The [Tr $\Vdash_{\alpha}$ ce] Modality*

Dominic Steinhöfel[0000−0003−4439−7129] and Reiner Hähnle[0000−0001−8000−7613]

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

Abstract. We propose the *trace modality*, a concept to uniformly express a wide range of program verification problems. To demonstrate its usefulness, we formalize several program verification problems in it: Functional Verification, Information Flow Analysis, Temporal Model Checking, Program Synthesis, Correct Compilation, and Program Evolution. To reason about the trace modality, we translate programs *and* specifications to *regular symbolic traces* and construct simulation relations on first-order symbolic automata. The idea with this uniform representation is that it helps to identify synergy potential—theoretically and practically—between so far separate verification approaches.

1 Introduction

Since the foundational work on program verification during the 1960s [19], the program verification tasks that were studied have very much broadened beyond mere functional (partial or total) correctness. Basic variations include termination [16], reachability [27], and program synthesis [18]. Starting in the early 2000s, verification of *relational* properties of programs, such as information flow [10], correct compilation [23], or correctness of program transformations (refactoring) [14] has been in the focus of interest. Relational properties compare two programs having identical or similar behavior. It is even more challenging to reason about programs having related, but intentionally *differing* behavior, such as in program evolution [15].

For all these tasks *dedicated* verification approaches were developed: dynamic logic [17], Hoare quadruples [33], self composition [4,10], product programs [3], etc. Usually, the verification problem to be solved is stated informally, and then the problem is *directly* formalized in the approach to be used for its solution. Hence, the formalism that a problem is stated in and the formalism where it is solved, are *conflated*. We consider this problematic for two reasons:

(1) **Premature commitment to a specific solution approach.** If one has invested to master a specific methodology, the temptation to solve *any* problem by modifying or extending the familiar is considerable, even if a different approach would have been more efficient, flexible, or easily extensible: The well-known "for a hammer the world consists of nails" effect.

* This work was funded by the Hessian LOEWE initiative within the Software-Factory 4.0 project.

(2) **Hard to detect commonalities and to transfer results.** One of the most powerful scientific stratagems is to detect structural similarity among different problem areas. This makes it possible to transfer insights and solutions from one problem space to another. In formal verification, this additionally opens the road to re-use of software tools for new tasks. To be able to spot commonalities, it is essential to know which aspects of a problem are genuinely new and hence require a novel approach. However, if a problem is *already formalized* in terms of a specific solution method, it is hard to identify commonality and analogy.

Experience with various software verification problems [16,10,1,31] let us realize that a small number of principles occur time and again in dedicated verification approaches: (1) *abstraction* of program runs in the sense of abstract interpretation [9]; (2) *approximation* of a set of program runs by a superset; (3) the capability to handle *schematic* programs, i.e. programs with unspecified parts. Abstraction makes it possible to compare programs written in different languages via a suitable abstraction of their traces. Approximation is needed to focus on a specific property and "forget" irrelevant information.Finally, to reason about program transformation (synthesis, compilation, refactoring, etc.) it is essential to be able to specify a program fragment in an unknown context. We propose a framework based on these principles that lets one express a wide variety of verification problems in a uniform, comparable manner.

We make only one assumption about the programs under verification: they must have a *trace semantics*, i.e. for an initial execution state s and a program p we can obtain the set of all traces ("program runs") that are possible when p is started in s. Our framework builds on the semantic notion of a *trace modality* $[\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$, and a reasoning system based on *regular symbolic traces* and simulations on symbolic automata. The expression $[\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$ is *valid* if the traces arising from the *implementation* $\mathcal{C}_l$ are *approximated* by the traces of the *specification* $\mathcal{C}_r$ after the *abstraction* step defined by α, where implementation and specification may be either programs (potentially containing *abstract contexts*) or formulas (e.g., in first-order or temporal logic). Symbolic traces approximate concrete traces. Our reasoning system translates implementation and specification to symbolic traces, transforms these to symbolic automata, and finally shows language inclusion by constructing a simulation relation "up to subsumption".

The paper is structured as follows. Sect. 2 defines elementary notions used in the paper. The semantics of the trace modality is described in Sect. 3, where we also formalize various verification tasks using the trace modality to demonstrate its expressiveness. In Sect. 4, we describe our reasoning system based on symbolic traces. Finally, Sect. 5 discusses related work, and Sect. 6 concludes the paper and describes future work opportunities. For space reasons, we moved some details of the paper, (detailed examples, longer remarks) to an appendix, available at `www.key-project.org/papers/trace-modality/`.

2 Programs, Logic, Traces and Abstractions

We assume an imperative programming language $\mathcal{L}$ with the usual sequencing and assignment operators ";", "=". Programs may contain *schematic* statements, denoted with capital letters $\mathtt{P}$, $\mathtt{Q}$, etc. A program without schematic statements is called *concrete*. The set of concrete programs is $\mathcal{L}_0$. A program $\mathtt{p}$ with schematic statements represents the set $\mathrm{Concr}(\mathtt{p})$ of all well-formed $\mathcal{L}_0$-programs obtained by replacing each schematic statement in $\mathtt{p}$ with an arbitrary concrete statement.

At each point during the execution of a program $\mathtt{p} \in \mathcal{L}_0$ it is in a *state* $s \in \mathcal{S}$, mapping program variables to domain values. To model failed assertions, we distinguish a state $\bot$. We write $\mathcal{S}^\bot$ for $\mathcal{S} \cup \{\bot\}$. A *trace* τ is a possibly infinite sequence of states, denoted $s_0 s_1 \cdots s_n$ or $s_0 s_1 \cdots$ (the latter being infinite). For the empty trace we write ε and $\mathit{Traces} = (\mathcal{S}^\bot)^* \cup (\mathcal{S}^\bot)^\omega$ for the set of all traces. Predicate $\mathit{finite}(\tau)$ holds for finite traces and $\mathit{first}(\tau)$, $\mathit{last}(\tau)$ select a trace's first and final state (the latter is only defined for finite traces).

We assume a *trace semantics* $\mathrm{Tr}_s(\mathtt{p})$ that maps a *concrete* program $\mathtt{p} \in \mathcal{L}_0$ and initial state $s \in \mathcal{S}$ into the set of possible traces when $\mathtt{p}$ is started in s. When $\mathcal{L}_0$ is deterministic, this is a singleton. We define the set of all traces of $\mathtt{p} \in \mathcal{L}_0$ as $\mathrm{Tr}(\mathtt{p}) = \{\mathrm{Tr}_s(\mathtt{p}) \mid s \in \mathcal{S}\}$. If $\mathcal{P}$ is a set of concrete programs, then $\mathrm{Tr}(\mathcal{P}) = \{\mathrm{Tr}(\mathtt{p}) \mid \mathtt{p} \in \mathcal{P}\}$, similar for $\mathrm{Tr}_s(\mathcal{P})$. Now we define the semantics of a *schematic* program $\mathtt{p} \in \mathcal{L}$ as $\mathrm{Tr}(\mathrm{Concr}(\mathtt{p}))$ and $\mathrm{Tr}_s(\mathrm{Concr}(\mathtt{p}))$, respectively.

Let PVar denote the set of program variables and "$\circ$" the usual function composition operator. We define *abstraction operators* $\alpha : 2^{\mathit{Traces}} \to 2^{\mathit{Traces}}$ (in the sense of abstract interpretation [9]) on sets of traces $\mathcal{T}$:

Big-step abstraction: $\alpha_{big}(\mathcal{T}) = \{(s_0, s_n) \mid s_0 \cdots s_n \in \mathcal{T}\}$, i.e. the set of all pairs of the first and last state of any finite trace in $\mathcal{T}$. Observe that for infinite traces in $\mathcal{T}$, there is no corresponding pair in the abstracted set.

Observation abstraction: Let $\mathrm{obs} \subseteq \mathrm{PVar}$, $s \in \mathcal{S}$, then $s \downarrow \mathrm{obs}$ is the state s restricted to values from obs. We define the *observation abstraction* relative to obs as $\alpha_{\mathrm{obs}}(\mathcal{T}) = \{(s_0 \downarrow \mathrm{obs})(s_1 \downarrow \mathrm{obs}) \cdots \mid s_0 s_1 \cdots \in \mathcal{T}\}$. For a concrete set of variables, for instance $\{\mathtt{x}\}$, we write $\alpha_{\{\mathtt{x}\}}$.

Data abstraction: Let α_d be an abstract operator on data types in $\mathtt{p}$ [9]. We define the *data abstraction* of a set of traces as $\alpha_d(\mathcal{T}) = \{(\alpha_d(s_0)\alpha_d(s_1) \cdots \mid s_0 s_1 \cdots \in \mathcal{T}\}$, where the state $\alpha_d(s)(\mathtt{x}) = \alpha_d(s(\mathtt{x}))$ is defined pointwise.

Combination: Combine two abstractions α_1, α_2 by composition $\alpha_1 \circ \alpha_2$.

We use a standard first-order language with equality. It contains the usual propositional connectives and first-order quantifiers. Terms and formulas are standard, but we permit trace modality formulas $[\mathcal{T}_l \Vdash_\alpha \mathcal{T}_r]$ as atomic formulas. With Trm and Fml we denote the sets of all terms and formulas. The semantics of a formula is provided by a first-order structure K and a state $s \in \mathcal{S}$ that define the *validity* relation $K, s \models \varphi$ for each $\varphi \in \mathrm{Fml}$. For example, $K, s \models \varphi \to \psi$ iff *either* $K, s \not\models \varphi$, *or* $K, s \models \varphi$ and $K, s \models \psi$. Given $s \in \mathcal{S}$, write $s \models \varphi$ and say φ is *valid* for s if $K, s \models \varphi$ for all K. Write $\models \varphi$ and say that φ is *valid* if $s \models \varphi$ for all $s \in \mathcal{S}$. While K interprets *rigid* first-order functions and predicates, a state s

assigns values to program variables that may *change* during execution. For the failure state $\bot$ we set $K, \bot \not\models \varphi$ for *any* K, φ.

We need $\texttt{assert}(\varphi)$ and $\texttt{assume}(\varphi)$ statements for asserting and assuming a first-order formula φ *in a program*. Our use cases are *program synthesis* for `assert` and *invariant reasoning* (appendix) for `assume` statements. For `assert`, we define $\mathrm{Tr}_s(\texttt{assert}(\varphi)) = \{s\}$ if $s \models \varphi$ and $\{\bot\}$ otherwise. The semantics of `assume` is defined as $\mathrm{Tr}_s(\texttt{assume}(\varphi)) = \{s\}$ if $s \models \varphi$ and $\emptyset$ otherwise. We define a full trace semantics for a simple $\mathcal{L}_0$-language in the appendix.

3 The Trace Modality

We define the *trace modality* $[\mathcal{T}_l \Vdash_\alpha \mathcal{T}_r]$, where the *implementation* $\mathcal{T}_l$ and *specification* $\mathcal{T}_r$ both are (possibly infinite) trace sets and α is a trace abstraction. It expresses that the specification is an approximation of the implementation relative to α. Its semantics is that the modality is *valid*, written $\models [\mathcal{T}_l \Vdash_\alpha \mathcal{T}_r]$, if $\alpha(\mathcal{T}_l) \subseteq \alpha(\mathcal{T}_r)$. We use *lifting functions* lift_l, lift_r that convert elements from the verification domain, such as programs or formulas, to trace sets. Formally:

Definition 1. *Let $\alpha : 2^{\mathit{Traces}} \to 2^{\mathit{Traces}}$ be a trace abstraction and $\mathcal{C}_l$, $\mathcal{C}_r$ elements of domains D_l, D_r with associated lifting functions $\mathit{lift}_{l/r} : \mathcal{S} \to D_{l/r} \to 2^{\mathit{Traces}}$. Then the trace modality $[\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$ is valid in $s \in \mathcal{S}$, written $s \models [\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$, iff $\alpha(\mathit{lift}_l(s)(\mathcal{C}_l)) \subseteq \alpha(\mathit{lift}_r(s)(\mathcal{C}_r))$.*

For readability, we omit lifting functions in the presentation and assume that verification domains have fixed lifting functions. We omit α if it is the identity.

Relation to Modal and Dynamic Logic Like in modal and dynamic logic we can define a dual modality as follows: $\langle \mathcal{C}_l \Vdash_\alpha \mathcal{C}_r \rangle \coloneqq \neg[\mathcal{C}_l \Vdash_\alpha \overline{\mathcal{C}_r}]$, where $\overline{\mathcal{C}_r}$ is the complement of $\mathcal{C}_r$. The semantics of this *diamond* trace modality can be phrased as: $s \models \langle \mathcal{C}_l \Vdash_\alpha \mathcal{C}_r \rangle$ iff there is a trace in $\alpha(\mathit{lift}_l(s)(\mathcal{C}_l))$ that is not in $\alpha(\overline{\mathit{lift}_r(s)(\mathcal{C}_r)})$. The axioms **N** (necessitation rule) and **K** (distribution axiom) of modal logic follow from Def. 1; when using programs (of the domain $D_{\mathcal{L}_0}$) as implementation and formulas (of the domain D_{Fml}) as specification, applicable axioms of Propositional Dynamic Logic (PDL) [17] also hold. We refer to our discussion in the appendix (Remark 1) for these observations. Note that, despite those similarities, the trace modality is strictly more general than Dynamic Logic (DL). Our specifications can originate from different verification domains, like temporal logic; even a *program* can be a specification (we show several examples for this case later on). Also, we are not restricted to big-step reasoning. To emphasize these differences of the trace modality to standard DL, we chose the notation also encapsulating the specification inside the box.

In the following, several verification tasks are described and formalized with the trace modality. We define suitable verification domains D, lifting functions lift_D, as well as abstractions. A summary table is in the appendix.

3.1 Functional Verification

In functional verification one shows a given program $\mathtt{p} \in \mathcal{L}_0$ to satisfy a postcondition *Post* provided that a precondition *Pre* holds initially. The problem is frequently formalized with *Hoare triples* $\{Pre\}\mathtt{p}\{Post\}$ [19]. In DL [1,17], one writes $Pre \to [\mathtt{p}]Post$. We distinguish *partial correctness*, where *Post* is asserted to hold *if* p terminates, from *total correctness*, where it is also shown *that* p terminates. For the latter one can use the dual modality $\langle\mathtt{p}\rangle Post$.

Functional correctness is over the domain $D_{\mathcal{L}_0}$ for programs and D_{Fml} for first-order postconditions. We define lifting functions $lift_{\mathcal{L}_0}(s)(\mathtt{p}) \coloneqq \mathrm{Tr}_s(\mathtt{p})$ and $lift_{\mathrm{Fml}}(s)(\varphi) \coloneqq \{\tau \in Traces : finite(\tau) \wedge first(\tau) = s \wedge last(\tau) \models \varphi\}$ (the set of all traces starting in s whose final state satisfies φ). Using the big-step abstraction, we can formalize (partial) functional correctness as $Pre \to [\mathtt{p} \Vdash_{\alpha_{big}} Post]$. Total correctness for *deterministic* programs is expressed as $Pre \to \langle\mathtt{p} \Vdash_{\alpha_{big}} Post\rangle$.

Example 1. Let $\mathtt{p} \coloneqq$ `j=i*i; while (i<j) { i=i*2; }`. It diverges iff the initial value of `i` is negative. One can prove postcondition $even(\mathtt{i})$ (with the obvious meaning) for p *if it terminates* (partial correctness). Thus, $s \models [\mathtt{p} \Vdash_{\alpha_{big}} even(\mathtt{i})]$ must hold in all $s \in \mathcal{S}$, i.e. $\alpha_{big}(lift_{\mathcal{L}_0}(s)(\mathtt{p})) \subseteq \alpha_{big}(lift_{\mathrm{Fml}}(s)(even(\mathtt{i})))$. If $s(\mathtt{i}) < 0$ then the set $lift_{\mathcal{L}_0}(s)(\mathtt{p})$ contains a single infinite trace. Therefore, $\alpha_{big}(lift_{\mathcal{L}_0}(s)(\mathtt{p})) = \emptyset$ and the subset relation holds. If $s(\mathtt{i}) \geq 0$, p has a single finite trace whose final state assigns an even value to `i` (either because $s(\mathtt{i})$ is 0, or because it is greater than 0, and the initial value was multiplied by 2 a number of times in the loop's body). Hence, $\alpha_{big}(lift_{\mathcal{L}_0}(s)(\mathtt{p}))$ contains a single pair (s, s_f) where s_f satisfies $even(\mathtt{i})$. It is in $\alpha_{big}(lift_{\mathrm{Fml}}(s)(even(\mathtt{i})))$ by defn. of $lift_{\mathrm{Fml}}$. We cannot show $s \models \langle\mathtt{p} \Vdash_{\alpha_{big}} even(\mathtt{i})\rangle$ for any s with $s(\mathtt{i}) < 0$, because then $\alpha_{big}(lift_{\mathcal{L}_0}(s)(\mathtt{p}))$ is empty and so cannot contain a trace not in $\alpha_{big}(\overline{lift_{\mathrm{Fml}}(s)(even(\mathtt{i}))})$. However, $\models \mathtt{i} \geq 0 \to \langle\mathtt{p} \Vdash_{\alpha_{big}} even(\mathtt{i})\rangle$ is true. ◊

3.2 Information Flow Analysis

To prove that a given program treats secret inputs (for example, a password) confidentially, i.e. it does not inadvertently leak secret information, one can formally prove that it satisfies an *information flow policy*. In the simplest case such policies partition program variables into *low*-security variables that hold observable values and *high*-security ones whose values are secret. A policy imposes restrictions on the flow of values from *high* to *low* variables. A standard and very strong policy is *non-interference*: "Whenever two instances of the same program are run with equal *low* values and arbitrary *high* values then the resulting *low* values are equal in the final state". This ensures that an attacker cannot learn anything about secret values by running the program with observable values. For simplicity, assume a program p contains exactly one low variable `l` and one high variable `h`, written $\mathtt{p}(\mathtt{l},\mathtt{h})$. Using *self composition* [4,10], this is formalized as a Hoare triple: If we can prove $\{\mathtt{l} \doteq \mathtt{l}'\}\mathtt{p}(\mathtt{l},\mathtt{h});\mathtt{p}(\mathtt{l}',\mathtt{h}')\{\mathtt{l} \doteq \mathtt{l}'\}$, p satisfies non-interference. It can also be directly expressed with the trace modality: $\models [\mathtt{p}(\mathtt{l},\mathtt{h}) \Vdash_{\alpha_{\{1\}} \circ \alpha_{big}} \mathtt{p}(\mathtt{l},\mathtt{h}')]$. Note that the renaming of `l` to $\mathtt{l}'$ is then not

necessary since programs are not composed, but evaluated separately. In the appendix, we discuss how *declassification* can be encoded with the trace modality.

Example 2. Let $\mathtt{p} \coloneqq \mathtt{l=42;}$ **if** $\mathtt{(h>20)}$ $\mathtt{\{l=17;\}}$. This program does not satisfy non-interference, because the final value of the observable variable $\mathtt{l}$ depends on the initial value of $\mathtt{h}$. We prove that indeed, $\models [\mathtt{p}(\mathtt{l},\mathtt{h}) \Vdash_{\alpha_{\{\mathtt{l}\}} \circ \alpha_{big}} \mathtt{p}(\mathtt{l},\mathtt{h}')]$ does *not* hold, by showing that there is a state $s \in \mathcal{S}$ for which

$$(\alpha_{\{\mathtt{l}\}} \circ \alpha_{big})(\mathit{lift}_{\mathcal{L}_0}(s)(\mathtt{p}(\mathtt{l},\mathtt{h}))) \subseteq (\alpha_{\{\mathtt{l}\}} \circ \alpha_{big})(\mathit{lift}_{\mathcal{L}_0}(s)(\mathtt{p}(\mathtt{l},\mathtt{h}')))$$

is not true. Let s be such that $s(\mathtt{l}) = 0$, $s(\mathtt{h}) = 0$ and $s(\mathtt{h}') = 30$. Then the trace set of the implementation is $\{(\{\mathtt{l} \mapsto 0\}, \{\mathtt{l} \mapsto 42\})\}$ which is not contained in the set for the specification $\{(\{\mathtt{l} \mapsto 0\}, \{\mathtt{l} \mapsto 17\})\}$. ◊

3.3 Software Model Checking

Software Model Checking (SMC) [21] describes a wide range of techniques for analyzing *safety* or *liveness* properties of programs. Those techniques have in common that they focus on *automation* at cost of expressivity. Frequently, the goal is not to prove correctness relative to a specification, but rather to quickly uncover bugs or to generate high-coverage test cases. Recently, there has been a *convergence* between *model checking* and *deductive verification* techniques [29], as more mechanisms traditionally known from the latter field, such as abstraction [32], symbolic execution [25], etc., are integrated to achieve greater expressivity. On the other side, Bounded Model Checking (BMC) approaches, which limit state space exploration by a user-defined upper bound on loop unwindings, are well-known and successful, and finite space checkers such as SPIN [20] continue being used, e.g. in protocol verification. Properties of interest to SMC (e.g., the absence of memory faults) can usually be formalized in Temporal Logic (TL).

We introduce the domain D_{TL} for Linear Temporal Logic (LTL) formulas, lift_{TL} is the standard trace semantics for temporal logic (e.g., $\mathit{lift}_{TL}(s)(\Box p)$ is the set of all traces starting in s where p always holds). We exemplarily instantiate the trace modality to *Finite Space* MC. Finite space model checkers like SPIN [20] exhaustively explore the state space of an abstract program model. This implies that the analysis starts from *a concrete input state* s and that no unbounded data structures are involved. We can formalize this problem as $s \models [\mathtt{p} \Vdash \varphi]$, where φ is an LTL formula. In the appendix, we show how to instantiate the trace modality to Bounded MC, Abstraction-Based MC and Symbolic Execution-Based MC. Model Checking tools for bug finding can be formalized with the diamond trace modality: They eagerly try to show $\models \langle \mathtt{p} \Vdash \neg\varphi \rangle$, i.e. there is a trace of $\mathtt{p}$ violating φ. Such a trace constitutes a counterexample which can be used to fix the program, and/or to create a useful test case.

So far, we considered *concrete* programs $\mathtt{p} \in \mathcal{L}_0$. The two subsequently discussed verification tasks are over *schematic* programs in $\mathcal{L}$.

3.4 Program Synthesis

Automated program synthesis starts with a specification of programs at a higher level than executable code. The latter is created (semi-)automatically from the specification. In [30], for instance, the user supplies a *scaffold* consisting of a functional specification $(Pre, Post)$, domain constraints defining the domains of expressions and guards, and a *schematic program* (called "flowgraph template") of the form $\bullet|*(T)|T;T$. Here, $\bullet$ is an acyclic fragment, T again a schematic program and $*(T)$ a loop with body T. The synthesizer infers *synthesis conditions*. These are satisfiable whenever there exists a valid program for the scaffold.

We encode $\bullet$ of the flowgraph template by programs $\mathtt{p} \in \mathcal{L}$ with schematic statements $\mathtt{P}$, $\mathtt{Q}$, ..., and define a new verification domain $D_{\mathcal{L}}$ with $\mathit{lift}_{\mathcal{L}}(s)(\mathtt{p}) \coloneqq \mathrm{Tr}_s(\mathrm{Concr}(\mathtt{p}))$. Synthesis conditions are included in the intermediate program as suitable $\texttt{assert}(\varphi)$ statements. When refining an intermediate program $\mathtt{p}$ to a more concrete program $\mathtt{p}'$, the property to show is $\models \mathit{Pre} \to [\mathtt{p}' \Vdash_{\alpha_{big}} \mathtt{p}]$: that $\mathtt{p}'$ is indeed a refinement of $\mathtt{p}$. In the appendix, we provide an example for the synthesis of a program computing integer square roots.

3.5 Correct Compilation

A compiler translates a program $\mathtt{p}$ in a source language into a program $\mathtt{c}$ of a target language, preserving the behavior of $\mathtt{p}$. The translation can introduce new program variables. Then preservation of behavior is typically restricted to a set of *observable* variables *obs*. In *modular* compilation a program $\mathtt{p}$ is given within an unspecified context. In this case both $\mathtt{p}$ and $\mathtt{c}$ are abstract. Correctness of compilation can be expressed as $\models [\mathtt{p} \Vdash_{\alpha_{obs} \circ \alpha_{big}} \mathtt{c}]$. If we want to exclude behavior of $\mathtt{c}$ that cannot be observed in $\mathtt{p}$, we can—for deterministic languages—use the diamond modality instead. For non-deterministic languages, we can *additionally* prove the reverse direction $\models [\mathtt{c} \Vdash_{\alpha_{obs} \circ \alpha_{big}} \mathtt{p}]$ to achieve this.

The formalization makes the similarity to program synthesis explicit. Indeed, one could create a scaffold by extracting synthesis conditions from $\mathtt{p}$, and then try to infer $\mathtt{c}$ automatically. For example, in [31], a symbolic execution tree of the source program is "mined" to extract the target program. It is related to proof mining techniques used in program synthesis.

3.6 Program Evolution & Bug Fixing

Sometimes, the behavior of the "implementation" should intentionally be *not* preserved. This situation occurs in program evolution, e.g., after manual or automatic bug fixing [24]: the patched program is supposed to exhibit the bug no longer, but no new bug is to be introduced. Similarly, in fault propagation analysis, an injected fault typically *will* change the behaviour of a program, but not arbitrarily. This problem is most naturally expressed as $\models [\mathtt{p}_{bug} \Vdash_{\alpha_{bug} \circ \alpha_{big}} \mathtt{p}_{\mathit{fixed}}]$, where behavioral differences are conveyed for a suitable abstraction α_{bug} suppressing buggy traces or relating them to corrected ones. We can go a step further and not just exclude buggy traces, but encode the *fix* by an abstraction

α_{patch}. This is likely to produce a more reliable result asserting that apart from the fix, the programs behave equivalently, even for the formerly buggy paths. In the appendix, we discuss two alternative formalizations with an example.

4 Reasoning about the Trace Modality

We propose a reasoning algorithm based on *symbolic traces* for the trace modality. The idea is to lift all verification domains to a common language over symbolic traces that over-approximates the set of concrete traces produced by each lifting function. Abstractions are generalized to symbolic traces. Validity of the trace modality can then be established by *symbolic trace subsumption*. The symbolic traces we propose are a *regular* language. Hence, programs generally have to be over-approximated, for example, by loop invariant reasoning or bounded loop unwinding. Not all properties can be encoded in a regular symbolic trace language, such as complex Computation Tree Logic properties. Even so, the language is expressive enough to represent the problems formalized in Sect. 3, and problems encoded in it can be solved effectively (if the underlying first-order problems can be solved). We define symbolic stores, states and traces as follows:

Definition 2. *Let* $\mathtt{x} \in \mathrm{PVar}$, $t \in \mathrm{Trm}$, $\varphi \in \mathrm{Fml}$, *and* $\mathtt{P}$ *a schematic statement. The sets* SymSto *of* Symbolic Stores, SymState *of* Symbolic States, *and* SymTr *of* Symbolic Traces *are defined as follows in extended BNF:*

$$\begin{aligned} \mathrm{SymSto} &::= \mathtt{x}\ \text{“:=”}\ t \mid \mathit{sto}_{\mathtt{P}} \mid \mathrm{SymSto}\ \text{“||”}\ \mathrm{SymSto} \mid \text{“\{”}\ \mathrm{SymSto}\ \text{“\}”}\ \mathrm{SymSto} \\ \mathrm{SymState} &::= \varphi \mid \text{“(”}\ \mathrm{SymSto}\ \text{“,”}\ \varphi\ \text{“)”} \\ \mathrm{SymTr} &::= \mathrm{SymState} \mid \mathrm{SymTr}\,(\text{“;”} \mid \text{“+”})\,\mathrm{SymTr} \mid \varphi\ \text{“}^{!}\text{”} \mid \mathrm{SymTr}\ \text{“}^{*}\text{”} \end{aligned}$$

Here an abstract *store* $\mathit{sto}_{\mathtt{P}}$ *represents an unknown state transition induced by a schematic* $\mathcal{L}$*-statement* $\mathtt{P}$. *The sets* SymSto_0, $\mathrm{SymState}_0$ *and* SymTr_0 *are defined as above, but do not contain abstract stores.*

Symbolic stores record changes to program variables. Elementary stores $\mathtt{x} := t$ represent states where the variable $\mathtt{x}$ attains the valuation of the (symbolic) term t. Symbolic stores sto_1, sto_2 are combined to a parallel store $\mathit{sto}_1 \,||\, \mathit{sto}_2$. If both assign a value to the same variable, the later assignment (in sto_2) “wins”. A symbolic store sto_1 can be *applied* to a symbolic store sto_2, written $\{\mathit{sto}_1\}\mathit{sto}_2$. Left-hand sides in sto_2 are then evaluated in the states represented by sto_1. Combining two stores into one works by the *store concatenation operator* “$\circ$”, defined as $\mathit{sto}_1 \circ \mathit{sto}_2 = \mathit{sto}_1 \,||\, \{\mathit{sto}_1\}\mathit{sto}_2$. We permit the application of symbolic stores to terms and formulas, with similar semantics. We write $\overrightarrow{\mathtt{x}} := \overrightarrow{t}$ for the store $\mathtt{x}_1 := t_1 \,||\, \ldots \,||\, \mathtt{x}_n := t_n$, where $\mathtt{x}_i$, t_i are the i-th components of $\overrightarrow{\mathtt{x}}$, $\overrightarrow{t}$.

Symbolic states consist of an (optional) symbolic store *sto* and path condition φ representing concrete states satisfying both φ and, if present, the assignments in *sto*. A symbolic trace is in the simplest case a sequence of symbolic states. The choice operator + models nondeterministic choice as well as case distinctions for deterministic programs, depending on the path conditions of the argument traces. The trace $\varphi^{!}$, primarily used to model assertions, represents the empty

$$
\begin{aligned}
\mathsf{tval}_K :&\quad \mathrm{SymSto}_0 \to (\mathcal{S} \to \mathcal{S}) \\
\mathsf{tval}_K(\mathrm{x} := t)(s)(\mathrm{y}) &\coloneqq \mathsf{val}(K, s; t) \text{ if } \mathrm{y} = \mathrm{x},\ s(\mathrm{y}) \text{ otherwise} \\
\mathsf{tval}_K(sto_1 \,||\, sto_2)(s) &\coloneqq \mathsf{tval}_K(sto_2)(\mathsf{tval}_K(sto_1)(s)) \\
\\
\mathsf{tval}_K :&\quad \mathrm{SymState}_0 \to 2^{\mathcal{S}} \\
\mathsf{tval}_K(\varphi) &\coloneqq \{s \in \mathcal{S} \mid K, s \models \varphi\} \\
\mathsf{tval}_K(sto, \varphi) &\coloneqq \{\mathsf{tval}_K(sto)(s) \mid K, s \models \varphi,\ s \in \mathcal{S}\} \\
\\
\mathsf{tval}_K :&\quad \mathrm{SymTr}_0 \to 2^{Traces} \\
\mathsf{tval}_K(\tau_1; \tau_2) &\coloneqq \{\tau_1^0 \tau_2^0 \mid \tau_1^0 \in \mathsf{tval}_K(\tau_1) \setminus \{\cdots \bot\},\ \tau_2^0 \in \mathsf{tval}_K(\tau_2)\} \\
&\quad \cup \{\bot \mid \tau_1^0 \bot \in \mathsf{tval}_K(\tau_1)\} \\
\mathsf{tval}_K(\tau_1 + \tau_2) &\coloneqq \mathsf{tval}_K(\tau_1) \cup \mathsf{tval}_K(\tau_2) \\
\mathsf{tval}_K(\varphi^!) &\coloneqq \begin{cases} \varepsilon & \text{if } \forall s \in \mathcal{S} :\ K, s \models \varphi \\ \bot & \text{otherwise} \end{cases} \\
\mathsf{tval}_K(\tau^*) &\coloneqq \{\tau_1^0 \tau_2^0 \cdots \tau_n^0 : \tau_i \in \mathsf{tval}_K(\tau),\ n \in \mathbb{N}\} \cup \{\varepsilon\}
\end{aligned}
$$

Fig. 1: The Valuation Function tval_K

trace if φ holds in the current state and the failure state otherwise. The traces τ^* represent all finite concrete traces in which all states satisfy τ. For instance, true^* represents the set of all finite concrete traces. We do not include an operator τ^ω for infinite traces which would significantly complicate validity checking: One would have to separate terminating from non-terminating traces—which is undecidable—or consider only non-terminating runs.

A formal semantics for symbolic traces is based on a first-order structure K with domain D and interpretation I, as well as $s \in \mathcal{S}$. The *valuation function* $\mathsf{val}(K, s; \cdot)$ assigns to terms a value in D, to formulas true or false. We write equivalently $K, s \models \varphi$ or $\mathsf{val}(K, s; \varphi) = \mathrm{true}$, as well as $K, s \not\models \varphi$ or $\mathsf{val}(K, s; \varphi) = \mathrm{false}$. The function $\mathsf{val}(K, s; \cdot)$ is defined as usual, except for the application of symbolic stores and the valuation of program variables. For $\mathrm{x} \in \mathrm{PVar}$, we define $\mathsf{val}(K, s; \mathrm{x}) = s(\mathrm{x})$. If $t \in \mathrm{Trm}$ and $sto \in \mathrm{SymSto}$, we define $\mathsf{val}(K, s; \{sto\}t) \coloneqq \mathsf{val}(K, s'; t)$, where $s' = \mathsf{val}(K, s; sto)(s)$ (similarly for formulas). We define the trace valuation function tval_K first on *concrete* symbolic traces SymTr_0. It is parametric in a structure K that fixes the values of uninterpreted constant, function, and predicate symbols. The cumulative valuation function tval is canonically defined as $\mathsf{tval}(\tau) \coloneqq \bigcup_K \mathsf{tval}_K(\tau)$.

Definition 3. *We inductively define the valuation function* tval_K*, overloaded for symbolic stores, states and traces, as in Fig. 1.*

Symbolic traces SymTr_0 are created for concrete programs $\mathcal{L}_0$. The symbolic evaluation of schematic programs in $\mathcal{L}$ creates abstract stores sto_{P} and path conditions C_{P} (details below). Intuitively, they represent *all possible symbolic*

stores and path conditions that may arise from concrete program execution. We define their semantics by the union of the semantics of possible instantiations.

Definition 4. *Let* $\tau \in \mathrm{SymTr}$ *be a symbolic trace with occurrences of abstract stores* $sto_{\mathtt{P}_1}, \ldots, sto_{\mathtt{P}_n}$ *and path conditions* $C_{\mathtt{P}_1}, \ldots, C_{\mathtt{P}_n}$ *(with possibly multiple occurrences of each* $sto_{\mathtt{P}_i}$*,* $C_{\mathtt{P}_i}$*). We define* $\mathsf{tval}(\tau)$ *as the union* $\bigcup \mathsf{tval}(\tau_0)$ *of all* $\tau_0 \in \mathrm{SymTr}_0$ *that are obtained by instantiating all occurrences of* $sto_{\mathtt{P}_i}$*,* $C_{\mathtt{P}_i}$ *with concrete stores* $sto^0_{\mathtt{P}_i} \in \mathrm{SymSto}_0$ *and path conditions* $C^0_{\mathtt{P}_i} \in \mathrm{Fml}$*.*

Abstractions α are generalized to symbolic traces in the obvious manner, e.g., the big-step abstraction α_{big} takes the first and all final states of a trace. Symbolic representations of the lifting functions require more work. For a lifting function *lift*, we denote by *slift* its symbolic version. Like *lift*, *slift* takes a *symbolic* state and a verification domain construct and produces a *symbolic* trace.

Definition 5. *A symbolic lifting function slift is* correct *relative to lift if, for all* $s \in \mathrm{SymState}$ *and* $\sigma \in \mathsf{tval}(s)$*,* $lift(\sigma)(C) \subseteq \mathsf{tval}(slift(s)(C))$*.*

Symbolic lifting functions for first-order formulas are straightforward to define: $slift_{\mathrm{Fml}}(s)(\varphi) \coloneqq \mathrm{true}^*; \varphi$. For LTL formulas, $slift_{TL}(s)$ maps (1) φ to "φ", (2) $\Box\varphi$ to "φ^*", (3) $\Diamond\varphi$ to "$\mathrm{true}^*; \varphi; \mathrm{true}^*$", and (4) $\varphi \,\mathcal{U}\, \psi$ to "$\varphi^*; \psi; \mathrm{true}^*$".

Defining symbolic lifting for programs means encoding symbolic execution. E.g., one can extract symbolic traces from a symbolic execution tree. Symbolic *traces* are more flexible, though, since they can encode non tree-like structures. The lifting function $slift_{\mathcal{L}_0}$ is defined as follows for assignments, if-else, assume and assert, and sequential composition (for those, it coincides with $slift^k_{\mathcal{L}_0}$ for BMC). W.l.o.g., we assume symbolic states to be of the form (sto, φ).

$$
\begin{aligned}
slift_{\mathcal{L}_0}(sto, \varphi)(\mathtt{x=e}) &\coloneqq (sto \circ (\mathtt{x} := e), \varphi) \\
slift_{\mathcal{L}_0}(sto, \varphi)(\mathtt{if(g)}\ p_1\ \mathtt{else}\ p_2) &\coloneqq (slift_{\mathcal{L}_0}(sto, \varphi \wedge \{sto\}g))(p_1) \,+ \\
&\qquad\quad (slift_{\mathcal{L}_0}(sto, \varphi \wedge \neg\{sto\}g))(p_2) \\
slift_{\mathcal{L}_0}(sto, \varphi)(\mathtt{assume}(\psi)) &\coloneqq (sto, \varphi \wedge \psi) \\
slift_{\mathcal{L}_0}(sto, \varphi)(\mathtt{assert}(\psi)) &\coloneqq (\varphi \rightarrow \{sto\}\psi)^! \\
slift_{\mathcal{L}_0}(sto, \varphi)(p_1\,;\, p_2) &\coloneqq \{\tau_1; \tau_2 : \tau_1 \in slift_{\mathcal{L}_0}(sto, \varphi)(p_1), \\
&\qquad\quad \tau_2 \in slift_{\mathcal{L}_0}(last(\tau_1))(p_2)\}
\end{aligned}
$$

Symbolic lifting is more complex for loops, as usual in symbolic execution. Possible approaches are *loop unwinding* which generally does not terminate for loops with symbolic guards, *bounded unwinding* with a fixed upper bound on the number of unwinding steps, and *loop invariants*. In the appendix, we provide a more detailed discussion and define symbolic lifting for those cases.

To define $slift_{\mathcal{L}}$, we have to encode schematic statements $\mathtt{P}$. We choose to do this with *abstract stores* $sto_{\mathtt{P}}$ that model state changes caused by schematic statements. We also admit *abstract formulas* $C_{\mathtt{P}}$ to model (unknown) path condition constraints arising from an abstract program $\mathtt{P}$. We define:

$$slift_{\mathcal{L}}(sto, \varphi)(\mathtt{P}) \coloneqq \mathrm{true}^*; (sto \circ sto_{\mathtt{P}}, \varphi \wedge \{sto \circ sto_{\mathtt{P}}\}C_{\mathtt{P}}) \quad \text{for all } \mathtt{P}.$$

The Algorithm for Checking Validity of Trace Modalities. When presented with a problem $[\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$ and a symbolic state s_0, the algorithm evaluates in three phases whether $\alpha(\mathit{lift}_l(\sigma_0)(\mathcal{C}_l)) \subseteq \alpha(\mathit{lift}_r(\sigma_0)(\mathcal{C}_r))$ holds for all $\sigma_0 \in \mathsf{tval}(s_0)$:

(1) Convert $\mathcal{C}_{l/r}$ to symbolic traces $\tau^s_{l/r}$ using symbolic lifting functions $\mathit{slift}_{l/r}$ (as described above for first-order and LTL formulas, as well as $\mathcal{L}$-programs).
(2) Construct Symbolic Finite Automata (SFAs) $\mathsf{SFA}_{l/r}$ accepting the languages $\mathsf{tval}(\tau^s_{l/r})$, i.e. concrete traces represented by symbolic ones.
(3) Check whether the language accepted by SFA_l is *included* in the language accepted by SFA_r through construction of a *simulation relation.*

Transitions in an SFA are labeled with symbolic states that may represent infinitely many concrete states. For example, a transition labeled with "true" models a transition for *any* concrete state. Formally, we define SFA as:

Definition 6. *A* Symbolic Finite Automaton *is a tuple $A = (Q, \Sigma, \delta, q_0, F)$ of a finite set of states Q, an alphabet $\Sigma \subseteq \mathcal{S}$, a finite transition relation $\delta \subseteq Q \times \text{SymState} \times Q$, an initial state $q_0 \in Q$ and a set of accepting states $F \subseteq Q$. Automaton A* accepts *a concrete trace $\sigma_1\sigma_2\cdots\sigma_n$ if there is a path $q_1 \xrightarrow{s_1} q_2 \xrightarrow{s_2} \cdots q_n \xrightarrow{s_n} q_{n+1}$ in A such that $q_{n+1} \in F$ and for each $i = 1, \ldots, n$ it holds that $\sigma_i \in \mathsf{tval}(s_i)$. The* language $L(A)$ *of an SFA A is the set of all accepted traces.*

The construction of an SFA from symbolic traces (step (2)) is shown in Algo. 2 in the appendix. Lem. 1 states the soundness of the algorithm.

Lemma 1. *Function* createSFA *in Algo. 2 is correct: $L(\text{createSFA}(\tau)) = \mathsf{tval}(\tau)$ holds for all $\tau \in \text{SymTr}$.*

Simulation relations on automata for checking language inclusion [26] and the complexity of crating them [12] have been studied before. Our notion is non-standard, though, since we use symbolic automata with *first-order* transitions. It is not sufficient to relate edges with identical labels or to use existing *propositional* symbolic approaches. Instead, we try to *prove* that an edge in the specification automaton *subsumes* an edge in the implementation automaton. We define symbolic state subsumption as follows.

Definition 7. *Let $s_i = (\mathit{sto}_i, \varphi_i)$, $i = 1, 2$ be symbolic states. Let $\overrightarrow{\mathrm{x}_i}$ be the left-hand sides of sto_i, subst be a substitution of abstract symbols in s_2 not occurring in s_1 with concrete symbols; i.e. uninterpreted constants, function symbols, abstract stores, abstract path conditions are replaced with terms, stores, and formulas. Let P be a fresh predicate with arity $|\overrightarrow{\mathrm{x}_2}|$. Then s_2* subsumes *s_1 iff*

(SUB1) *all variables in $\overrightarrow{\mathrm{x}_2}$ are also contained in $\overrightarrow{\mathrm{x}_1}$ and*
(SUB2) *there is a substitution subst such that:*
$$\models \varphi_1 \wedge \{\mathit{sto}_1\}P(\overrightarrow{\mathrm{x}_2}) \rightarrow \mathit{subst}\big(\big(\{\mathit{sto}_1\}\varphi_2 \wedge \{\mathit{sto}_2\}P(\overrightarrow{\mathrm{x}_2})\big)\big) \ .$$

For states without stores omit the $\{\mathit{sto}_i\}$. In the following, we write $s_1 \sqsubseteq s_2$ if s_2 subsumes s_1, and $s_1 \sqsubseteq_{\mathit{subst}} s_2$ to make the substitution subst for (SUB2) explicit.

Example 3. Let $s_1 = (\mathrm{x} := 17 \,||\, \mathrm{y} := 42 \,||\, \mathrm{z} := 2, \mathrm{true})$. It is subsumed by $s_2 = (\mathrm{x} := c, c \geq 0)$, since (SUB2) holds for $subst := (c \mapsto 17)$:

$$\begin{aligned} & \models \{\mathrm{x} := 17\}P(\mathrm{x}) \rightarrow (c \mapsto 17)(\{\mathrm{x} := 17\}c \geq 0 \wedge \{\mathrm{x} := c\}P(\mathrm{x})) \\ \text{follows from } & \models \{\mathrm{x} := 17\}P(\mathrm{x}) \rightarrow (\{\mathrm{x} := 17\}17 \geq 0 \wedge \{\mathrm{x} := 17\}P(\mathrm{x})) \\ \text{follows from } & \models P(17) \rightarrow (17 \geq 0 \wedge P(17)) \end{aligned}$$

which is true (w.l.o.g. we omit parts of the store of s_1 that do not occur in the target formula). Two more small examples are in Example 8 (appendix). ◊

Lemma 2. *For s_1, $s_2 \in$ SymState, $s_1 \sqsubseteq s_2$ implies $\mathsf{tval}(s_1) \subseteq \mathsf{tval}(s_2)$.*

Subsumption can also be used to establish whether, for a concrete state σ and symbolic state s, it holds that $\sigma \in \mathsf{tval}(s)$ which is needed for the acceptance criterion of SFAs (Def. 6): for the symbolic state $s' = (\overrightarrow{\mathrm{x_s}} := \sigma(\overrightarrow{\mathrm{x_s}}), \mathrm{true})$, where $\overrightarrow{\mathrm{x_s}}$ are the left-hand sides of the store of s, it is sufficient to prove $s' \sqsubseteq s$.

Now we can define the notion of *a Subsumption Simulation Relation (SSR)*, a simulation relation on SFAs based on subsumption.

Definition 8. *A* Subsumption Simulation Relation *between SFAs $A_i = (Q_i, \Sigma, \delta_i, q_0^i, F_i)$, $i = 1, 2$, is any relation $R \subseteq Q_1 \times Q_2$ satisfying*

$$\begin{aligned} &\text{(SR1)} \quad \forall q_1 \in Q_1, q_2 \in Q_2, s, q_1' \in Q_1, \\ &\qquad \big((R(q_1, q_2) \wedge (q_1, s, q_1') \in \delta_1) \implies \\ &\qquad\quad \exists q_2' \in Q_2, s', (R(q_1', q_2') \wedge (q_2, s', q_2') \in \delta_2 \wedge \boxed{s \sqsubseteq s'}\big) \\ &\text{(SR2)} \quad (q_0^1, q_0^2) \in R \end{aligned}$$

Def. 8 equals the "safety simulation relation" of [12], except for the highlighted conjunct $s \sqsubseteq s'$ in (SR1). Constructing an SSR additionally requires to find a suitable substitution and to call a prover showing subsumption. Since SSRs are closed under union and (SR2) is monotone, one can compute R by repeatedly deleting pairs from $Q_1 \times Q_2$ that locally do not satisfy (SR1), and then check whether the result satisfies (SR2) [12]). For each local check, we might have to substitute abstract symbols in the specification automaton. The subsequent lemma, also stated in [12] for their very similar notion, establishes a sufficient condition between simulation relations and language inclusion.

Lemma 3. *If there is an SSR between SFAs A_1 and A_2, then $L(A_1) \subseteq L(A_2)$.*

Our top-level algorithm EVALUATE is shown in Algo. 1. In the final step it tries to find an SSR. Only if this was successful, it returns YES. Function FIND-SSR (Algo. 1) starts with an "initial simulation" produced by function INITSIM (Algo. 3, appendix) instead of the cross product to save expensive subsumption checks. During the filtering to derive an SSR, it maintains a *set* of substitutions *substs*, since there might be multiple options. Function SUBSUMPTION$(s, s', substs)$ (Algo. 4, appendix) tries to find compatible *extensions subst'* of the substitutions

Algorithm 1 Evaluation of a Trace Modality Formula using SSRs

function EVALUATE(s_0, $[\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]$)
  $\tau_l \leftarrow \alpha(\mathit{slift}_l(s_0)(\mathcal{C}_l))$, $\tau_r \leftarrow \alpha(\mathit{slift}_r(s_0)(\mathcal{C}_r))$ ▷ Step (1)
  $A_l \leftarrow$ CREATESFA(τ_l), $A_r \leftarrow$ CREATESFA(τ_r) ▷ Step (2)
  if $(q_0^l, q_0^r) \in$ FINDSSR(A_l, A_r) **then return** YES ▷ Step (3)
  else return `UNKNOWN` **end if**
end function

function FINDSSR($(Q_l, \Sigma, \delta_l, q_0^l, F_l)$, $(Q_r, \Sigma, \delta_r, q_0^r, F_r)$)
  $R \leftarrow$INITSIM(Q_l, Q_r, δ_l, δ_r), $\mathit{substs} \leftarrow \{\lambda x.x\}$, $\mathit{changed} \leftarrow \mathit{true}$
  while $\mathit{changed} = \mathit{true}$ **do**
    $\mathit{changed} \leftarrow \mathit{false}$
    for all $(q_l, q_r) \in R$, $(q_l, s, q_l') \in \delta_l$ **do**
      if $\exists (q_r, s', q_r') \in \delta_r$ s.t. SUBSUMPTION(s, s', substs) $\neq \emptyset$) **then**
      $\mathit{substs} \leftarrow$ SUBSUMPTION(s, s', substs) ▷ (for all such s')
      else $R \leftarrow R \setminus (q_l, q_r)$, $\mathit{changed} \leftarrow \mathit{true}$ **end if**
    end for
  end while
  return R
end function

in *substs* by first applying an existing substitution and then finding another one for yet uninstantiated abstract symbols. If there is no such substitution, e.g., since one would have to instantiate the same abstract symbol with different values, the original substitution is dropped. We do not further specify the process of finding substitutions; a naive approach could try to instantiate abstract symbols with all combinations of terms occurring as right-hand sides in the store of s. An example application of Algo. 1 is shown in the appendix. Lem. 4 below states correctness of the FINDSSR. The subsequent main theorem follows from Lems. 1 to 4 and the usage of correct symbolic lifting functions (Def. 5).

Lemma 4. *Function* FINDSSR *(Algo. 1) is correct: For SFAs A_1, A_2, it holds that any SSR R found by* FINDSSR(A_1, A_2) *satisfies (SR1).*

Theorem 1. *Function* EVALUATE *(Algo. 1) is correct: For all $s_0 \in$* SymState*,* EVALUATE$(s_0, [\mathcal{C}_1 \Vdash_\alpha \mathcal{C}_2]) =$ YES *only if, for all $\sigma \in \mathsf{tval}(s_0)$, $\sigma \models [\mathcal{C}_1 \Vdash_\alpha \mathcal{C}_2]$.*

5 Related Work

We compare our work to (1) logics based on traces and (2) approaches unifying program verification techniques. De Giacomo & Vardi [11] propose a Regular Temporal Specification language RE_f that is syntactically similar to our symbolic traces, but ranges over *propositional* formulas while our atoms are first-order symbolic states. They show that RE_f has the same expressiveness as Monadic Second-order Logic (MSO) and is strictly more expressive than LTL on finite traces. They define Linear-time Dynamic Logic LDL_f, having the same

expressivity as RE_f, but allowing logical connectives like negation. Reasoning in LDL_f is also translated to automata. They mention, but do not detail, the possibility to "capture finite executions of programs [...] (in a propositional variant [...])", which is exactly what we do—but not restricted to a propositional variant. In addition, we incorporate *abstract programs* to reason about *classes* of programs. It would be interesting to investigate whether we could use a variant of LDL_f to embed symbolic traces conveniently into logic formulas.

Beckert & Bruns [5] combine dynamic logic and first-order temporal logic to a *Dynamic Trace Logic*. They have a trace-based semantics for a while language and provide a sequent calculus to reason about temporal properties (not preceded by symbolic lifting). The calculus rules depend on the top-level operator of the first-order LTL post condition. This leads quite complex loop invariant rules. Also, the approach is not directly applicable to other verification domains, e.g., relational verification. Our approach is more flexible, because there is no syntactic constraint between the left and right-hand side of the trace modality.

Din et al. [13] propose a trace semantics for the actor-based concurrent language ABS. Traces are "locally abstract, globally concrete": at the local (e.g., method) level, symbolic traces are used. These are primarily a *semantic* notion, facilitating a modular semantics for a concurrent language, while our symbolic traces are *syntactic* entities. The authors briefly sketch a program logic with trace formulas, but leave the notion of trace formulas abstract.

Regarding area (2), Kamburjan [22] proposes the *behavioral modality* aiming to integrate existing analyses and sharing some aspects with the trace modality. It asserts that a statement in a concurrent language meets a behavioral specification consisting of a *type* and a *translation* of the type into an MSO formula. This is the case if that formula holds for all traces generated by the statement. Important differences to our approach include: (a) The behavioral modality *syntactically* integrates analyses on the *same program class*, while the trace modality is mainly a *general semantic framework*, (b) the "translation" of [22] projects to MSO and is thus less expressive than lifting to arbitrary trace sets. The trace modality can also be used to combine verification techniques. Two specifications can semantically be combined by forming the intersection of the trace sets. For reasoning about combinations, we could use product constructions on SFAs.

Some systems do not aim to provide a common semantics for verification domains, but provide a framework to *implement* different analyses. They usually represent verification problems in an Intermediate Language (IL) and interface to different provers. Boogie [2] and Why3 [7] both are an IL and tool for deductive program verification. They are used as backends by verifiers for languages like C and Java. Our "IL" are the regular symbolic trace language, which, compared to Boogie and WhyML, is less usable for direct programming, more abstract and less expressive (e.g., we cannot directly write loops, but have to use invariants). Yet, the syntactic notion of symbolic traces is closely related to the semantic notion of the trace modality, allowing to *formalize* and *prove* a problem in a closely related framework. Moreover, the trace modality can easily express other prob-

lems than "standard" post condition verification. Our algorithm also interfaces to different provers: Which one to use in the subsumption step is left open.

6 Conclusion and Outlook

We presented the trace modality, a novel formalism for expressing many practical problems of sequential program verification. It relates two elements of the same or different domains, e.g., programs, first-order assertions, or temporal logic formulas. Programs can be abstract and represent classes of concrete programs. We demonstrate the usefulness of the trace modality by providing formalizations of various verification problems: Functional Verification, Information Flow Analysis, Model Checking, Program Synthesis, Compilation, and Program Evolution. Our uniform reasoning system translates programs and formulas to regular symbolic traces and then reduces the problem to the construction of simulation relations between finite automata with symbolic transitions. Similar to the semantics of the trace modality, this approach is parametric in the translation to symbolic traces and the abstraction operator. Although regular symbolic traces have already been proposed before as both a specification mechanism and semantic representation, our work is the first we know of connecting both aspects. This facilitates flexible reasoning about programs and specifications in different combinations: A program can even serve as the specification of a formula.

We hope that our uniform formalization helps to uncover synergy potential between so far separate areas in the field of program verification. Moreover, the practical potential of a system based on symbolic traces supporting different verification techniques, for example, program synthesis and deductive verification, is huge. For instance, after a failed proof attempt of a postcondition, one could try synthesis techniques for stepwise refinement of the postcondition to an abstract program. MC and deductive verification techniques could work hand in hand to treat loops, by unwinding, k-induction, abstract interpretation-based techniques, etc. Finally, the idea of "patch abstraction" for program evolution could help in proof reuse, by applying the patches also to existing proofs.

Apart from investigating these ideas, we plan to implement our reasoning algorithm for symbolic traces and to examine different existing trace languages, like linear-time dynamic logic, which might lead to more intuitive or more expressive representations. Also, we project to extend our framework to non-deterministic, in particular, to concurrent programming languages.

References

1. Ahrendt, W., Beckert, B., et al. (eds.): Deductive Software Verification – The KeY Book, LNCS, vol. 10001. Springer (2016)
2. Barnett, M., Chang, B.Y.E., et al.: Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In: Intern. Symp. on FMCO. pp. 364–387. Springer (2005)
3. Barthe, G., Crespo, J.M., et al.: Relational Verification Using Product Programs. In: Butler, M.J., Schulte, W. (eds.) Proc. 17th FM. pp. 200–214. Springer (2011)

4. Barthe, G., D'Argenio, P.R., et al.: Secure Information Flow by Self-Composition. In: Proc. CSFW-17. pp. 100–114. IEEE Computer Society (2004)
5. Beckert, B., Bruns, D.: Dynamic Logic with Trace Semantics. In: CADE-24 (2013)
6. Biere, A., Cimatti, A., et al.: Bounded Model Checking. Advances in Computers **58**, 117–148 (2003)
7. Bobot, F., Filliâtre, J.C., et al.: Why3: Shepherd Your Herd of Provers. In: Boogie 2011: First International Workshop on IVL. pp. 53–64 (2011)
8. Clarke, E.M., Kroening, D., et al.: A Tool for Checking ANSI-C Programs. In: Proc. TACAS 2004. pp. 168–176 (2004)
9. Cousot, P., Cousot, R.: Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In: 4th Symp. of POPL. pp. 238–252. ACM Press (Jan 1977)
10. Darvas, Á., Hähnle, R., et al.: A Theorem Proving Approach to Analysis of Secure Information Flow. In: Proc. 2nd Intern. Conf. on SPC. pp. 193–209 (2005)
11. De Giacomo, G., Vardi, M.Y.: Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In: Proc. 23rd IJCAI. pp. 854–860 (2013)
12. Dill, D.L., Hu, A.J., et al.: Checking for Language Inclusion Using Simulation Preorders. In: CAV '91. pp. 255–265. LNCS, Springer (1991)
13. Din, C.C., Hähnle, R., et al.: Locally Abstract, Globally Concrete Semantics of Concurrent Programming Languages. In: TABLEAUX 2017. pp. 22–43 (2017)
14. Garrido, A., Meseguer, J.: Formal Specification and Verification of Java Refactorings. In: Proc. 6th SCAM. pp. 165–174. IEEE Computer Society (2006)
15. Godlin, B., Strichman, O.: Regression Verification: Proving the Equivalence of Similar Programs. Softw. Test., Verif. Reliab. **23**(3), 241–258 (2013)
16. Hähnle, R., Heisel, M., et al.: An Interactive Verification System based on Dynamic Logic. In: Siekmann, J.H. (ed.) 8th CADE, pp. 306–315. Springer (1986)
17. Harel, D., Tiuryn, J., et al.: Dynamic Logic. MIT Press (2000)
18. Heisel, M.: Formalizing and Implementing Gries' Program Development Method in Dynamic Logic. Sci. Comput. Program. **18**(1), 107–137 (1992)
19. Hoare, C.A.R.: An Axiomatic Basis for Computer Programming. Communications of the ACM **12**(10), 576–580 (1969)
20. Holzmann, G.J.: The Model Checker SPIN. IEEE Trans. SE **23**(5) (1997)
21. Jhala, R., Majumdar, R.: Software Model Checking. ACM Comput. Surv. **41**(4), 21:1–21:54 (2009)
22. Kamburjan, E.: Behavioral Program Logic. In: Proc. 28th TABLEAUX (2019)
23. Leroy, X.: Formal Verification of a Realistic Compiler. Comm. ACM **52**(7) (2009)
24. Monperrus, M.: Automatic Software Repair: A Bibliography. ACM Comput. Surv. **51**(1), 17:1–17:24 (2018)
25. Pasareanu, C.S., Visser, W.: Verification of Java Programs Using Symbolic Execution and Invariant Generation. In: Proc. 11th Intern. SPIN Workshop on Model Checking Software. pp. 164–181 (2004)
26. Rauch Henzinger, M., Henzinger, T.A., et al.: Computing Simulations on Finite and Infinite Graphs. In: Proc. 36th Symp. on FoCS. pp. 453–462. IEEE (1995)
27. Reps, T.W., Horwitz, S., et al.: Precise Interprocedural Dataflow Analysis via Graph Reachability. In: Proc. 22nd POPL. pp. 49–61 (1995)
28. Sabelfeld, A., Myers, A.C.: A Model for Delimited Information Release. In: Proc. 2nd Intern. Symp. on Software Security - Theories and Systems. pp. 174–191 (2003)
29. Shankar, N.: Combining Model Checking and Deduction. In: Handbook of Model Checking., pp. 651–684. Springer (2018)
30. Srivastava, S., Gulwani, S., et al.: From Program Verification to Program Synthesis. In: Proc. 37th POPL. pp. 313–326 (2010)

31. Steinhöfel, D., Hähnle, R.: Modular, Correct Compilation with Automatic Soundness Proofs. In: Margaria, T., Steffen, B. (eds.) Proc. 8th ISOLA. LNCS (2018)
32. Visser, W., Havelund, K., et al.: Model Checking Programs. Autom. Softw. Eng. **10**(2), 203–232 (2003)
33. Yang, H.: Relational Separation Logic. Theoretical CS **375**(1-3), 308–334 (2007)

Appendix

We provide additional information and explanations that had to be left out of the paper for space reasons. Numbered (sub-)sections relate to those of the paper.

1 Introduction

Remark. A striking example for the difficulty of knowledge transfer in formal methods is the sparse interaction between the deductive verification and the abstract interpretation community, given the significant methodological overlap.

Remark. One can view *approximation* as a special case of *abstraction*. Since approximation can be expressed by a subset relation alone, it is unnatural to conflate them, however.

2 Programs, Logic, Traces and Abstractions

For completeness, we subsequently define a trace semantics for a simple deterministic while language. In Sect. 2, only the semantics for the more exotic cases of the **assert** and **assume** statements was defined. In the following definition, x represents program variables and e expressions. We write $[\![e]\!]_s$ for the semantics of expression e in a state s. For appending a trace τ_2 to a (finite) trace τ_1, we simply write $\tau_1\tau_2$. The empty trace ε is the neutral element of this operation. The predicate $\mathit{failed}(\tau)$ holds for a trace τ if it contains the failure state $\bot$.

$$
\begin{aligned}
\mathrm{Tr}_s(\texttt{x=e}) &\coloneqq \{s[\texttt{x} \mapsto [\![e]\!]_s]\} \\
\mathrm{Tr}_s(\texttt{if}(e)\ p_1\ \texttt{else}\ p_2) &\coloneqq \begin{cases} \mathrm{Tr}_s(p_1) & \text{if } [\![e]\!]_s = \mathit{true} \\ \mathrm{Tr}_s(p_2) & \text{otherwise} \end{cases} \\
\mathrm{Tr}_s(p_1\ ;\ p_2) &\coloneqq \{\tau \in \mathrm{Tr}_s(p_1) : \neg \mathit{finite}(\tau) \wedge \neg \mathit{failed}(\tau)\} \cup \\
&\quad \{\tau_1\tau_2 : \tau_1 \in \mathrm{Tr}_s(p_1) \wedge \mathit{finite}(\tau_1) \wedge \neg \mathit{failed}(\tau_1), \\
&\qquad \tau_2 \in \mathrm{Tr}_{\mathit{last}(\tau_1)}(p_2)\} \cup \\
&\quad \{\bot : \tau_1 \in \mathrm{Tr}_s(p_1) \wedge \mathit{failed}(\tau_1)\} \\
\mathrm{Tr}_s(\texttt{while}(e)\ p) &\coloneqq \begin{cases} \mathrm{Tr}_s(p;\ \texttt{while}(e)\ p) & \text{if } [\![e]\!]_s = \mathit{true} \\ \{\varepsilon\} & \text{otherwise} \end{cases} \\
\mathrm{Tr}_s(\texttt{assert}(\varphi)) &\coloneqq \begin{cases} \{s\} & \text{if } s \models \varphi \\ \{\bot\} & \text{otherwise} \end{cases} \\
\mathrm{Tr}_s(\texttt{assume}(\varphi)) &\coloneqq \begin{cases} \{s\} & \text{if } s \models \varphi \\ \emptyset & \text{otherwise} \end{cases} \\
\mathrm{Tr}_s(\texttt{havoc}) &\coloneqq \mathcal{S}
\end{aligned}
$$

Example 4 (Trace Semantics). The $\mathcal{L}$-Program $p = \texttt{x=-1;}\ \texttt{assume}(\texttt{x} \geq 0)$ evaluates to the empty trace set ($\mathrm{Tr}(p) = \emptyset$) because of the definition of the semantics for the sequencing operator. There is no $\tau_2 \in \mathrm{Tr}_s(\texttt{assume}(\texttt{x} \geq 0))$ since $\mathrm{Tr}_{s[\texttt{x} \mapsto -1]}(\texttt{assume}(\texttt{x} \geq 0))$ returns the empty set. Conversely, $p' = \mathrm{Tr}(\texttt{x=-1;}$

assume(x $\leq$ 0)) evaluates to the set $\{s[\mathrm{x} \mapsto -1] : s \in \mathcal{S}\}$, since now, the empty trace is produced for the **assume** statement, which is the neural element of trace concatenation. The trace modality formula $[p \Vdash \mathit{spec}]$ thus holds for any specification *spec*, since the empty set is a subset of any set; this is not the case for p', where satisfying the specification would be more meaningful. Consider now the following program: $p'' =$ x=-1; **assume**(y $\geq$ 0). There, we make an assumption about a variable y that is not mentioned before. It evaluates to

$$\mathrm{Tr}(p'') = \{s' : s \in \mathcal{S} \wedge s' = s[\mathrm{x} \mapsto -1] \wedge [\![\mathrm{y}]\!]_{s'} \geq 0\},$$

i.e. the set of all states that after setting x to -1 satisfy the property of y being positive. This is the actual use case of assumptions: To assume facts that *cannot* be otherwise established locally in the current program context.

The property of assumptions to trivially satisfy any specification if they are invalid motivated the introduction of the failure state $\bot$ for *assertions*. For the trace modality, a failed assertion in the implementation should only yield a provable result if there is also a failed assertion in the specification. The program $q =$ x=-1; **assert**(x $\geq$ 0) (similar to p, but with an assertion instead of an assumption) evaluates to $\{\bot\}$, while the program $q' =$ x=-1; **assert**(x $\leq$ 0) evaluates to the same trace set as p'. The program $q'' =$ x=-1; **assert**(y $\geq$ 0), on the other hand, evaluates to

$$\begin{aligned}\mathrm{Tr}(q'') = \{&s' : s \in \mathcal{S} \wedge s' = s[\mathrm{x} \mapsto -1] \wedge \quad [\![\mathrm{y}]\!]_{s'} \geq 0\} \cup \\ &\{\bot : s \in \mathcal{S} \wedge s' = s[\mathrm{x} \mapsto -1] \wedge \neg[\![\mathrm{y}]\!]_{s'} \geq 0\}\end{aligned}$$

It also contains a set of traces for all concrete traces that after the assignment satisfy the assertion, but also a trace with the failure state if there is *any* state which does not satisfy the assertion. The use case for assertions is, in contrast to assumptions, to verify a fact that is *assumed to be provable* locally in the current program context.

The **havoc** command is used in loop invariant reasoning (see Example 10). For instance, $r =$ x=-1; **assert**(*Inv*); **havoc**; **assume**(*Inv*) evaluates to

$$\begin{aligned}\mathrm{Tr}(r) = \{&s \in \mathcal{S} : [\![\mathit{Inv}]\!]_s\} \cup \\ &\{\bot : s \in \mathcal{S} \wedge s' = s[\mathrm{x} \mapsto -1] \wedge \neg[\![\mathit{Inv}]\!]_{s'}\}\end{aligned} \qquad \Diamond$$

3 The Trace Modality

Table 1 (page III) provides a quick summary of our formalizations for the considered verification tasks in Sect. 3.

Remark 1 (Axioms of Modal and Dynamic Logic). The necessitation rule (axiom **N**) and distribution axiom (axiom **K**) of modal logic are theorems of the trace modality: If φ is a theorem, then $[\mathrm{p} \Vdash_{\alpha_{big}} \varphi]$ holds for all p (that do not produce a failure $\bot$) since $\mathit{lift}_{\mathrm{Fml}}(s)(\varphi)$ contains *all* traces starting in s (axiom **N**). It is also straightforward to show that $[\mathrm{p} \Vdash_{\alpha_{big}} \varphi \to \psi]$ implies $[\mathrm{p} \Vdash_{\alpha_{big}} \varphi] \to [\mathrm{p} \Vdash_{\alpha_{big}} \psi]$

Task	Problem	D_l	D_r	Solution Techniques (excerpt)
Partial Correctness	$\models [\mathtt{p} \Vdash_{\alpha_{big}} \mathit{Post}]$	$D_{\mathcal{L}_0}$	D_{Fml}	Symbolic execution, weakest precondition reasoning, Hoare calculus
Total Correctness	$\models \langle \mathtt{p} \Vdash_{\alpha_{big}} \mathit{Post} \rangle$	$D_{\mathcal{L}_0}$	D_{Fml}	ditto; plus reasoning about variant / ranking function
Information Flow	$\models [\mathtt{p}(\mathtt{l},\mathtt{h}) \Vdash_{\alpha_{\{\mathtt{l}\}} \circ \alpha_{big}} \mathtt{p}(\mathtt{l},\mathtt{h}')]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	Security type systems, Hoare calculus, symbolic execution
Information Flow with Declassification	$\models \bigwedge_{i=1}^{n}(e_i(\mathtt{l},\mathtt{h}) \doteq e_i(\mathtt{l},\mathtt{h}')) \rightarrow [\mathtt{p}(\mathtt{l},\mathtt{h}) \Vdash_{\alpha_{\{\mathtt{l}\}} \circ \alpha_{big}} \mathtt{p}(\mathtt{l},\mathtt{h}')]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	ditto
Finite Space MC	$s \models [\mathtt{p} \Vdash \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Automata constructions
Bounded MC	$\models [\mathtt{p} \Vdash \varphi]$	$D^k_{\mathcal{L}_0}$	D_{TL}	SMT solvers for checking encoded program paths
Abstraction-Based MC	$\models [\mathtt{p} \Vdash_{\alpha_d} \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Overapproximation techniques, CEGAR loops
Symbolic Execution-Based MC	$\models [\mathtt{p} \Vdash \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Invariant generation, k-induction
Bug Finding	$\models \langle \mathtt{p} \Vdash \neg\varphi \rangle$	$D_{\mathcal{L}_0}$	D_{TL}	All MC techniques; can be integrated with all abstractions
Program Synthesis	$\models [\mathtt{p}' \Vdash_{\alpha_{big}} \mathtt{p}]$	$D_{\mathcal{L}}$	$D_{\mathcal{L}}$	Proof-theoretic synthesis, proof mining
Correct compilation	$\models [\mathtt{p} \Vdash_{\alpha_{obs} \circ \alpha_{big}} \mathtt{c}]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}}$	Simultaneous symbolic execution, compiler extraction from executable HOL specifications
Program evolution / Bug fixing	$\models [\mathtt{p}_{\mathit{buggy}} \Vdash_{\alpha_{patch} \circ \alpha_{big}} \mathtt{p}_{\mathit{fixed}}]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	Manual program refinement, automatic software repair

Table 1: Modeling Different Verification Tasks with the Trace Modality

(axiom **K**). As an example for an axiom of PDL [17], we consider the axiom for a PDL "test" $\psi?$, where $\psi?$ corresponds to our **`assume`** ψ: The assertion $[\texttt{assume}\ \psi \Vdash_{\alpha_{big}} \varphi]$ is equivalent to $\psi \to \varphi$. If ψ does not hold for a state s, the trace set for the **`assume`** statement is the empty set which is trivially contained in any set, and the premise ψ of the implication is false and the implication therefore holds. If, however, s *does* satisfy ψ, then the abstracted trace set for the **`assume`** contains the trace starting and ending in s which is only contained in the set for the specification if s also satisfies φ, and similar for the implication.

Remark 2 (Linearization). As a small exercise, we can prove a theorem corresponding to the "linearization" axiom $[\mathtt{p};\mathtt{q}]\varphi \leftrightarrow [\mathtt{p}][\mathtt{q}]\varphi$ also for the trace modality. The trace modality version of this axiom is

$$[\mathtt{p};\mathtt{q} \Vdash_{\alpha_{big}} \varphi] \leftrightarrow [\mathtt{p} \Vdash_{\alpha_{big}} [\mathtt{q} \Vdash_{\alpha_{big}} \varphi]] \tag{1}$$

To give this a meaning, we define a lifting function $\mathit{lift}_{D_{mod}}$ assigning trace sets to trace modality formulas as follows:

$$\begin{aligned}&\mathit{lift}_{D_{mod}}(s)([\mathcal{C}_l \Vdash_\alpha \mathcal{C}_r]) \coloneqq \\ &\quad \left\{s\tau_1 : \mathit{finite}(\tau_1) \wedge \exists \tau_2 \in \mathit{lift}_{D_l}(\mathit{last}(s\tau_1))(\mathcal{C}_l);\ \alpha(s\tau_1\tau_2) \subseteq \alpha(\mathit{lift}_{D_r}(s)(\mathcal{C}_r))\right\} \\ &\qquad\qquad \cup \left\{s\tau : \neg\mathit{finite}(\tau) \wedge \alpha(s\tau) \subseteq \alpha(\mathit{lift}_{D_r}(s)(\mathcal{C}_r))\right\}\end{aligned}$$

In other words, the set of (i) all finite prefixes τ_1 starting in s which can completed by the traces τ_2 for the implementation starting in the last state of τ_1 such that the result meets the specification after α-abstraction, and (ii) all infinite traces starting in s which meet the specification after α-abstraction. Part (ii) may seem strange since the implementation $\mathcal{C}_l$ does not occur there. The idea is that for a nonterminating program p in Eq. (1), q is never evaluated neither on the left nor on the right-hand side of the equality. If φ is a first-order post condition, it will in any case only be lifted to finite traces, which is why (ii) does not apply then. We prove the "$\to$" direction of Eq. (1). For simplicity, we assume that p, q terminate normally for all inputs. Our hypothesis is that for all states s, the set $\alpha_{big}(\{\tau_1\tau_2\})$, where τ_1 corresponds to an execution of p starting in s and τ_2 to an execution of q starting in the final state of τ_1, is contained in the set of all pairs (s, f) where f satisfies φ. We have to show that the $\alpha_{big}(\{\tau_1\})$, where τ_1 is as before, is contained in the set of all pairs (s, f'), where f' is the final state of a prefix trace that can be completed by execution of q to a trace the final state of which satisfies φ. Since from the hypothesis, we already know that when execution q after p, the final state satisfies φ, f' can be instantiated to a final state of τ_1. Direction "$\to$" is similar.

We point out that the trace lifting of trace modality formulas in Remark 2 is interesting in its own, since it closely resembles the definition of the semantics of PDL modality formulas in [17]. There, the semantics of $[\mathtt{p}]\varphi$ is the set of all states s for which, when executing p starting in s, the final states satisfy φ. Our definition is similar, only that we generalize from single initial states to

whole "prefix traces". Another consideration is that we could have regarded trace modality formulas (in the specification side of a trace modality formula) simply as an atom of first-order logic. Then, $\mathit{lift}_{D_{\mathrm{Fml}}}(s)([\mathtt{q} \Vdash_{\alpha_{big}} \varphi])$ is the set of all finite traces starting in s whose final states satisfy $[\mathtt{q} \Vdash_{\alpha_{big}} \varphi]$, which is equivalent to the first part of the union in Remark 2–only infinite traces are not considered.

3.1 Functional Verification

Remark 3. Termination only is expressed as $\langle \mathtt{p} \Vdash_{\alpha_{big}} \mathrm{true} \rangle$: There has to be a α_{big}-abstracted trace for p starting in s which is *not* contained in the set $\alpha_{big}(\overline{\mathit{lift}_{\mathrm{Fml}}(s)(\mathrm{true})})$ consisting of all *infinite* traces starting in s (and all traces not starting in s).

3.2 Information Flow Analysis

Remark 4. In Example 2, we could have omitted α_{big}; however, we cannot do so in general, if we want to allow *intermediate* violation of the policy. If we, for instance, added a statement `l=42;` to the program, it would be safe w.r.t. the big-step formalization, although in between, `l` attains a different value according to the value of `h`.

Declassification, such as *delimited information release* [28], can be easily encoded via preconditions. Assume e is an expression of $\mathcal{L}_0$ we want to declassify. We extend $\mathcal{L}_0$ by expressions **`declassify`**(e), as in [28], which evaluate to e while permitting flow of e to the *low* level. As for programs, write $e(\mathtt{l},\mathtt{h})$ to make the variables occurring in e explicit. Then non-interference with declassification is formalized as:

$$\models e(\mathtt{l},\mathtt{h}) \doteq e(\mathtt{l},\mathtt{h}') \rightarrow [\mathtt{p}(\mathtt{l},\mathtt{h}) \Vdash_{\alpha_{\{\mathtt{l}\}} \circ \alpha_{big}} \mathtt{p}(\mathtt{l},\mathtt{h}')] \ .$$

Example 5 (Declassification). We consider the classic PIN example, where a *low* variable `OK` is set to `true` depending on whether a *high* input `inp` equals a *high* variable `pin` containing a PIN. Let

```
p := if (declassify(pin==inp)) { OK=true } else { OK=false }
```

be this program. If we do not give special semantics to the **`declassify`** expression, there is a state s where $s(\mathtt{pin}) = s(\mathtt{inp})$, but $s(\mathtt{pin}') \neq s(\mathtt{inp}')$; for this state, the subset relation does not hold, which is why p would be classified as insecure. The additional precondition $\mathtt{pin} \doteq \mathtt{inp} \leftrightarrow \mathtt{pin}' \doteq \mathtt{inp}'$, however, rules this choice out, and we can classify the program as secure w.r.t. the delimited release semantics. ◊

3.3 Software Model Checking

In the following, we formalize four popular Software Model Checking approaches using the trace modality.

Finite Space MC Finite space model checkers (SPIN [20] is a prominent representative) exhaustively explore the state space of an abstract program model. This implies that the analysis starts from *a concrete input state s* and that no unbounded data structures are involved. We can formalize this problem as $s \models [\mathtt{p} \Vdash \varphi]$, where φ is an LTL formula.

Bounded MC (BMC) BMC [6,8] handles unbounded data structures, but restricts the search space according to a predefined upper bound on the number of loop executions. This problem can simply be expressed as $\models [\mathtt{p} \Vdash \varphi]$ when using a domain $D^k_{\mathcal{L}_0}$ with lifting function $\mathit{lift}^k_{\mathcal{L}_0}$ that only produces traces up to a fixed number k of loop executions (and recursive method calls).

Abstraction-Based MC This variant of SMC applies data abstraction to limit the search space. We can express it as $\models [\mathtt{p} \Vdash_{\alpha_d} \varphi]$, where α_d is an abstract interpretation of the data types of `p`.

Symbolic Execution-Based MC This variant of SMC is similar to functional verification (Sect. 3.1). They mainly differ in the used abstraction (identity vs. big-step) and that in MC less complex properties are proved: $\models [\mathtt{p} \Vdash \varphi]$.

3.4 Program Synthesis

Example 6 discusses our formalization of the program synthesis problem within the trace modality along an example from the literature computing integer square roots.

Example 6. We consider the square root example from [30]. Given a user-defined specification $\mathit{Pre} \coloneqq \mathtt{x} \geq 1$, $\mathit{Post} \coloneqq \mathtt{i}^2 \leq \mathtt{x} < (\mathtt{i}+1)^2$ and scaffold program $\bullet; *(\bullet); \bullet$, the synthesizer should generate a program `IntSqrt(int x)` satisfying the specification (i.e., computing the integer square root of a strictly positive variable `x`) and matching the structure of the scaffold. An additional user-defined constraint is that, apart from `x` and `i`, there must be at most one additional variable `v`, also of integer type. Listing 1 shows a concrete program matching the specification. To apply our formalization, we first translate the scaffold in a schematic $\mathcal{L}$-Program: `P;while(b){Q};R`. Let now $\mathit{Syn}_{\mathtt{P/Q/R}}$ be synthesis conditions for `P`, `Q` and `R` inferred by the synthesizer. Note that $\mathit{Syn}_{\mathtt{Q}}$ is an inductive invariant for `Q`. A concrete instantiation for $\mathit{Syn}_{\mathtt{Q}}$ is $\mathtt{v} \doteq \mathtt{i}^2 \wedge \mathtt{x} \geq (\mathtt{i}-1)^2 \wedge \mathtt{i} \geq 1$. The scaffold annotated by assert statements for the synthesis conditions is depicted in Listing 2 (we write $\mathtt{x}'$ for the value of `x` before the execution of a schematic statement). Suppose that now we refine the scaffold `sc` to a program `p` by replacing `Q` *and* the following **`assert`** statement by the following program `q`: `v=v+2i+1;i++;`. To prove this correct, we have to show $\models \mathtt{i} \geq 1 \rightarrow [\mathtt{p} \Vdash_{\alpha_{big}} \mathtt{sc}]$. Since the traces for `sc` include one trace for each possible instantiation of `Q` satisfying $\mathit{Syn}_{\mathtt{Q}}$ in the given context, and `q` also satisfies $\mathit{Syn}_{\mathtt{Q}}$ in this context, this is true. We point out that we cannot instead show $\mathit{Pre} \rightarrow [\mathtt{q} \Vdash_{\alpha_{big}} \mathit{Syn}_{\mathtt{Q}}]$, since the program before the insertion position, i.e. already substituted concrete programs as well as asserted synthesis conditions, also has to be considered. Here, in particular, it is important that `v` and `i` initially are 1 for the invariant to hold. ◊

```
v = 1; i = 1;

while (v<=x) {
  v = v+2i+1;
  i++;
}
i = i-1;
```

Listing 1: IntSqrt

```
P;
assert(v ≐ 1 ∧ i ≐ 1 ∧ x ≐ x′);
while (v<=x) {
  Q;
  assert(v ≐ i² ∧ x ≥ (i − 1)² ∧ i ≥ 1);
}
R;
assert(i² ≤ x < (i + 1)²);
```

Listing 2: Annotated scaffold for IntSqrt

3.5 Correct Compilation

Remark 5. A popular approach realizing correct compilation is the specification of the compiler within the executable fragment of an interactive proof assistant like Isabelle or Coq, as done in CompCert [23]. We proposed in earlier work a rule-based technique using *simultaneous* Symbolic Execution [31] with a *dual modality*, which can be seen as a specialization of the trace modality. The interesting property of this framework is that compilation rules can be proven automatically based on Symbolic Execution calculi for the source and target language. We think that a similar technique might be applicable to different verification tasks.

3.6 Program Evolution & Bug Fixing

There are (at least) four formalizations of the problem of program evolution / but fixing, of which two variants ("bug abstractions" and "patch abstractions") already have been presented in Sect. 3.6. Alternatives are:

- Like in declassification (Sect. 3.2), an added precondition Pre_{safe} excludes buggy traces: $\models Pre_{safe} \rightarrow [\mathtt{p}_{bug} \Vdash_{\alpha_{big}} \mathtt{p}_{fixed}]$.
- One could prove the buggy and fixed program *in isolation* (as in functional correctness, Sect. 3.1); then, though, one cannot use techniques of relational program verification to exploit similarities between the two program versions. Also, the existence of full separate functional specifications is required.

The following example demonstrates the application of the formalizations using additional preconditions, "bug abstractions" and "patch abstractions".

Example 7. We explain the mentioned techniques along a simple example. The program $\mathtt{p}_{buggy} \coloneqq$ `if (x<-1) {x=-x;}` should compute the absolute of a given integer x; i.e., after execution of the program, x should be positive. However, the program contains a bug: The programmer misspelled the "<" operator which should be a "<=" instead. For the input -1, a wrong result is thus produced. Let $\mathtt{p}_{fixed}$ be the corrected program. We choose $Pre_{safe} \coloneqq \mathtt{x} \neq -1$, which excludes the buggy path. Then, $\models Pre_{safe} \rightarrow [\mathtt{p}_{buggy} \Vdash_{\alpha_{big}} \mathtt{p}_{fixed}]$ can be proven, since apart from that path, the traces of the programs coincide. Note that Pre_{safe} is in fact the negation of the path condition for the buggy path. Choosing the

second formalization, we can define $\alpha_{bug}(\mathcal{T}) \coloneqq \{\tau \in \mathcal{T} | \mathit{first}(\tau) \models \mathit{Pre}_{safe}\}$ and show $\models [\mathtt{p}_{buggy} \Vdash_{\alpha_{bug} \circ \alpha_{big}} \mathtt{p}_{fixed}]$, which is in this case equivalent. Indeed, the latter formalization is more flexible than the former and allows for a more systematic approach which not simply excluding buggy paths, but rather encoding the correction as a "patch". Let $\alpha_{patch}(\mathcal{T}) \coloneqq \{\mathit{patch}(\tau) | \tau \in \mathcal{T}\}$, where

$$\mathit{patch}(\tau) \coloneqq \begin{cases} (\mathit{first}(\tau), \mathit{last}(\tau)[\mathtt{x} \mapsto -\mathit{first}(\tau)(\mathtt{x})]) & \text{if } \mathit{first}(\tau)(\mathtt{x}) = -1 \\ \tau & \text{otherwise} \end{cases} \qquad \lozenge$$

As demonstrated by the example, the abstraction approach makes program evolution more explicit by describing the applied patch. Also, just excluding the buggy path would be too easy, since apart from that path, the buggy program is likely to be equivalent to the original one—we could show the correctness of a "fixed" program where no fix was applied at all, or the fix introduced new wrong behavior for the buggy path. We therefore choose the "patch abstraction" as the canonical representation for program evolution within the trace modality.

4 Reasoning about the Trace Modality

Example 8. Let, as in Example 3, $s_1 = (\mathtt{x} := 17 \,||\, \mathtt{y} := 42 \,||\, \mathtt{z} := 2, \mathrm{true})$. It is subsumed by $s_3 = (\mathit{sto}_{\mathrm{P}} \circ \mathtt{x} := c, C_{\mathrm{P}} \wedge c \geq 0)$ containing an abstract store and path condition: We can prove

$$\models \{\mathtt{x} := 17\} P(\mathtt{x}) \rightarrow \mathit{subst}(\{\mathtt{x} := 17\}(C_{\mathrm{P}} \wedge 17 \geq 0) \wedge \{\mathit{sto}_{\mathrm{P}} \circ \mathtt{x} := 17\} P(\mathtt{x}))$$

for any *subst* replacing C_{P} with true and $\mathit{sto}_{\mathrm{P}}$ with any concrete store. Subsumption can*not* be shown, for instance, for $s_4 = (\mathtt{x} := 17 \,||\, \mathtt{w} := c', c' \geq 0)$, since (SUB1) is violated. The symbolic state $s_5 = (\mathtt{x} \leq 0)$ does not subsume s_1, since the first conjunct under *subst* in (SUB2), $\{\mathtt{x} := 17\}\mathtt{x} \leq 0$, does not hold. $\lozenge$

Our algorith for the creation of SFAs from symbolic traces is shown in Algo. 2. For an alphabet Σ, write Σ_ε for its extension by instantaneous ε-transitions. We use standard ε-elimination to convert an SFA on Σ_ε to Σ. During SFA construction, we maintain a map L for assertion labels, mapping states to assertions that should hold when arriving at them. Labels are, in the post processing step addAssertionEdges, transformed to assertion edges leading to a failure state for input states not satisfying them.

Algos. 3 and 4 show the auxiliary algorithms for Algo. 1 computing over-approximating initial simulation relations and checking two symbolic states for subsumption. The algorithms are explained in Sect. 4.

The approach of deriving SSRs from starting from the cross product by repeated filtering is polynomial in the size of the automata; we refer to [26] for more efficient approaches.

Symbolic Lifting of Loops Our symbolic trace language is a *regular* language, therefore it generally is not possible to encode loops with full precision. For this, it would be necessary to maintain a mutable state for tracking changes made

Algorithm 2 Creation of SFA from Symbolic Traces

function createSFA(τ : SymTr)
  $q_0 \leftarrow$ fresh state, $\Sigma \leftarrow \mathcal{S}$
  $(L, (Q, \Sigma_\varepsilon, \delta, q_0, F)) \leftarrow$ extendSFA$(\tau, q_0, \{q_0\}, \Sigma_\varepsilon, q_0)$
  $(L', (Q', \Sigma, \delta', q_0, F')) \leftarrow$ eliminateEpsilonTransitions$(L, (Q, \Sigma_\varepsilon, \delta, q_0, F))$
    ▷ Standard, but preserve labels
  return addAssertionEdges$(L', (Q', \Sigma, \delta', q_0, F'))$
end function

Ensure: Returns a pair $(L, (Q', \Sigma_\varepsilon, \delta, q_0, F))$ of a set of labels and an SFA for τ, where $L \subseteq Q' \times \mathrm{Fml}$, $Q' \supseteq Q$, $\delta \subseteq Q' \times \mathrm{SymState} \times Q'$, $F \subseteq Q'$. Σ_ε and q_0 are not changed.
function extendSFA(τ, q, Q, Σ_ε, q_0)
  if $\tau = s$ **then**
    $q' \leftarrow$ fresh state, $q' \notin Q$
    $Q' \leftarrow Q \cup \{q'\}$
    $\delta \leftarrow \{(q, s, q')\}$
    $F \leftarrow \{q'\}$
    return $(L, (Q', \Sigma_\varepsilon, \delta, q_0, F))$
  else if $\tau = \tau_1; \tau_2$ **then**
    $(L_1, (Q_1, \Sigma_\varepsilon, \delta_1, q_0, F_1)) \leftarrow$ extendSFA$(\tau_1, q, Q, \Sigma_\varepsilon, q_0)$
    $Q' \leftarrow Q_1$, $\delta \leftarrow \delta_1$, $F \leftarrow \emptyset$, $L \leftarrow L_1$
    for all $q' \in F_1$ **do**
      $(L_2, (Q_2, \Sigma_\varepsilon, \delta_2, q_0, F_2)) \leftarrow$ extendSFA$(\tau_2, q', Q', \Sigma_\varepsilon, q_0)$
      $Q' \leftarrow Q' \cup Q_2$, $\delta \leftarrow \delta \cup \delta_2$, $F \leftarrow F \cup F_2$, $L \leftarrow L \cup L_2$
    end for
    return $(L, (Q', \Sigma_\varepsilon, \delta, q_0, F))$
  else if $\tau = \tau_1 + \tau_2$ **then**
    $(L_i, (Q_i, \Sigma_\varepsilon, \delta_i, q_0, F_i)) \leftarrow$ extendSFA$(\tau_i, q, Q, \Sigma_\varepsilon, q_0)$, $i = 1, 2$
    return $(L_1 \cup L_2, (Q_1 \cup Q_2, \Sigma_\varepsilon, \delta_1 \cup \delta_2, q_0, F_1 \cup F_2))$
  else if $\tau = \varphi^!$ **then**
    return $(L \cup \{(q, \varphi)\}, (Q, \Sigma_\varepsilon, \emptyset, q_0, \emptyset))$
  else if $\tau = (\tau')^*$ **then**
    $(L, (Q', \Sigma_\varepsilon, \delta, q_0, F)) \leftarrow$ extendSFA$(\tau', q, Q, \Sigma_\varepsilon, q_0)$
    $\delta' \leftarrow \delta \cup \{(q', \varepsilon, q) : q' \in F\}$
    return $(L, (Q', \Sigma_\varepsilon, \delta', q_0, F \cup \{q\}))$
  end if
end function

function addAssertionEdges$(L, (Q, \Sigma, \delta, q_0, F))$
  $q_{fail} \leftarrow$ fresh state, $\delta' \leftarrow \delta$
  for all $(q, \varphi) \in L$ **do**
    for all $(q', (sto, \varphi'), q) \in \delta$ **do** ▷ similarly for φ'-only transitions
      $\delta' \leftarrow \delta' \setminus \{(q', (sto, \varphi'), q)\}$
      $s \leftarrow (sto, \varphi' \wedge \varphi)$
      $\delta' \leftarrow \delta' \cup \{(q', s, q), (q', \neg\varphi, q_{fail})\}$
    end for
  end for
  return $(Q \cup \{q_{fail}\}, \Sigma, \delta', q_0, F \cup \{q_{fail}\})$
end function

Algorithm 3 Construction of Initial Simulation Relation

function INITSIM(Q_l, Q_r, δ_l, δ_r)
  $R \leftarrow Q_l \times Q_r$, *changed* $\leftarrow$ *true*
  while *changed* = *true* **do**
    changed $\leftarrow$ *false*
    for all $(q_l, q_r) \in R$, $(q_l, s, q_l') \in \delta_l$ **do**
      if $\neg\exists (q_r, s', q_r') \in \delta_r$ **then** $R \leftarrow R \setminus (q_l, q_r)$, *changed* $\leftarrow$ *true* **end if**
    end for
  end while
  return R
end function

Algorithm 4 Subsumption Checking

function SUBSUMPTION(s, s', *substs*)
  substs$' \leftarrow \emptyset$
  for all *subst* $\in$ *substs* **do**
    substs$' \leftarrow$ *substs*$' \cup \{$*subst* $\circ$ *subst*$' \mid$ *subst*$'$ such that $s \sqsubseteq_{subst'}$ *subst*$(s')\}$
  end for
  return *substs*$'$
end function

inside the loop body, which then could be evaluated to decide whether to continue or leave the loop. Basically, we can within our framework apply the same solutions as known from *symbolic execution*: (Bounded) loop unwinding and invariant reasoning. The simplest solution is loop unwinding, by which symbolic lifting can be defined as follows:

$$\mathit{slift}_{\mathcal{L}_0}(\mathit{sto}, \varphi)(\texttt{while(b) } p) \coloneqq \mathit{slift}_{\mathcal{L}_0}(\mathit{sto}, \varphi)(\texttt{if(b) } p \texttt{ ; while(b) } p)$$

This, however, does generally not terminate for loops with symbolic guards. In the context of BMC, the bounded lifting function $\mathit{slift}^k_{\mathcal{L}_0}$ would unwind the loop as above exactly k times and then remove the loop statement. A standard approach in deductive program verification is to use *loop invariants*. Let $\mathit{Inv} \in$ Fml. Then, we can replace a loop `while(b)` p by the following program:

assert *Inv*; **havoc**; **assume** *Inv*; **if** (b) { p; **assert** *Inv* }

The **havoc** statement erases the state: $\mathit{slift}_{\mathcal{L}_0}(\mathit{sto}, \varphi)(\texttt{havoc}) \coloneqq \text{true}$. The replacement as above is sound for the left side of the trace modality: If *Inv* is not an invariant, the program evaluates to a failure trace, which can only be part of the traces for the right side if there also occurs a failed assertion. Otherwise, the new program evaluates to *at least* the same traces as the old one. It would also be sound *terminate* the trace after the assertion in the **if**, either by an **exit** statement or by wrapping the remaining program in an **else** block.

Reasoning with Symbolic Traces In the following, we provide two examples (Examples 9 and 10) to establish an intuition about how to solve the problem of symbolic trace subsumption. This is meant to support understanding the prin-

ciples behind Algo. 1, although the algorithm itself is not directly used there. Example 11 after that applies Algo. 1 to the problem of Example 9.

For the subsequent example, we assume that our term language has a conditional operator $\varphi\,?\,t_1 : t_2$, intuitively evaluating to the value of t_1 if φ holds and otherwise to that of t_2.

Example 9 (Functional Verification). Consider the following program p computing the difference of two integers a and b:

"`res=0; if (b < a) { tmp=a; a=b; b=tmp; } res=b-a`"

For this program, we want to show the post condition $\varphi \coloneqq \texttt{res} \geq 0$, i.e., $\models [\mathtt{p} \Vdash_{\alpha_{big}} \varphi]$. We first compute the symbolic traces by symbolic lifting, starting from an initial store $st_0 \coloneqq \mathtt{a} := a_0 \,||\, \mathtt{b} := b_0 \,||\, \mathtt{res} := res_0 \,||\, \mathtt{tmp} := tmp_0$:

$$
\begin{aligned}
&\mathit{slift}_{\mathcal{L}_0}(st_0, \text{true})(\mathtt{p}) = \\
&\quad (st_0 \circ (\,\mathtt{res} := 0\,), \text{true}); \\
&\quad \big(\quad ((st_0 \circ (\mathtt{res} := 0),\, b_0 < a_0\,); \\
&\qquad (st_0 \circ (\mathtt{res} := 0 \,||\, \mathtt{tmp} := \mathtt{a}\,), b_0 < a_0); \\
&\qquad (b := b_0 \,||\, \mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0\,, b_0 < a_0); \\
&\qquad (\mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0 \,||\, \mathtt{b} := a_0\,, b_0 < a_0); \\
&\qquad (\mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0 \,||\, \mathtt{b} := a_0 \,||\, \mathtt{res} := a_0 - b_0\,, b_0 < a_0)) \\
&\quad + ((st_0 \circ (\mathtt{res} := 0),\, b_0 \geq a_0\,); \\
&\qquad (\mathtt{a} := a_0 \,||\, \mathtt{b} := b_0 \,||\, \mathtt{tmp} := tmp_0 \,||\, \mathtt{res} := b_0 - a_0\,, b_0 \geq a_0))\big)
\end{aligned}
$$

During symbolic lifting, we simultaneously simplified the symbolic state, e.g., the update $st_0 \circ (\mathtt{res} := 0 \,||\, \mathtt{tmp} := \mathtt{a}) \circ (\mathtt{a} := \mathtt{b})$ is simplified to $\mathtt{b} := b_0 \,||\, \mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0$. Note that alternatively, one could use state merging to create an equivalent trace with only one final state:

$$
\begin{aligned}
\mathit{slift}_{\mathcal{L}_0}(st_0, \text{true})(\mathtt{p}) = (st_0 \circ (\mathtt{res} := 0), \text{true}); \\
\big(\quad ((st_0 \circ (\mathtt{res} := 0), b_0 < a_0); \\
(st_0 \circ (\mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0), b_0 < a_0); \\
(st_0 \circ (\mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0), b_0 < a_0); \\
(\mathtt{res} := 0 \,||\, \mathtt{tmp} := a_0 \,||\, \mathtt{a} := b_0 \,||\, \mathtt{b} := a_0, b_0 < a_0)) \\
+ (st_0 \circ (\mathtt{res} := 0), b_0 \geq a_0)\big); \\
(\mathtt{tmp} := (b_0 < a_0)\,?\,a_0 : tmp_0 \,|| \\
\mathtt{a} := (b_0 < a_0)\,?\,b_0 : a_0 \,|| \\
\mathtt{b} := (b_0 < a_0)\,?\,a_0 : b_0 \,||
\end{aligned}
$$

$$\texttt{res} := (b_0 < a_0)\,?\,a_0 - b_0 : b_0 - a_0, \text{true})$$

Thus, it is always possible (for deterministic programs) to create symbolic traces with exactly one final state. Symbolically lifting the post condition leads to $\mathit{slift}_{\text{Fml}}(st_0, \text{true})(\varphi) = \text{true}^*; (\texttt{res} \geq 0)$. Applying big step abstraction, we have to show that the symbolic trace $\texttt{res} \geq 0$ subsumes both traces $(\texttt{tmp} := a_0 \,||\, \texttt{a} := b_0 \,||\, \texttt{b} := a_0 \,||\, \texttt{res} := a_0 - b_0, b_0 < a_0)$ and $(\texttt{res} := b_0 - a_0, b_0 \geq a_0)$. We can do so by evaluating the following formula in a theorem prover or SMT solver:

$$\begin{aligned}(b_0 < a_0 \rightarrow \{\texttt{tmp} := a_0 \,||\, \texttt{a} := b_0 \,||\, \texttt{b} := a_0 \,||\, \texttt{res} := a_0 - b_0\}\texttt{res} \geq 0) \,\wedge\\ (b_0 \geq a_0 \rightarrow \{\texttt{res} := b_0 - a_0\}\texttt{res} \geq 0)\end{aligned}$$

which is equivalent to $(b_0 < a_0 \rightarrow a_0 - b_0 \geq 0) \wedge (b_0 \geq a_0 \rightarrow b_0 - a_0 \geq 0)$ and therefore valid, which is why $\models [\texttt{p} \Vdash_{\alpha_{big}} \varphi]$ also holds. In this example, we also could have shown the assertion $\Box\varphi$, i.e. that the trace $\mathit{slift}_{TR}(st_0, \text{true})(\Box\varphi) = \varphi^*$ subsumes the symbolic trace of p, since φ holds for all intermediate states. ◊

The above example demonstrated how to match symbolic traces for a program and a post condition after big step abstraction. In the following, we investigate the situation for two *programs* in a program synthesis setting.

Example 10 (Synthesis). Consider `IntSqrt` and its scaffold for synthesis from Listings 1 and 2. We aim to show that the program p is a specialization of the scaffold sc, i.e., that $\models [\texttt{p} \Vdash \texttt{sc}]$. Since we already have an invariant at hand, we can use invariant reasoning to handle loops. Before, we mentioned that this is generally unsound for the right-hand side of the trace modality, since it constitutes an abstraction. However, we apply the same abstraction on the left and right-hand side, which is why this technique is admissible. Additionally, one has to check that the invariant is not unsatisfiable, since otherwise, both sides evaluate to the failure state and the property can be shown trivially. Let $st_0 = (\texttt{x} := x_0 \,||\, \texttt{v} := v_0 \,||\, \texttt{i} := i_0)$ be an initial store. The invariant is $\mathit{Inv}(\texttt{v}, \texttt{i}, \texttt{x}) := \texttt{v} = \texttt{i}^2 \wedge \texttt{x} \geq (\texttt{i} - 1)^2 \wedge \texttt{i} \geq 1$. We obtain the following symbolic traces (to ease the presentation, we omit abstract path conditions for schematic statements):

$$\begin{aligned}&\mathit{slift}_{\mathcal{L}}(st_0, x_0 \geq 1)(\texttt{p}) =\\
&\quad (st_0 \circ (\,\texttt{v} := 1\,), x_0 \geq 1);\\
&\quad (\texttt{x} := x_0 \,||\, \texttt{v} := 1 \,||\, \texttt{i} := 1, x_0 \geq 1);\\
&\quad (x_0 \geq 1 \rightarrow 1 = 1^2 \wedge x_0 \geq (1-1)^2 \wedge 1 \geq 1)^!;\\
&\quad (\texttt{x} := x_1 \,||\, \texttt{v} := v_1 \,||\, \texttt{i} := i_1, x_0 \geq 1);\\
&\quad (\texttt{x} := x_1 \,||\, \texttt{v} := v_1 \,||\, \texttt{i} := i_1, x_0 \geq 1 \wedge \mathit{Inv}(v_1, i_1, x_1));\\
&\quad \big(((\texttt{x} := x_1 \,||\, \texttt{v} := v_1 \,||\, \texttt{i} := i_1, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1);\\
&\qquad (\texttt{x} := x_1 \,||\, \texttt{i} := i_1 \,||\, \texttt{v} := v_1 + 2i_1 + 1, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1);\end{aligned}$$

$$
\begin{array}{l}
(\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 + 2i_1 + 1 \,||\, \mathtt{i} := i_1 + 1 \,, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1); \\
((\mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \mathit{Inv}(v_1 + 2i_1 + 1, i_1 + 1, x_1))^!) \; + \\
((\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1, \mathit{Inv}(v_1, i_1, x_1) \wedge \; v_1 > x_1 \;); \\
(\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1 - 1 \,, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1)))
\end{array}
$$

$$
\begin{array}{l}
\mathit{slift}_{\mathcal{L}}(st_0, x_0 \geq 1)(\mathtt{sc}) = \\
\quad \text{true}^*; (st_0 \circ \mathit{sto}_{\mathtt{P}} \,, x_0 \geq 1); \\
\quad (x_0 \geq 1 \rightarrow \{st_0 \circ \mathit{sto}_{\mathtt{P}}\}(\mathtt{v} \doteq 1 \wedge \mathtt{i} \doteq 1 \wedge \mathtt{x} \doteq 1))^!; \\
\quad (x_0 \geq 1 \rightarrow \{st_0 \circ \mathit{sto}_{\mathtt{P}}\}\mathit{Inv}(\mathtt{v}, \mathtt{i}, \mathtt{x}))^!; \\
\quad (\mathtt{x} := \; x_1 \; || \, \mathtt{v} := \; v_1 \; || \, \mathtt{i} := \; i_1 \;, x_0 \geq 1); \\
\quad (\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1, x_0 \geq 1 \wedge \; \mathit{Inv}(v_1, i_1, x_1) \;); \\
\quad \big(((\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1, \mathit{Inv}(v_1, i_1, x_1) \wedge \; v_1 \leq x_1 \;); \\
\qquad \text{true}^*; ((\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1) \circ \; \mathit{sto}_{\mathtt{Q}} \,, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1); \\
\qquad ((\mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \\
\qquad\quad \{(\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1) \circ \mathit{sto}_{\mathtt{Q}}\}\mathit{Inv}(\mathtt{v}, \mathtt{i}, \mathtt{x}))^!; \\
\qquad ((\mathit{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \\
\qquad\quad\; \{(\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1) \circ \mathit{sto}_{\mathtt{Q}}\}\mathit{Inv}(\mathtt{v}, \mathtt{i}, \mathtt{x}))^!) \; + \\
\quad ((\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1, \mathit{Inv}(v_1, i_1, x_1) \wedge \; v_1 > x_1 \;); \\
\qquad \text{true}^*; \\
\qquad ((\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1) \circ \; \mathit{sto}_{\mathtt{R}} \,, \mathit{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1); \\
\qquad (((\mathit{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1) \rightarrow \\
\qquad\quad \{(\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1) \circ \mathit{sto}_{\mathtt{R}}\}(\; \mathtt{i}^2 \leq \mathtt{x} < (\mathtt{i} + 1)^2 \;))^!))
\end{array}
$$

We point out that we performed the same anonymization $\mathtt{x} := x_1 \,||\, \mathtt{v} := v_1 \,||\, \mathtt{i} := i_1$ for the **`havoc`** statement in both traces; this corresponds to explicit loop coupling. Alternatively, we could have left the job of finding suitable substitutions to the subsumption checker. Now, we have to decide whether $(\mathit{slift}_{\mathcal{L}}(st_0, x_0 \geq 1)(\mathtt{p})$ is subsumed by $\mathit{slift}_{\mathcal{L}}(st_0, x_0 \geq 1)(\mathtt{sc}))$. We do not explicitly construct SFAs for this example, but follow the symbolic trace of `p`, mapping it to that of `sc` in a "lock-step" approach. This corresponds to an "on-the-fly" simulation construction process.

The state $(st_0 \circ (\mathtt{v} := 1), x_0 \geq 1)$, for instance, is subsumed by true^*, and the state $(\mathtt{x} := x_0 \,||\, \mathtt{v} := 1 \,||\, \mathtt{i} := 1, x_0 \geq 1)$ by $(st_0 \circ \mathit{sto}_{\mathtt{P}}, x_0 \geq 1)$, since $\mathit{sto}_{\mathtt{P}}$ is yet uninstantiated and can therefore be instantiated to $\mathtt{v} := 1 \,||\, \mathtt{i} := 1$. Next are the

assertions before the loop: Since based on the instantiation of $sto_{\mathtt{p}}$ all of those are satisfied, the failure trace is neither in the left, nor in the right trace set, which is why we can continue. This way, we process the whole symbolic trace and finally find a simulation relation. ◊

Example 11 (Applying Algo. 1). We consider the example from Example 9. Algo. 1 first lifts the implementation p and specification φ to symbolic traces; this is already done in Example 9 (we consider the version with state merging for p). Then, it creates SFAs for p and φ (Algo. 2). Those automata are shown in Figs. 2 and 3. Now, we have to find the simulation relation (Algo. 1). The initial simulation produced by INITSIM is $(\{q_0, q_1, q_2, q_3, q_4, q_5\} \times \{q_6\}) \cup \{(q_5, q_7)\}$. The pair (q_3, q_7), for instance, is not contained, since there is an outgoing edge from q_3, but not from q_7. If the pair (q_0, q_6) was not contained in the initial simulation, we could stop here since property (SR2) would not be satisfied. The subsequent subsumption checking steps performed by FINDSSR do not eliminate any pair from this relation, since all symbolic states in the implementation are subsumed by true in the specification, and the state on the transition from q_3 to q_5 satisfies the post condition $\mathtt{res} \geq 0$. The algorithm has found a relation also satisfying (SR2) and returns YES. If, however, the program would set res simply to, say, -1 in a last step, then the symbolic state in the last transition, $(\mathtt{res} := -1, \mathrm{true})$ would not be subsumed by $\mathtt{res} \geq 0$ and the algorithm would return UNKNOWN. ◊

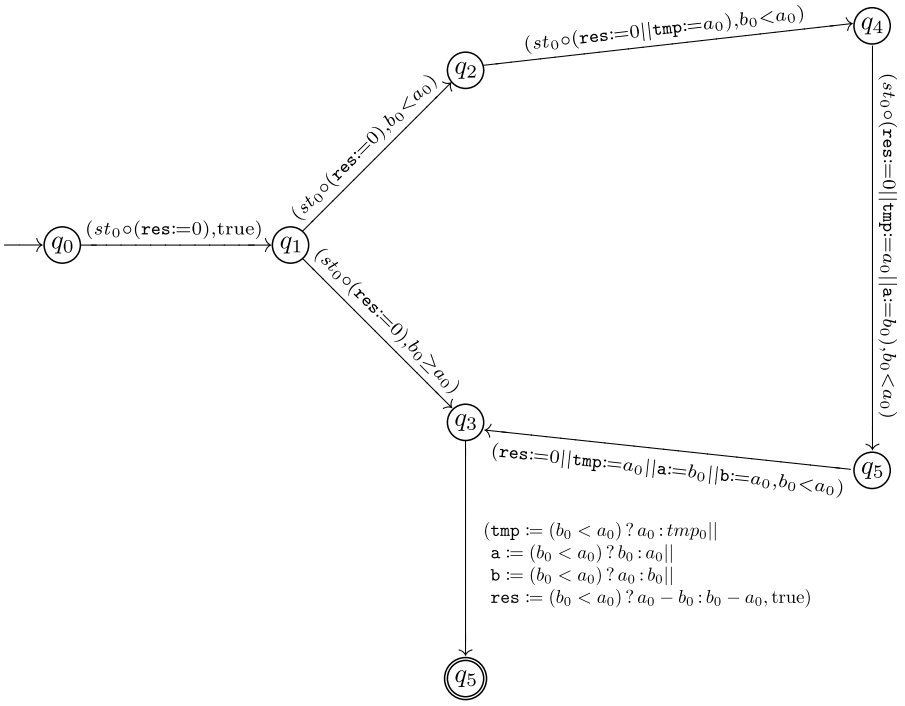

Fig. 2: Implementation SFA for Example 9

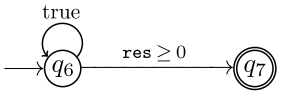

Fig. 3: Specification SFA for Example 9