leq j' \wedge p'_2 \wedge q_2$
\spadesuit	$\langle \langle FIND \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St \rangle_{st} \wedge pre \Rightarrow \langle \langle \cdot \rangle_k \langle Env \rangle_{env} \rangle_{th} \langle St' \rangle_{st} \wedge post$
Corresponding proofs given by kcheck	
<i>s</i> (<i>i</i>)	\triangleq <code>[CaseAnalysis], ([SymbolicStep]) ∨ ([SymbolicStep]), [CircularHypothesis](<i>i</i>)</code>
\clubsuit	<code>[SymbolicStep], [CaseAnalysis], [Implication] ∨ (s(\clubsuit), [Implication])</code>
\diamond	<code>[SymbolicStep], [CaseAnalysis], [Implication] ∨ (s(\diamond), [Implication])</code>
\spadesuit	<code>[SymbolicStep] × 5, [CircularHypothesis](1), [CircularHypothesis](2), [Implication]</code>

Figure 5.6: RL formulas necessary to verify FIND. We use **a**, **i**, **j**, **oddtop**, **eventop**, **N**, **k** to denote program variables, **a**, **i**, **j**, **o**, **e**, **N**, **k** to denote locations, and *a*, *i*, *j*, *o*, *e*, *N*, *k* for variables values. We also use *s*(*i*) to denote a common sequence in the proofs of (\clubsuit) and (\diamond). `CaseAnalysis` splits the proof in two goals separated by \vee , while `CircularHypothesis`(*i*) represents the application of the formula (*i*) as a circularity. `[SymbolicStep] × n` is the equivalent of applying `[SymbolicStep]` *n* times.

particular symbolic execution was fruitfully used for the initial testing of programs before they were verified.

5.3.2 Verifying the Knuth-Morris-Pratt string matching algorithm: KMP

The Knuth-Morris-Pratt algorithm [63] searches for occurrences of a word P , usually called *pattern*, within a given text T by making use of the fact that when a mismatch occurs, the pattern contains sufficient information to determine where the next search should begin. A detailed description of the algorithm, whose CinK code is shown in Figure 5.7, can be found in [30].

The KMP algorithm optimises the naive search of a pattern into a given string by using some additional information collected from the pattern. For instance, let us consider $T = \text{ABADABCD A}$ and $P = \text{ABAC}$. It can be easily observed that **ABAC** does not match **ABADABCD A** starting with the first position because there is a mismatch on the fourth position, namely $C \neq D$.

The KMP algorithm uses a *failure function* π , which, for each position j in P , returns the length of the longest proper prefix of the pattern which is also a suffix of it. For our example, $\pi[3] = 1$ and $\pi[j] = 0$ for $j = 1, 2, 4$. In the case of a mismatch between the position i in T and the position j in P , the algorithm proceeds with the comparison of the positions i and $\pi[j]$. For the above mismatch, the next comparison is between the **B** in **ABAC** and the first instance of **D** in **ABADABCD A**, which saves a comparison of the characters preceding them, since the algorithm "already knows" that they are equal (here, they are both **A**).

An implementation of KMP is shown in Figure 5.7. The comments include the specifications for preconditions, postconditions, and invariants, which will be explained later in this section (briefly, they are syntactic sugar for RL formulas, which are automatically generated from them). The program can be run either using the \mathbb{K} semantics of CinK or the `g++` GNU compiler. The `compute_prefix` function computes the failure function π for each component of the pattern and stores it in a table, called `pi`. The `kmp_matcher` searches for all occurrences of the pattern in the string comparing characters one by one; when a mismatch is found on positions i in the string and q in the pattern, the algorithm shifts the search to the right as many positions as indicated by `pi[q]`, and initiates a

new search. The algorithm stops when the string is completely traversed.

For the proof of KMP we use the original algorithm as presented in [30]. Another formal proof of the algorithm is given in [40] by using Why3 [41]. There, the authors collapsed the nested loops into a single one in order to reduce the number of invariants they have to provide. They also modified the algorithm to stop when the first occurrence of the pattern in the string was found. By contrast, we do not modify the algorithm from [30]. We also prove that KMP finds *all* the occurrences of the pattern in the string, not only the first one. We let $P[1..i]$ denote the prefix of P of size i , and $P[i]$ denote its i -th element.

Definition 5.3.1 *Let P be a pattern of size $m \geq 1$ and T a string of characters of size $n \geq 1$. We define the following functions and predicate:*

- $\pi(i)$ is the length of the longest proper prefix of $P[1..i]$ which is also a suffix for $P[1..i]$, for all $1 \leq i \leq m$;
- $\theta(i)$ is the length of the longest prefix of P that matches T on the final position i , for all $1 \leq i \leq n$;
- $allOcc(Out, P, T, i)$ holds iff the list Out contains all the occurrences of P in $T[1..i]$.

The specification of the `kmp_matcher` function is the following RL formula:

$$\left\langle \begin{array}{l} \langle \text{kmp_matcher}(\mathbf{p}, \mathbf{t}, m, n); \rangle_{\mathbf{k}} \langle \cdot \rangle_{\text{out}} \\ \langle \mathbf{p} \mapsto l_1 \ \mathbf{t} \mapsto l_2 \rangle_{\text{env}} \langle l_1 \mapsto P \ l_2 \mapsto T \rangle_{\text{store}} \ \dots \end{array} \right\rangle_{\text{cfg}} \wedge n \geq 1 \wedge m \geq 1$$

$$\Rightarrow$$

$$\langle \langle \cdot \rangle_{\mathbf{k}} \langle Out \rangle_{\text{out}} \langle \dots \rangle_{\text{env}} \langle \dots \rangle_{\text{store}} \dots \rangle_{\text{cfg}} \wedge allOcc(Out, P, T, n)$$

This formula says that from a configuration where the program variables \mathbf{p} and \mathbf{t} are bound to the values P , T , respectively, the output cell is empty, and the `kmp_matcher` function has to be executed, one reaches a configuration where the function has been executed and the output cell contains all the occurrences of P in T . Note that we passed the symbolic values m and n as actual parameters to the function which are the sizes of P , and T , respectively. An advantage of RL with respect to Hoare

```

/*@pre: m>=1 */
void compute_prefix(char p[],
                   int m, int pi[])
{
  int k, q;
  k = 0;
  pi[1] = 0;
  q = 2;
  while(q <= m) {
    /*@inv: 0<=k /\ k<q /\ q<=m+1 /\
      (forall u:1..k)(p[u]=p[q-k+u]) /\
      (forall u:1..q-1)(pi[u]=Pi(u)) /\
      Pi(q)<=k+1 */
    while (k > 0 && p[k+1] != p[q]) {
      /*@inv: 0<=k /\ k<q /\ q<=m /\
        (forall u:1..k)(p[u]=p[q-k+u]) /\
        (forall u:1..q-1)(pi[u]=Pi(u)) /\
        (forall u:1..m)(0<=Pi(u)<u) /\
        Pi(q)<=k+1 */
      k = pi[k];
    }
    if (p[k + 1] == p[q]) {
      k = k + 1;
    }
    pi[q] = k;
    q++;
  }
}
/*@post: (forall u:1..m)(pi[u]=Pi(u)) */

/*@pre: m>=1 /\ n>=1 */
void kmp_matcher(char p[], char t[], int m, int n)
{
  int q = 0, i = 1, pi[m];
  compute_prefix(p, m, pi);
  while (i <= n) {
    /*@inv: 1<=m /\ 0<=q<=m /\ 1<=i<=n+1 /\
      (forall u:1..q-1)(pi[u]=Pi(u)) /\
      (exists v)(forall u:v+1..i-1)(Theta(u)<m /\
        allOcc(Out,p,t,v)) /\
      (forall u:1..q)(p[u]=t[i-1-q+u]) /\
      Theta(i)<=q+1 */
    while (q > 0 && p[q + 1] != t[i]) {
      /*@inv: 1<=m /\ 0<=q /\ q<=m /\
        (forall u:1..q-1)(pi[u]=Pi(u)) /\
        (exists v)(forall u:v+1..i-1)(Theta(u)<m /\
          allOcc(Out,p,t,v)) /\
        (forall u:1..q)(p[u]=t[i-1-q+u]) /\
        (forall u:1..i-1)(Theta(u)<m) /\
        Theta(i)<=q+1 */
      q = pi[q];
    }
    if (p[q + 1] == t[i]) { q = q + 1; }
    if (q == m) {
      cout << "shift: " << (i - m) << endl;
      q = pi[q];
    }
    i++;
  }
}
/*@post: allOcc(Out, p, t, n) */

```

Figure 5.7: The KMP algorithm annotated with pre-/post-conditions and invariants: failure function (left) and the main function (right). Note that we used Pi , Theta , and allOcc to denote functions π and θ , and predicate allOcc , respectively.

Logic is, in addition to language independence, the fact that RL formulas may refer to all the language’s configuration, whereas Hoare Logic formulas may only refer to program variables. A Hoare Logic formula for the `kmp_matcher` function would require the addition of assignments to a new variable playing the role of our output cell.

There are some additional issues concerning the way users write the RL formulas. These may be quite large depending on the size of the \mathbb{K} configuration of the language. To handle that, we have created an interactive tool for generating such formulas. Users can annotate their programs with preconditions and postconditions and then use our tool to generate RL formulas from those annotations. The above specification for KMP is generated from the annotations:

```
//@pre: m >= 1 /\ n >= 1
kmp_matcher(p, t, m, n);
//@post: allOcc(Out, p, t, n)
```

Loops can be annotated with invariants as shown below:

```
while (COND) {
  //@inv: INV
  k = pi[k];
}
```

For each annotated loop, the tool generates a RL formula which states that by starting with a configuration where the entire loop remains to be executed and `INV` holds, one reaches a configuration where the loop was completely executed and `INV \wedge \neg COND` holds.

From the annotations shown in Figure 5.7 the tool generates all the RL formulas that we need to prove KMP. Since KMP has four loops and two pairs of pre/post-conditions, the tool generates and proves a total number of six RL formulas. In the annotations we use the program variables (e.g. `pi`, `p`, `m`) and a special variable `Out` which is meant to refer the content of the $\langle \rangle_{\text{out}}$ cell. This variable gives us access to the output cell, which is essential in proving that the algorithm computes all the occurrences of the pattern.

Finally, since every particular verification problem requires problem-specific constructions and properties about them, for verifying KMP we have enriched the symbolic definition of CinK with functional symbols for π , θ , and `allOcc`, and some of their properties shown in Figure 5.8.

1. $0 \leq k \leq m \vdash 0 \leq \pi(k) < k$.
2. $0 \leq q \leq n \vdash 0 \leq \theta(q) \leq m$.
3. $(\forall u : 1..k)(P[u] = P[q - k + u]) \wedge \pi(q) \leq k + 1 \wedge P[k + 1] \neq P[q] \vdash \pi(q) \leq \pi(k) + 1$.
4. $(\forall u : 1..k)(P[u] = P[q - k + u]) \wedge \pi(q) \leq k + 1 \wedge P[k + 1] = P[q] \vdash \pi(i) = k + 1$.
5. $(\forall u : 1..q)(P[u] = T[i - 1 - q + u]) \wedge \theta(i) \leq q + 1 \wedge P[q + 1] \neq T[i] \vdash \theta(i) \leq \pi(q) + 1$.
6. $(\forall u : 1..q)(P[u] = T[i - 1 - q + u]) \wedge \theta(i) \leq q + 1 \wedge P[q + 1] = T[i] \vdash \theta(i) = q + 1$.
7. $(\exists v)(\forall u : v+1..i-1)(allOcc(Out, P, T, v) \wedge \theta(u) < m) \wedge \theta(i) = m \wedge i < n \vdash (\exists v)(\forall u : v+1..i)(allOcc(Out, P, T, v) \wedge \theta(u) < m)$.
8. $(\exists v)(\forall u : v+1..i)(allOcc(Out, P, T, v) \wedge \theta(u) < m) \wedge i = n \vdash allOcc(Out, P, T, v)$.

Figure 5.8: Properties needed to prove KMP.

Chapter 6

Conclusions

We have presented a formal and generic framework for the symbolic execution of programs in languages definable in an algebraic and term-rewriting setting. Symbolic execution is performed by applying the rules in the semantics of a language by so-called symbolic rewriting. We prove that the symbolic execution thus defined has the naturally expected properties with respect to concrete execution: *coverage*, meaning that to each concrete execution there is a feasible symbolic one on the same program path, and *precision*, meaning that each feasible symbolic execution has a concrete execution on the program same path. These properties are expressed in terms of mutual simulations.

In order to implement our symbolic rewriting-based approach in a setting where only standard execution is available, such as the \mathbb{K} framework, we define a transformation of language definitions \mathcal{L} into other language definitions \mathcal{L}^s , and show that concrete program execution \mathcal{L}^s , which uses standard rewriting, are in a covering&precise relationship with the symbolic execution of the corresponding programs in \mathcal{L} .

The incorporation of symbolic execution into a deductive system for program verification with respect to Reachability-Logic specifications is also presented in detail. Reachability Logic has a sound and relatively complete deduction system, which offers a lot of freedom (but very few guidelines) for constructing proofs. We came up with an alternative proof system, and we show that, under reasonable conditions on the semantics of programming languages and of the Reachability Logic formulas, a cer-

tain strategy executing our proof system is sound and weakly complete. This essentially means that, when it terminates, the strategy solves the Reachability-Logic verification problem: when presented with a valid input (set of RL formulas) it proves the formulas, and when presented with an invalid input it detects this invalidity.

Finally, we present the implementation of a prototype tool for symbolic execution based on the above theory, which is now a part of the \mathbb{K} framework. We illustrate the tool's capabilities on several programs written in different programming languages, in order to emphasise that our framework is not language dependent. Moreover, we show that the tool can be easily extended to perform other types of analyses, e.g. bounded model checking, LTL model-checking, etc. We also introduce a prototype for Reachability-Logic verification based on our proof system, which is implemented on top of our symbolic execution. We used this prototype by proving several non-trivial programs.

The work presented in this dissertation is important for at least three reasons: first, we formalise and prove two important properties of symbolic execution which allows one to consistently transfer properties of symbolic execution to concrete executions; second, we have shown that this framework is suitable for language independent program verification; finally, we provide an implementation which shows that language independent symbolic execution based on formal semantics is possible.

6.1 Future Work

We are planning to expand our tool, to make it able to seamlessly perform a wide range of program analyses, from testing and debugging to formal verifications, following ideas presented in related work, but with the added value of being language independent and grounded in formal methods. For this, we shall develop a rich domain of symbolic values, able to handle various kinds of data types. Formalising the interaction of symbolic-domain computations with symbolic execution is also a matter for future work.

Another future research direction is specifically targeted at our RL-formulas verifier, and aims at certifying its executions. The idea is to generate proof scripts for the Coq proof assistant [73], in order to obtain

certificates that, despite any (inevitable) bugs in our tool, the proofs it generates are indeed correct. This amounts to, firstly, encoding our RL proof system in Coq, and proving its soundness with respect to the original proof system of RL (which have already been proved sound in Coq [93]). Secondly, our verifier must be enhanced to return, for any successful execution, the rules of our system it has applied and the substitutions it has used. From this information a Coq script is built that, if successfully run by Coq, generates a proof term that constitutes a correctness certificate for the verifier's original execution. A longer-term objective is to turn our verifier into an external proof tactic for Coq, resulting in a powerful mixed interactive/automatic program verification tool.

Bibliography

- [1] W. Ahrendt. The KeY tool. *Software and Systems Modeling*, 4:32–54, 2005.
- [2] Krzysztof R. Apt, Frank de Boer, and Ernst-Rüdiger Olderog. *Verification of Sequential and Concurrent Programs*. Springer Verlag, 3rd edition, 2009.
- [3] Alessandro Armando, Massimo Benerecetti, and Jacopo Mantovani. Model checking linear programs with arrays. In *Proceedings of the Workshop on Software Model Checking*, volume 144 - 3, pages 79 – 94, 2006.
- [4] Andrei Arusoai. Engineering hoare logic-based program verification in K framework. In *SYNASC*, pages 177–184, 2013.
- [5] Andrei Arusoai, Dorel Lucanu, and Vlad Rusu. A generic framework for symbolic execution. In *6th International Conference on Software Language Engineering*, volume 8225 of *LNCS*, pages 281–301. Springer Verlag, 2013. Also available as a technical report <http://hal.inria.fr/hal-00853588>.
- [6] Andrei Arusoai, Dorel Lucanu, and Vlad Rusu. Towards a semantics for OCL. *Electronic Notes in Theoretical Computer Science*, 304(0):81 – 96, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).
- [7] Andrei Arusoai, Traian Florin Șerbănuță, Chucky Ellison, and Grigore Roșu. Making Maude definitions more interactive. In

- Rewriting Logic and Its Applications, WRLA 2012*, volume 7571 of *Lecture Notes in Computer Science*. Springer, 2012.
- [8] Irina Mariuca Asavoae. *K Semantics for Abstractions*. PhD thesis, Alexandru Ioan Cuza, Univerisity of Iași, September 2012. PhD Thesis.
 - [9] Irina Mariuca Asavoae and Mihail Asavoae. Collecting semantics under predicate abstraction in the K framework. In *WRLA*, pages 123–139, 2010.
 - [10] Irina Mariuca Asavoae, Mihail Asavoae, and Dorel Lucanu. Path directed symbolic execution in the K framework. In *SYNASC*, pages 133–141, 2010.
 - [11] Mihail Asavoae. *Semantics-Based WCET Analysis*. PhD thesis, Alexandru Ioan Cuza, Univerisity of Iași, September 2012. PhD Thesis.
 - [12] Mihail Asavoae, Irina Mariuca Asavoae, and Dorel Lucanu. On abstractions for timing analysis in the K framework. In *Proceedings of the Second International Conference on Foundational and Practical Aspects of Resource Analysis, FOPARA’11*, pages 90–107, Berlin, Heidelberg, 2012. Springer-Verlag.
 - [13] Irina Măriuca Asăvoae. Abstract semantics for alias analysis in k. *Electronic Notes in Theoretical Computer Science*, 304(0):97 – 110, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).
 - [14] Mihail Asăvoae. K semantics for assembly languages: A case study. *Electronic Notes in Theoretical Computer Science*, 304(0):111 – 125, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).
 - [15] Mike Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: an overview. In *Proc. 2004 international conference on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices, CASSIS’04*, pages 49–69, 2005.

- [16] Patrick Baudin, Jean C. Filliâtre, Thierry Hubert, Claude Marché, Benjamin Monate, Yannick Moy, and Virgile Prevosto. *ACSL: ANSI C Specification Language (preliminary design V1.2)*, preliminary edition, May 2008.
- [17] Denis Bogdănaş. Java semantics in \mathbb{K} . <https://github.com/kframework/java-semantics>.
- [18] Sylvie Boldo and Jean-Christophe Filliâtre. Formal verification of floating-point programs. In *IEEE Symposium on Computer Arithmetic*, pages 187–194. IEEE Computer Society, 2007.
- [19] 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*, pages 234–245, New York, NY, USA, 1975. ACM.
- [20] Stefan Bucur, Johannes Kinder, and George Candea. Prototyping symbolic execution engines for interpreted languages. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '14*, pages 239–254, New York, NY, USA, 2014. ACM.
- [21] Stefan Bucur, Johannes Kinder, and George Candea. Prototyping symbolic execution engines for interpreted languages. *SIGPLAN Not.*, 49(4):239–254, February 2014.
- [22] Jacob Burnim and Koushik Sen. Heuristics for scalable dynamic test generation. In *ASE*, pages 443–446. IEEE, 2008.
- [23] 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*, pages 209–224, Berkeley, CA, USA, 2008. USENIX Association.
- [24] Cristian Cadar, Vijay Ganesh, Peter M. Pawlowski, David L. Dill, and Dawson R. Engler. EXE: Automatically generating inputs of death. *ACM Trans. Inf. Syst. Secur.*, 12(2), 2008.

- [25] Vitaly Chipounov, Volodymyr Kuznetsov, and George Candea. S2E: A platform for in-vivo multi-path analysis of software systems. *SIGPLAN Not.*, 47(4):265–278, March 2011.
- [26] Vitaly Chipounov, Volodymyr Kuznetsov, and George Candea. The S2E platform: Design, implementation, and applications. *ACM Trans. Comput. Syst.*, 30(1):2:1–2:49, February 2012.
- [27] Edmund Clarke and Daniel Kroening. Hardware verification using ANSI-C programs as a reference. In *Proceedings of the 2003 Asia and South Pacific Design Automation Conference, ASP-DAC '03*, pages 308–311, New York, NY, USA, 2003. ACM.
- [28] Lori A. Clarke. A program testing system. In *Proceedings of the 1976 Annual Conference*, ACM '76, pages 488–491, New York, NY, USA, 1976. ACM.
- [29] Alberto Coen-Porisini, Giovanni Denaro, Carlo Ghezzi, and Mauro Pezzé. Using symbolic execution for verifying safety-critical systems. *SIGSOFT Softw. Eng. Notes*, 26(5):142–151, 2001.
- [30] Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition, 2001.
- [31] Heming Cui, Gang Hu, Jingyue Wu, and Junfeng Yang. Verifying systems rules using rule-directed symbolic execution. In *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '13*, pages 329–342, New York, NY, USA, 2013. ACM.
- [32] Pascal Cuoq, Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles, and Boris Yakobowski. Frama-C: A software analysis perspective. In *Proceedings of the 10th International Conference on Software Engineering and Formal Methods, SEFM'12*, pages 233–247, Berlin, Heidelberg, 2012. Springer-Verlag.
- [33] Jakub Daniel, Pavel Parizek, and Corina S. Pasareanu. Predicate abstraction in Java Pathfinder. *ACM SIGSOFT Software Engineering Notes*, 39(1):1–5, 2014.

- [34] Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS'08*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
- [35] Chucky Ellison and Grigore Roşu. An executable formal semantics of C with applications. In *Proceedings of the 39th Symposium on Principles of Programming Languages (POPL'12)*, pages 533–544. ACM, 2012.
- [36] Chucky Ellison, Traian Florin Şerbănuţă, and Grigore Roşu. A rewriting logic approach to type inference. In *Recent Trends in Algebraic Development Techniques*, volume 5486 of *Lecture Notes in Computer Science*, pages 135–151. Springer, 2009. Revised Selected Papers from the 19th International Workshop on Algebraic Development Techniques (WADT'08).
- [37] Santiago Escobar, José Meseguer, and Ralf Sasse. Variant narrowing and equational unification. *Electr. Notes Theor. Comput. Sci.*, 238(3):103–119, 2009.
- [38] PEX: Automated exploratory testing for .NET.
- [39] Daniele Filaretti and Sergio Maffei. An executable formal semantics of PHP. In *Proceedings of European Conference on Object-Oriented Programming*, 2014.
- [40] Jean Christophe Filliâtre. Proof of KMP string searching algorithm. <http://toccata.lri.fr/gallery/kmp.en.html>.
- [41] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 — where programs meet provers. In Matthias Felleisen and Philippa Gardner, editors, *Proceedings of the 22nd European Symposium on Programming*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128. Springer, March 2013.
- [42] Robert W. Floyd. Assigning meanings to programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages

- 19–32, Providence, Rhode Island, 1967. American Mathematical Society.
- [43] Werner Gabrisch and Wolf Zimmermann. A Hoare-style verification calculus for control state ASMs. In *Proceedings of the Fifth Balkan Conference in Informatics*, BCI '12, pages 205–210, New York, NY, USA, 2012. ACM.
- [44] Milos Gligoric, Darko Marinov, and Sam Kamin. Codese: Fast deserialization via code generation. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis*, ISSSTA '11, pages 298–308, New York, NY, USA, 2011. ACM.
- [45] Patrice Godefroid, Nils Klarlund, and Koushik Sen. DART: directed automated random testing. In *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '05, pages 213–223, New York, NY, USA, 2005. ACM.
- [46] Patrice Godefroid, Michael Y. Levin, and David A. Molnar. Sage: whitebox fuzzing for security testing. *Commun. ACM*, 55(3):40–44, 2012.
- [47] Joseph A. Goguen and José Meseguer. Order-sorted algebra I: Equational deduction for multiple inheritance, overloading, exceptions and partial operations. *Theor. Comput. Sci.*, 105(2):217–273, 1992.
- [48] Adele Goldberg and David Robson. *Smalltalk-80: The Language and Its Implementation*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1983.
- [49] Mike Gordon and Hélène Collavizza. Forward with Hoare. In A. W. Roscoe, Cliff B. Jones, and Kenneth R. Wood, editors, *Reflections on the Work of C.A.R. Hoare*, History of Computing, chapter 5, pages 101–121. Springer London, London, 2010.
- [50] K. Goshi, P. Wray, and Yong Sun. An intelligent tutoring system for teaching and learning Hoare logic. In *Computational Intelligence and Multimedia Applications, 2001. ICCIMA 2001. Proceedings. Fourth International Conference on*, pages 293–297, 2001.

- [51] Dwight Guth. A formal semantics of Python 3.3. Master's thesis, University of Illinois at Urbana-Champaign, July 2013.
- [52] David Harel, Dexter Kozen, and Jerzy Tiuryn. Dynamic logic. In *Handbook of Philosophical Logic*, pages 497–604. MIT Press, 1984.
- [53] Mark Hills and Grigore Roşu. KOOL: An Application of Rewriting Logic to Language Prototyping and Analysis. In *Proceedings of the 18th International Conference on Rewriting Techniques and Applications (RTA '07)*, volume 4533 of *LNCIS*, pages 246–256. Springer, 2007.
- [54] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969.
- [55] Gerard J. Holzmann. *The SPIN Model Checker - primer and reference manual*. Addison-Wesley, 2004.
- [56] W.E. Howden. Symbolic testing and the dissect symbolic evaluation system. *Software Engineering, IEEE Transactions on*, SE-3(4):266–278, July 1977.
- [57] ISO-IEC. Standard for Programming Language C++. Working Draft. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3797.pdf>.
- [58] J-Algo. The algorithm visualization tool, 2007.
- [59] Bart Jacobs, Jan Smans, and Frank Piessens. A quick tour of the verifast program verifier. In *Proceedings of the 8th Asian conference on Programming languages and systems, APLAS'10*, pages 304–311, Berlin, Heidelberg, 2010. Springer-Verlag.
- [60] Joxan Jaffar, Vijayaraghavan Murali, Jorge A. Navas, and Andrew E. Santosa. Tracer: a symbolic execution tool for verification. In *Proc. 24th international conference on Computer Aided Verification, CAV'12*, pages 758–766. Springer-Verlag, 2012.
- [61] Sarfraz Khurshid, Corina S. Pasareanu, and Willem Visser. Generalized symbolic execution for model checking and testing. In Hubert

- Garavel and John Hatcliff, editors, *TACAS*, volume 2619 of *Lecture Notes in Computer Science*, pages 553–568. Springer, 2003.
- [62] James C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976.
- [63] Donald E. Knuth, J.H. Morris, and Vaughan R. Pratt. Fast pattern matching in strings. *SIAM Journal of Computing*, 6(2):323–350, 1977.
- [64] Volodymyr Kuznetsov, Johannes Kinder, Stefan Bucur, and George Candea. Efficient state merging in symbolic execution. *SIGPLAN Not.*, 47(6):193–204, June 2012.
- [65] Chris Lattner and Vikram Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO'04)*, Palo Alto, California, Mar 2004.
- [66] David Lazar. K definition of Haskell, 2013.
- [67] David Lazar, Andrei Arusoai, Traian Florin Serbanuta, Chucky Ellison, Radu Mereuta, Dorel Lucanu, and Grigore Roşu. Executing formal semantics with the K tool. In *FM*, volume 7436 of *LNCS*, pages 267–271. Springer, 2012.
- [68] David Lazar and Chucky Ellison. K definition of the llvm assembly language, 2012.
- [69] Dorel Lucanu and Vlad Rusu. Program equivalence by circular reasoning. In *IFM*, Lecture Notes in Computer Science, pages 362–377. Springer, 2013.
- [70] Dorel Lucanu and Traian Florin Serbanuta. CinK - an exercise on how to think in K. Technical Report TR 12-03, Version 2, Alexandru Ioan Cuza University, Faculty of Computer Science, December 2013.
- [71] Dorel Lucanu, Traian Florin Şerbănuţă, and Grigore Roşu. The K Framework distilled. In *9th International Workshop on Rewriting Logic and its Applications*, volume 7571 of *Lecture Notes in Computer Science*, pages 31–53. Springer, 2012. Invited talk.

- [72] Kin-Keung Ma, Khoo Yit Phang, Jeffrey S. Foster, and Michael Hicks. Directed symbolic execution. In *Proceedings of the 18th International Conference on Static Analysis, SAS'11*, pages 95–111, Berlin, Heidelberg, 2011. Springer-Verlag.
- [73] Reference Manual. The Coq proof assistant. <http://coq.inria.fr/refman/>.
- [74] The \mathbb{K} tool. Github repository. <https://github.com/kframework/k>.
- [75] Daniel D. McCracken and Edwin D. Reilly. Backus-aur form (bnf). In *Encyclopedia of Computer Science*, pages 129–131. John Wiley and Sons Ltd., Chichester, UK, 2003.
- [76] Patrick O’Neil Meredith, Michael Katelman, José Meseguer, and Grigore Roşu. A formal executable semantics of Verilog. In *Eighth ACM/IEEE International Conference on Formal Methods and Models for Codesign (MEMOCODE’10)*, pages 179–188. IEEE, 2010.
- [77] José Meseguer. Rewriting logic and Maude: Concepts and applications. In Leo Bachmair, editor, *RTA*, volume 1833 of *LNCS*, pages 1–26. Springer, 2000.
- [78] José Meseguer and Prasanna Thati. Symbolic reachability analysis using narrowing and its application to verification of cryptographic protocols. *Higher-Order and Symbolic Computation*, 20(1-2):123–160, 2007.
- [79] José Meseguer. Rewriting as a unified model of concurrency. In J.C.M. Baeten and J.W. Klop, editors, *CONCUR ’90 Theories of Concurrency: Unification and Extension*, volume 458 of *Lecture Notes in Computer Science*, pages 384–400. Springer Berlin Heidelberg, 1990.
- [80] Daejun Park. K definition of Javascript, 2013.
- [81] Corina S. Pasareanu, Peter C. Mehlitz, David H. Bushnell, Karen Gundy-Burlet, Michael R. Lowry, Suzette Person, and Mark Pape. Combining unit-level symbolic execution and system-level concrete

- execution for testing nasa software. In Barbara G. Ryder and Andreas Zeller, editors, *ISSTA*, pages 15–26. ACM, 2008.
- [82] Corina S. Păsăreanu and Neha Rungta. Symbolic pathfinder: symbolic execution of Java bytecode. In *ASE*, pages 179–180, 2010.
- [83] Corina S. Pasareanu and Willem Visser. Verification of java programs using symbolic execution and invariant generation. In Susanne Graf and Laurent Mounier, editors, *SPIN*, volume 2989 of *Lecture Notes in Computer Science*, pages 164–181. Springer, 2004.
- [84] Corina S. Pasareanu, Willem Visser, David H. Bushnell, Jaco Geldenhuys, Peter C. Mehltz, and Neha Rungta. Symbolic pathfinder: integrating symbolic execution with model checking for java bytecode analysis. *Autom. Softw. Eng.*, 20(3):391–425, 2013.
- [85] Grigore Roşu Patrick Meredith, Mark Hills. An Executable Rewriting Logic Semantics of K-Scheme. In Danny Dube, editor, *Proceedings of the 2007 Workshop on Scheme and Functional Programming (SCHEME’07)*, Technical Report DIUL-RT-0701, pages 91–103. Laval University, 2007.
- [86] Suzette Person, Guowei Yang, Neha Rungta, and Sarfraz Khurshid. Directed incremental symbolic execution. In *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’11*, pages 504–515, New York, NY, USA, 2011. ACM.
- [87] David A Ramos and Dawson R. Engler. Practical, low-effort equivalence verification of real code. In *Proceedings of the 23rd international conference on Computer aided verification, CAV’11*, pages 669–685, Berlin, Heidelberg, 2011. Springer-Verlag.
- [88] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science, LICS ’02*, pages 55–74, Washington, DC, USA, 2002. IEEE Computer Society.
- [89] Camilo Rocha, José Meseguer, and Cesar A. Munoz. Rewriting modulo SMT. Technical Report 218033, NASA, 2013.

- [90] Camilo Rocha, José Meseguer, and Cesar A. Munoz. Rewriting modulo smt and open system analysis. In *Proceedings of the 10th International Workshop on Rewriting Techniques and Applications*, WRLA'14. LNCS, 2014.
- [91] Grigore Roşu. CS322, Fall 2003 - Programming Language Design: Lecture Notes. Technical Report UIUCDCS-R-2003-2897, University of Illinois at Urbana-Champaign, Department of Computer Science, December 2003. Lecture notes of a course taught at UIUC.
- [92] Grigore Roşu and Andrei Ştefănescu. Matching logic rewriting: Unifying operational and axiomatic semantics in a practical and generic framework. Technical Report <http://hdl.handle.net/2142/28357>, University of Illinois, November 2011.
- [93] Grigore Roşu, Andrei Ştefănescu, Ştefan Ciobâcă, and Brandon M. Moore. One-path reachability logic. In *Proceedings of the 28th Symposium on Logic in Computer Science (LICS'13)*, pages 358–367. IEEE, June 2013.
- [94] Grigore Roşu, Chucky Ellison, and Wolfram Schulte. Matching logic: An alternative to Hoare/Floyd logic. In Michael Johnson and Dusko Pavlovic, editors, *Proceedings of the 13th International Conference on Algebraic Methodology And Software Technology (AMAST '10)*, volume 6486 of *Lecture Notes in Computer Science*, pages 142–162, 2010.
- [95] Grigore Roşu and Dorel Lucanu. Circular coinduction – a proof theoretical foundation. In *CALCO 2009*, volume 5728 of *LNCS*, pages 127–144. Springer, 2009.
- [96] Grigore Rosu, Wolfram Schulte, and Traian Florin Serbanuta. Runtime verification of C memory safety. In Saddek Bensalem and Doron A. Peled, editors, *Runtime Verification (RV'09)*, volume 5779 of *Lecture Notes in Computer Science*, pages 132–152, 2009.
- [97] Grigore Roşu and Traian Florin Şerbănuţă. An overview of the K semantic framework. *Journal of Logic and Algebraic Programming*, 79(6):397–434, 2010.

- [98] Grigore Rosu and Andrei Stefanescu. Checking reachability using matching logic. In *Proceedings of the 27th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '12)*, pages 555–574. ACM, 2012.
- [99] Grigore Roşu and Andrei Ştefănescu. Checking reachability using matching logic. Technical Report <http://hdl.handle.net/2142/33771>, UIUC, August 2012.
- [100] Grigore Rosu and Andrei Stefanescu. From hoare logic to matching logic reachability. In *Proceedings of the 18th International Symposium on Formal Methods (FM'12)*, volume 7436 of *Lecture Notes in Computer Science*, pages 387–402. Springer, 2012.
- [101] Grigore Rosu and Andrei Stefanescu. Towards a unified theory of operational and axiomatic semantics. In *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP'12)*, volume 7392 of *Lecture Notes in Computer Science*, pages 351–363. Springer, 2012.
- [102] Grigore Roşu and Traian Florin Şerbănuţă. Overview and SIMPLE case study. *Electronic Notes in Theoretical Computer Science*, 304(0):3 – 56, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).
- [103] Samir Sapra, Marius Minea, Sagar Chaki, Arie Gurfinkel, and Edmund M. Clarke. Finding errors in python programs using dynamic symbolic execution. In *ICTSS*, pages 283–289, 2013.
- [104] Ralf Sasse and José Meseguer. Java+ITP: A verification tool based on Hoare Logic and Algebraic Semantics. *Electron. Notes Theor. Comput. Sci.*, 176(4):29–46, July 2007.
- [105] Prateek Saxena, Pongsin Poosankam, Stephen McCamant, and Dawn Song. Loop-extended symbolic execution on binary programs. In *Proceedings of the ACM/SIGSOFT International Symposium on Software Testing and Analysis (ISSTA)*, Chicago, IL, July 2009.
- [106] Koushik Sen and Gul Agha. CUTE and jCUTE: Concolic unit testing and explicit path model-checking tools. In Thomas Ball and

- Robert B. Jones, editors, *CAV*, volume 4144 of *Lecture Notes in Computer Science*, pages 419–423. Springer, 2006.
- [107] Koushik Sen, Darko Marinov, and Gul Agha. CUTE: a concolic unit testing engine for C. 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*, pages 263–272, New York, NY, USA, 2005. ACM.
- [108] Traian Florin Şerbănuță. *A Rewriting Approach to Concurrent Programming Language Design and Semantics*. PhD thesis, University of Illinois at Urbana-Champaign, December 2010. <https://www.ideals.illinois.edu/handle/2142/18252>.
- [109] Traian Florin Serbanuta, Andrei Arusoai, David Lazar, Chucky Ellison, Dorel Lucanu, and Grigore Rosu. The k primer (version 3.3). In *Proceedings of International K Workshop (K'11)*, ENTCS. Elsevier, 2013. To appear.
- [110] Traian Florin Serbanuta and Grigore Rosu. A truly concurrent semantics for the k framework based on graph transformations. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski, and Grzegorz Rozenberg, editors, *ICGT*, volume 7562 of *Lecture Notes in Computer Science*, pages 294–310. Springer, 2012.
- [111] Traian Florin Şerbănuță, Grigore Roşu, and José Meseguer. A rewriting logic approach to operational semantics. *Information and Computation*, 207:305–340, 2009.
- [112] Cristian Zamfir and George Candea. Execution synthesis: A technique for automated software debugging. In *Proceedings of the 5th European Conference on Computer Systems, EuroSys '10*, pages 321–334, New York, NY, USA, 2010. ACM.
- [113] Traian Florin Şerbănuță. Rewriting semantics and analysis of concurrency features for a C-like language. *Electronic Notes in Theoretical Computer Science*, 304(0):167 – 182, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).

- [114] Traian Florin Șerbănuță, Andrei Arusoaiie, David Lazar, Chucky Ellison, Dorel Lucanu, and Grigore Roșu. The primer (version 3.3). *Electronic Notes in Theoretical Computer Science*, 304(0):57 – 80, 2014. Proceedings of the Second International Workshop on the K Framework and its Applications (K 2011).