

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 $\llbracket e \rrbracket_s$ 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 $failed(\tau)$ holds for a trace τ if it contains the failure state \perp .

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

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

assume($x \leq 0$) evaluates to the set $\{s[x \mapsto -1] : s \in \mathcal{S}\}$, since now, the empty trace is produced for the **assume** statement, which is the neutral element of trace concatenation. The trace modality formula $[p \Vdash spec]$ 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: $p'' = x=-1; \mathbf{assume}(y \geq 0)$. There, we make an assumption about a variable y that is not mentioned before. It evaluates to

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

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 \perp 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 $q = x=-1; \mathbf{assert}(x \geq 0)$ (similar to p , but with an assertion instead of an assumption) evaluates to $\{\perp\}$, while the program $q' = x=-1; \mathbf{assert}(x \leq 0)$ evaluates to the same trace set as p' . The program $q'' = x=-1; \mathbf{assert}(y \geq 0)$, on the other hand, evaluates to

$$\begin{aligned} \text{Tr}(q'') = & \{s' : s \in \mathcal{S} \wedge s' = s[x \mapsto -1] \wedge \llbracket y \rrbracket_{s'} \geq 0\} \cup \\ & \{\perp : s \in \mathcal{S} \wedge s' = s[x \mapsto -1] \wedge \neg \llbracket y \rrbracket_{s'} \geq 0\} \end{aligned}$$

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, $r = x=-1; \mathbf{assert}(Inv); \mathbf{havoc}; \mathbf{assume}(Inv)$ evaluates to

$$\begin{aligned} \text{Tr}(r) = & \{s \in \mathcal{S} : \llbracket Inv \rrbracket_s\} \cup \\ & \{\perp : s \in \mathcal{S} \wedge s' = s[x \mapsto -1] \wedge \neg \llbracket Inv \rrbracket_{s'}\} \quad \diamond \end{aligned}$$

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 $[p \Vdash_{\alpha_{big}} \varphi]$ holds for all p (that do not produce a failure \perp) since $lift_{\text{Fml}}(s)(\varphi)$ contains *all* traces starting in s (axiom **N**). It is also straightforward to show that $[p \Vdash_{\alpha_{big}} \varphi \rightarrow \psi]$ implies $[p \Vdash_{\alpha_{big}} \varphi] \rightarrow [p \Vdash_{\alpha_{big}} \psi]$

Task	Problem	D_l	D_r	Solution Techniques (excerpt)
Partial Correctness	$\models [p \Vdash_{\alpha_{big}} Post]$	$D_{\mathcal{L}_0}$	D_{Fmi}	Symbolic execution, weakest precondition reasoning, Hoare calculus
Total Correctness	$\models \langle p \Vdash_{\alpha_{big}} Post \rangle$	$D_{\mathcal{L}_0}$	D_{Fmi}	ditto; plus reasoning about variant / ranking function
Information Flow	$\models [p(\mathbb{1}, h) \Vdash_{\alpha_{\{1\}} \circ \alpha_{big}} p(\mathbb{1}, h')]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	Security type systems, Hoare calculus, symbolic execution
Information Flow with Declassification	$\models \bigwedge_{i=1}^n (e_i(\mathbb{1}, h) \doteq e_i(\mathbb{1}, h')) \rightarrow [p(\mathbb{1}, h) \Vdash_{\alpha_{\{1\}} \circ \alpha_{big}} p(\mathbb{1}, h')]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	ditto
Finite Space MC	$s \models [p \Vdash \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Automata constructions
Bounded MC	$\models [p \Vdash \varphi]$	$D_{\mathcal{L}_0}^k$	D_{TL}	SMT solvers for checking encoded program paths
Abstraction-Based MC	$\models [p \Vdash_{\alpha_d} \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Overapproximation techniques, CEGAR loops
Symbolic Execution- Based MC	$\models [p \Vdash \varphi]$	$D_{\mathcal{L}_0}$	D_{TL}	Invariant generation, k -induction
Bug Finding	$\models \langle p \Vdash \neg \varphi \rangle$	$D_{\mathcal{L}_0}$	D_{TL}	All MC techniques; can be integrated with all abstractions
Program Synthesis	$\models [p' \Vdash_{\alpha_{big}} p]$	$D_{\mathcal{L}}$	$D_{\mathcal{L}}$	Proof-theoretic synthesis, proof mining
Correct compilation	$\models [p \Vdash_{\alpha_{obs} \circ \alpha_{big}} c]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}}$	Simultaneous symbolic execution, compiler extraction from executable HOL specifications
Program evolution / Bug fixing	$\models [p_{buggy} \Vdash_{\alpha_{patch} \circ \alpha_{big}} p_{fixed}]$	$D_{\mathcal{L}_0}$	$D_{\mathcal{L}_0}$	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 $[\mathbf{assume} \ \psi \Vdash_{\alpha_{big}} \varphi]$ is equivalent to $\psi \rightarrow \varphi$. 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 $[p; q]\varphi \leftrightarrow [p][q]\varphi$ also for the trace modality. The trace modality version of this axiom is

$$[p; q] \Vdash_{\alpha_{big}} \varphi \leftrightarrow [p \Vdash_{\alpha_{big}} [q \Vdash_{\alpha_{big}} \varphi]] \quad (1)$$

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

$$\begin{aligned} lift_{D_{mod}}(s)([C_l \Vdash_{\alpha} C_r]) := \\ \{s\tau_1 : finite(\tau_1) \wedge \exists \tau_2 \in lift_{D_l}(last(s\tau_1))(C_l); \alpha(s\tau_1\tau_2) \subseteq \alpha(lift_{D_r}(s)(C_r))\} \\ \cup \{s\tau : \neg finite(\tau) \wedge \alpha(s\tau) \subseteq \alpha(lift_{D_r}(s)(C_r))\} \end{aligned}$$

In other words, the set of (i) all finite prefixes τ_1 starting in s which can be 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 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 “ \rightarrow ” 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 “ \rightarrow ” 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 $[p]\varphi$ 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, $\text{lift}_{D_{\text{Fml}}}(s)([q \Vdash_{\alpha_{\text{big}}} \varphi])$ is the set of all finite traces starting in s whose final states satisfy $[q \Vdash_{\alpha_{\text{big}}} \varphi]$, 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 $\langle p \Vdash_{\alpha_{\text{big}}} \text{true} \rangle$: There has to be a α_{big} -abstracted trace for p starting in s which is *not* contained in the set $\alpha_{\text{big}}(\overline{\text{lift}_{\text{Fml}}(s)(\text{true})})$ 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 $\mathbf{declassify}(e)$, as in [28], which evaluate to e while permitting flow of e to the *low* level. As for programs, write $e(l, h)$ to make the variables occurring in e explicit. Then non-interference with declassification is formalized as:

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

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 $s(\text{pin}) = s(\text{inp})$, but $s(\text{pin}') \neq s(\text{inp}')$; for this state, the subset relation does not hold, which is why p would be classified as insecure. The additional precondition $\text{pin} \doteq \text{inp} \leftrightarrow \text{pin}' \doteq \text{inp}'$, however, rules this choice out, and we can classify the program as secure w.r.t. the delimited release semantics. \diamond

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 $s \models [p \Vdash \varphi]$, 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 $\models [p \Vdash \varphi]$ when using a domain $D_{\mathcal{L}_0}^k$ with lifting function $lift_{\mathcal{L}_0}^k$ 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 $\models [p \Vdash_{\alpha_d} \varphi]$, 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: $\models [p \Vdash \varphi]$.

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 $Pre := x \geq 1$, $Post := i^2 \leq x < (i + 1)^2$ and scaffold program $\bullet; * (\bullet); \bullet$, the synthesizer should generate a program $\text{IntSqrt}(\text{int } x)$ 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: $P; \mathbf{while}(b)\{Q\}; R$. Let now $Syn_{P/Q/R}$ be synthesis conditions for P , Q and R inferred by the synthesizer. Note that Syn_Q is an inductive invariant for Q . A concrete instantiation for Syn_Q is $v \doteq i^2 \wedge x \geq (i - 1)^2 \wedge i \geq 1$. The scaffold annotated by assert statements for the synthesis conditions is depicted in Listing 2 (we write 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 $\models i \geq 1 \rightarrow [p \Vdash_{\alpha_{big}} sc]$. Since the traces for sc include one trace for each possible instantiation of Q satisfying Syn_Q in the given context, and q also satisfies Syn_Q in this context, this is true. We point out that we cannot instead show $Pre \rightarrow [q \Vdash_{\alpha_{big}} Syn_Q]$, 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. \diamond

```

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: $\models Pre_{safe} \rightarrow [p_{bug} \Vdash_{\alpha_{big}} p_{fixed}]$.
- 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 $p_{buggy} := \mathbf{if} (x < -1) \{x = -x;\}$ 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 p_{fixed} be the corrected program. We choose $Pre_{safe} := x \neq -1$, which excludes the buggy path. Then, $\models Pre_{safe} \rightarrow [p_{buggy} \Vdash_{\alpha_{big}} p_{fixed}]$ 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 $\alpha_{bug}(\mathcal{T}) := \{\tau \in \mathcal{T} \mid first(\tau) \models Pre_{safe}\}$ and show $\models [p_{buggy} \Vdash_{\alpha_{bug} \circ \alpha_{big}} P_{fixed}]$, 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 $\alpha_{patch}(\mathcal{T}) := \{patch(\tau) \mid \tau \in \mathcal{T}\}$, where

$$patch(\tau) := \begin{cases} (first(\tau), last(\tau)[x \mapsto -first(\tau)(x)]) & \text{if } first(\tau)(x) = -1 \\ \tau & \text{otherwise} \end{cases} \quad \diamond$$

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, $s_1 = (x := 17 \parallel y := 42 \parallel z := 2, true)$. It is subsumed by $s_3 = (sto_p \circ x := c, C_p \wedge c \geq 0)$ containing an abstract store and path condition: We can prove

$$\models \{x := 17\}P(x) \rightarrow subst(\{x := 17\}(C_p \wedge 17 \geq 0) \wedge \{sto_p \circ x := 17\}P(x))$$

for any *subst* replacing C_p with *true* and sto_p with any concrete store. Subsumption *cannot* be shown, for instance, for $s_4 = (x := 17 \parallel w := c', c' \geq 0)$, since (SUB1) is violated. The symbolic state $s_5 = (x \leq 0)$ does not subsume s_1 , since the first conjunct under *subst* in (SUB2), $\{x := 17\}x \leq 0$, does not hold. \diamond

Our algorithm 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( $\tau : \text{SymTr}$ )
   $q_0 \leftarrow$  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 $L \subseteq Q' \times \text{Fml}$, $Q' \supseteq Q$, $\delta \subseteq Q' \times \text{SymState} \times Q'$, $F \subseteq Q'$. Σ_ε and q_0 are not changed.

```

function EXTENDSFA( $\tau, q, Q, \Sigma_\varepsilon, q_0$ )
  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  $\varphi'$ -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( $Q_l, Q_r, \delta_l, \delta_r$ )
   $R \leftarrow Q_l \times Q_r$ ,  $changed \leftarrow true$ 
  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( $s, s', substs$ )
   $subst_s' \leftarrow \emptyset$ 
  for all  $subst \in substs$  do
     $subst_s' \leftarrow subst_s' \cup \{subst \circ subst' \mid subst' \text{ such that } s \sqsubseteq_{subst'} subst(s')\}$ 
  end for
  return  $subst_s'$ 
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:

$$slift_{\mathcal{L}_0}(sto, \varphi)(\mathbf{while}(b) p) := slift_{\mathcal{L}_0}(sto, \varphi)(\mathbf{if}(b) p; \mathbf{while}(b) p)$$

This, however, does generally not terminate for loops with symbolic guards. In the context of BMC, the bounded lifting function $slift_{\mathcal{L}_0}^k$ 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 $Inv \in \text{Fml}$. Then, we can replace a loop $\mathbf{while}(b) p$ by the following program:

$$\mathbf{assert} \text{ } Inv; \mathbf{havoc}; \mathbf{assume} \text{ } Inv; \mathbf{if} (b) \{ p; \mathbf{assert} \text{ } Inv \}$$

The **havoc** statement erases the state: $slift_{\mathcal{L}_0}(sto, \varphi)(\mathbf{havoc}) := true$. 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 $\varphi := \text{res} \geq 0$, i.e., $\models [p \Vdash_{\alpha_{\text{big}}} \varphi]$. We first compute the symbolic traces by symbolic lifting, starting from an initial store $st_0 := a := a_0 \parallel b := b_0 \parallel \text{res} := \text{res}_0 \parallel \text{tmp} := \text{tmp}_0$:

$$\begin{aligned}
\text{sift}_{\mathcal{L}_0}(st_0, \text{true})(p) = & \\
& (st_0 \circ (\text{res} := 0), \text{true}); \\
& (((st_0 \circ (\text{res} := 0), b_0 < a_0); \\
& \quad (st_0 \circ (\text{res} := 0 \parallel \text{tmp} := a), b_0 < a_0); \\
& \quad (b := b_0 \parallel \text{res} := 0 \parallel \text{tmp} := a_0 \parallel a := b_0, b_0 < a_0); \\
& \quad (\text{res} := 0 \parallel \text{tmp} := a_0 \parallel a := b_0 \parallel b := a_0, b_0 < a_0); \\
& \quad (\text{tmp} := a_0 \parallel a := b_0 \parallel b := a_0 \parallel \text{res} := a_0 - b_0, b_0 < a_0)) \\
& + ((st_0 \circ (\text{res} := 0), b_0 \geq a_0); \\
& \quad (a := a_0 \parallel b := b_0 \parallel \text{tmp} := \text{tmp}_0 \parallel \text{res} := b_0 - a_0, b_0 \geq a_0))
\end{aligned}$$

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

$$\begin{aligned}
\text{sift}_{\mathcal{L}_0}(st_0, \text{true})(p) = & (st_0 \circ (\text{res} := 0), \text{true}); \\
& (((st_0 \circ (\text{res} := 0), b_0 < a_0); \\
& \quad (st_0 \circ (\text{res} := 0 \parallel \text{tmp} := a_0), b_0 < a_0); \\
& \quad (st_0 \circ (\text{res} := 0 \parallel \text{tmp} := a_0 \parallel a := b_0), b_0 < a_0); \\
& \quad (\text{res} := 0 \parallel \text{tmp} := a_0 \parallel a := b_0 \parallel b := a_0, b_0 < a_0)) \\
& + (st_0 \circ (\text{res} := 0), b_0 \geq a_0); \\
& (\text{tmp} := (b_0 < a_0) ? a_0 : \text{tmp}_0 \parallel \\
& \quad a := (b_0 < a_0) ? b_0 : a_0 \parallel \\
& \quad b := (b_0 < a_0) ? a_0 : b_0 \parallel
\end{aligned}$$

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

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

$$(b_0 < a_0 \rightarrow \{\text{tmp} := a_0 \parallel a := b_0 \parallel b := a_0 \parallel \text{res} := a_0 - b_0\} \text{res} \geq 0) \wedge \\ (b_0 \geq a_0 \rightarrow \{\text{res} := b_0 - a_0\} \text{res} \geq 0)$$

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 $\models [p \Vdash_{\alpha_{\text{big}}} \varphi]$ also holds. In this example, we also could have shown the assertion $\Box\varphi$, i.e. that the trace $\text{slift}_{\text{TR}}(st_0, \text{true})(\Box\varphi) = \varphi^*$ subsumes the symbolic trace of p , since φ holds for all intermediate states. \diamond

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 $\models [p \Vdash sc]$. 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 $st_0 = (x := x_0 \parallel v := v_0 \parallel i := i_0)$ be an initial store. The invariant is $\text{Inv}(v, i, x) := v = i^2 \wedge x \geq (i - 1)^2 \wedge i \geq 1$. We obtain the following symbolic traces (to ease the presentation, we omit abstract path conditions for schematic statements):

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

$$\begin{aligned}
& (x := x_1 \parallel v := v_1 + 2i_1 + 1 \parallel i := i_1 + 1, \text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1); \\
& ((\text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \text{Inv}(v_1 + 2i_1 + 1, i_1 + 1, x_1))^! \quad + \\
& ((x := x_1 \parallel v := v_1 \parallel i := i_1, \text{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1)); \\
& (x := x_1 \parallel v := v_1 \parallel i := i_1 - 1, \text{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1))
\end{aligned}$$

$$\begin{aligned}
& \text{sift}_{\mathcal{L}}(st_0, x_0 \geq 1)(sc) = \\
& \text{true}^*; (st_0 \circ \text{sto}_P, x_0 \geq 1); \\
& (x_0 \geq 1 \rightarrow \{st_0 \circ \text{sto}_P\}(v \doteq 1 \wedge i \doteq 1 \wedge x \doteq 1))^!; \\
& (x_0 \geq 1 \rightarrow \{st_0 \circ \text{sto}_P\} \text{Inv}(v, i, x))^!; \\
& (x := x_1 \parallel v := v_1 \parallel i := i_1, x_0 \geq 1); \\
& (x := x_1 \parallel v := v_1 \parallel i := i_1, x_0 \geq 1 \wedge \text{Inv}(v_1, i_1, x_1)); \\
& (((x := x_1 \parallel v := v_1 \parallel i := i_1, \text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1); \\
& \text{true}^*; ((x := x_1 \parallel v := v_1 \parallel i := i_1) \circ \text{sto}_Q, \text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1); \\
& ((\text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \\
& \{ (x := x_1 \parallel v := v_1 \parallel i := i_1) \circ \text{sto}_Q \} \text{Inv}(v, i, x))^!; \\
& ((\text{Inv}(v_1, i_1, x_1) \wedge v_1 \leq x_1) \rightarrow \\
& \{ (x := x_1 \parallel v := v_1 \parallel i := i_1) \circ \text{sto}_Q \} \text{Inv}(v, i, x))^! \quad + \\
& ((x := x_1 \parallel v := v_1 \parallel i := i_1, \text{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1)); \\
& \text{true}^*; \\
& ((x := x_1 \parallel v := v_1 \parallel i := i_1) \circ \text{sto}_R, \text{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1); \\
& (((\text{Inv}(v_1, i_1, x_1) \wedge v_1 > x_1) \rightarrow \\
& \{ (x := x_1 \parallel v := v_1 \parallel i := i_1) \circ \text{sto}_R \} (i^2 \leq x < (i+1)^2))^!))
\end{aligned}$$

We point out that we performed the same anonymization $x := x_1 \parallel v := v_1 \parallel i := i_1$ 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 $(\text{sift}_{\mathcal{L}}(st_0, x_0 \geq 1)(p))$ is subsumed by $\text{sift}_{\mathcal{L}}(st_0, x_0 \geq 1)(sc)$. 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 $(st_0 \circ (v := 1), x_0 \geq 1)$, for instance, is subsumed by true^* , and the state $(x := x_0 \parallel v := 1 \parallel i := 1, x_0 \geq 1)$ by $(st_0 \circ \text{sto}_P, x_0 \geq 1)$, since sto_P is yet uninstantiated and can therefore be instantiated to $v := 1 \parallel i := 1$. Next are the

assertions before the loop: Since based on the instantiation of $stop$ 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. \diamond

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 $res \geq 0$. 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, $(res := -1, true)$ would not be subsumed by $res \geq 0$ and the algorithm would return UNKNOWN. \diamond




Fig. 2: Implementation SFA for Example 9




Fig. 3: Specification SFA for Example 9