Proving Reachability-Logic Formulas Incrementally

Vlad Rusu, Andrei Arusoaie

HAL Id: hal-01282379

https://inria.hal.science/hal-01282379

Submitted on 3 Mar 2016

Proving Reachability-Logic Formulas Incrementally

Vlad Rusu and Andrei Arusoaie
`First.Last@inria.fr`

Inria, Lille, France

Abstract. Reachability Logic (RL) is a formalism for defining the operational semantics of programming languages and for specifying program properties. As a program logic it can be seen as a language-independent alternative to Hoare Logics. Several verification techniques have been proposed for RL, all of which have a circular nature: the RL formula under proof can circularly be used as a hypothesis in the proof of another RL formula, or even in its own proof. This feature is essential for dealing with possibly unbounded repetitive behaviour (e.g., program loops). The downside of such approaches is that the verification of a set of RL formulas is monolithic, i.e., either all formulas in the set are proved valid, or nothing can be inferred about any of the formula's validity or invalidity. In this paper we propose a new, incremental method for proving a large class of RL formulas. The proposed method takes as input a given RL formula under proof (corresponding to a given program fragment), together with a (possibly empty) set of other valid RL formulas (e.g., already proved using our method), which specify sub-programs of the program fragment under verification. It then checks certain conditions are shown to be equivalent to the validity of the RL formula under proof. A newly proved formula can then be incrementally used in the proof of other RL formulas, corresponding to larger program fragments. The process is repeated until the whole program is proved. We illustrate our approach by verifying the nontrivial Knuth-Morris-Pratt string-matching program.

1 Introduction

Reachability Logic (RL) [1,2,3,4] is a language-independent logic for defining the operational semantics of programming languages and for specifying properties of programs. For instance, on the `sum` program in Fig. 1, the RL formula

$$\langle \mathtt{sum}, \mathtt{n} \mapsto a \,\rangle \wedge a \geq 0 \Rightarrow (\exists i, s)\langle \mathtt{skip}, \mathtt{n} \mapsto a\, \mathtt{i} \mapsto i\, \mathtt{s} \mapsto s \rangle\rangle \wedge s = sum(a) \qquad (1)$$

specifies that after the complete execution of the `sum` program from a configuration where the program variable `n` is bound to a non-negative value a, a configuration where `s` is bound to a value $s = sum(a)$ is reached. Here, $sum(a)$ is a mathematical definition of the sum of natural numbers up to a.

Existing RL verification tools [1,2,4,5,6] would typically verify formula (1) as follows. First, they would consider (1) together with, e.g., the following formula (2), where `while` denotes the program fragment consisting of the while-loop

```
i := 1;
s := 0;
while (i <= n) do
  s := s + i;
  i := i + 1
end
```

Fig. 1. Program sum.

in Figure 1. The formula (2) is intended to specify the while loop, just like (1) specifies the whole program, and can be seen as encoding a loop invariant.

$$\langle \texttt{while}, \texttt{n} \mapsto a \ \texttt{i} \mapsto i \ \texttt{s} \mapsto s \rangle \wedge 0 < i \leq a + 1 \wedge s = sum(i-1) \tag{2}$$
$$\Rightarrow (\exists i', s') \langle \texttt{skip}, \texttt{n} \mapsto a \ \texttt{i} \mapsto i' \ \texttt{s} \mapsto s' \rangle \wedge s' = sum(a)$$

Then, the tool would symbolically execute at least one instruction in the programs in the left-hand side of both (1) and (2) using the semantics of the instructions of the language (assumed to be also expressed as RL formulas[1]), and then execute the remaining programs in the left-hand sides of the resulting formulas *as if both (1) and (2) became new semantical rules of the language*. For example, when the program executed in (1) reaches the while loop, the rule (2) can be applied instead of the rule defining the semantics of the while instruction - that is, when proving (1), (2) is assumed to hold. Similarly, when the program in (2) completes one loop iteration, the left-hand side of (2) contains again the same while loop as initially, with other values mapped to the variables. Then, (2) is applied instead of the rule defining the semantics of the while instruction. Thus, it is assumed that (2) holds after having completed one loop iteration.

The *circular* reasoning illustrated in the above example is sound, in the sense that if such a proof succeeds, all the formulas under proof are (semantically) valid. However, if the proof does not succeed, nothing can be said about the validity of the formulas. In our example, (1) or (2) (or both) could be invalid.

Contribution. In this paper we propose a new method for proving a significant subset of RL formulas, which, unlike existing verification methods, is *incremental*. In our example, the proposed method would first prove (2), and then would prove (1) knowing for a fact (i.e., not assuming) that (2) is valid. Thus, if the proof of (1) fails for some reason, the user still knows that (2) holds, and can take action for fixing the proof based on this knowledge. Of course, for a simple program such as the above example the advantage of incremental RL verification is not obvious, but it turns out to make quite a difference when veriying more challenging programs, such as the KMP program illustrated later in the paper.

We first establish an equivalence between the validity of RL formulas and two technical conditions (one condition is an invariance property, and the other one regards the so-called capturing of terminal configurations). Then we propose a

---

[1] for the language of interest in this paper the rules are shown in Section 2.

graph-construction approach that takes a given RL formula under proof (corresponding to a given program fragment), together with a (possibly empty) set of other valid RL formulas (e.g., proved using a previous iteration of our approach, or by any other sound RL formula verification method). The latter formulas specify sub-programs of the program fragment currently under verification. The invariance and terminal-configuration capturing conditions are then checked on the graph, thus establishing the validity of the RL formula under proof. The newly proved formula can then be incrementally used in the proof of other RL formulas, corresponding to larger program fragments. The same process is then repeated until, eventually, the whole program is proved.

Of course, the proposed method has limitations, since verification of RL formulas is in general undecidable. The graph construction may not terminate, or the conditions to be checked on it may not hold. One situation that a purely incremental method cannot handle is mutually recursive function calls, in which none of the functions can be verified individually unless (coinductively) assuming that the other function's specifications hold. A natural solution here is to use an incremental method as much as possible, and to locally apply a circular approach only for subsets of formulas that the incremental method cannot handle.

In order to demonstrate the feasibility of our approach we illustrate it on the nontrivial Knuth-Morris-Pratt KMP string-matching program. The program is written in a simple imperative language, whose syntax and semantics is defined in Maude [7]. We chose Maude in order to benefit from its reflective capabilities, which turned out to be very useful for implementation purposes. We are using a specific version of Maude that has been interfaced with the Z3 solver [8], which is here used for simplifying conditions required for proving RL formula validity.

Paper organisation. After this introduction we present in Section 2 the Maude-based definition of a simple imperative programming language IMP+ that includes assignments, conditions, loops, and simple procedures operating on global variables. In Section 3 we present background notions: Reachability Logic, and how the language definition from the previous section fits in this framework (Section 3.1); and language-parametric symbolic execution, together with its implementation by rewriting based on transforming the semantical rules of a language (Section 3.2). In Section 4 we present the incremental RL-formula verification method. In Section 5 we illustrated our method on the KMP string-matching algorithm, and in Section 6 we conclude and present related and future work.

2 Defining a Simple Programming Language

In this section we define the language IMP+ in Maude. IMP+ is simple enough so that its Maude code is reasonably small (less than two hundred lines of code), yet expressive enough for programming algorithms on arrays such as the KMP. We assume Maude is familiar to readers; for details the standard reference is [7].

Datatypes. IMP+ computes over Booleans, integers, and integer arrays. We use the builtin Booleans and integers of Maude, and provide a standard algebraic definition of arrays. The constructor `array : Nat -> IntArray` creates an array of a given length. The operation `store : IntArray Nat Int -> IntArray` stores a given integer (third argument) at a given natural-number index. An operation `select : IntArray Nat -> Int` returns the element at the position given by the second argument. These functions are defined equationally. They return error values in case of attempts to access indices out of an array's bounds.

Syntax. The syntax on IMP+ consists of expressions (arithmetic and Boolean) and statements. Each of these syntactical categories is defined by a sort, i.e., `AExp`, `BExp`, and `Stmt`. Allowed arithmetical operations are addition, substraction, and array selector, denoted by `_++_`, `_--_`, and `_[_]` respectively, in order to avoid confusion with the corresponding Maude operations on the datatypes. In the same spirit, Boolean operations are less-or-equal-than (`_<==_`) and equality (`_===_`); negation `!`; and conjunction `_&&_`. Such expressions are built from identifiers (i.e., program variables) and constants (Maude integers and Booleans).

The statements of IMP+ are: assignments to integer variables and array elements (`_:=_`); conditional (`if_then_else_endif`); while loops (`while_do_end`); parameterless function declaration (`function_(){_}`) and call (`_()`); a `print` instruction; and finally, a sequencing `_;_` instruction that, for convenience, is declared associative with the "do-nothing" `skip` instruction as a neutral element.

Semantics. Semantical rules operate on *configurations*, which consist of a program to be executed, a mapping of integer variables to values and of function names to statements, and a list of integers denoting the output of the program. In Maude we write a constructor `<_,_,_,_> : Stmt Map Funs Ints -> Cfg`. Getters and setters for the `Map` and `Funs` maps are also equationally defined.

The semantics of IMP+ then consists in evaluating expressions (in a given map, assigning values to variables) and statements (in a given configuration, describing all the infrastructure required for statements to execute). Expressions are evaluated using equations, and statements are evaluated using rewrite rules.

Evaluating expressions. This amounts to writing a function `eval` and equations:

```
op eval : AExp Map  -> Int .
eq eval(I, M) = I .
eq eval(X, (M (X -> J))) = J .
eq eval(X[E], (M (X -> A))) = select(A,eval(E,(M (X -> A)))) .
...
op eval : BExp Map -> Bool .
eq eval(B,M) = B .
eq eval(Cnd1 && Cnd2, M) = eval(Cnd1,M) and eval(Cnd2,M) .
...
```

That is, `eval` goes through the structure of an expression and evaluates it in a given mapping of values to variables. Here, e.g., `M (X -> J)` denotes an associative-commutative map, constructed as the anonymous juxtaposition operation `__` of a map variable `M` with a map of the identifier `X` to the integer `J`.

Evaluating statements. This is performed by rewrite rules, some of which are:

```
rl [assign]: <((X := E) ; S), M, F, O > => < S, set(X, eval(E, M), M), F, O > .

crl [if-true]: <(if Cnd then S1 else S2 endif) ; S, M, F, O > => < S1 ; S, M, F, O >
               if  eval(Cnd,M) .

crl [if-false]: <(if Cnd then S1 else S2 endif) ; S, M, F, O > => < S2 ; S, M, F, O >
                if  not eval(Cnd,M) .

rl [while]: <(while Cnd do S1 end) ; S, M, F, O > =>
               <(if Cnd then S1 ; while Cnd do S1 end else skip endif) ; S, M, F, O > .

rl [print]: < (print E) ; S, M, F, O >   =>   < S, M, F, (O ; eval(E,M)) > .
```

The first rule deals with assigment to a program-variable `X` of an arithmetic expression `E`. It uses the `set` function on maps in order to update the map so that `X` is mapped to the value of `E`. Another rule, not shown here, deals with assignments to array elements. The following two rules describe the two possible outcomes of a conditional instruction, depending on the value of the condition. The rule for the while loop consists essentially in loop unrolling. The rule for the printing instruction appends the value of the instruction's argument to the list of integers (last argument of configurations) denoting the program's output.

3   Reachability Logic and Symbolic Execution

In this section we present background material used in the rest of the paper. We illustrate the concepts with examples from the IMP+ language.

3.1   Reachability Logic

Several versions of RL have been proposed in the last few years [1,2,3,4]. Moreover, RL is built on top of *Matching Logic* (ML), which also exists in several versions [9,10,11]. (The situation is somewhat similar to the relationship between rewriting logic and the equational logics underneath it.) We adopt the recent *all-paths* interpretation of RL [4], built upon a minimal ML that is enough to express typical practically-relevant properties about program configurations and is amenable to *symbolic execution* by rewriting, a key ingredient of our method.

The formulas of ML that we consider are called *patterns* and are defined as follows. Assume an algebraic signature Σ with a set S of sorts, including two distinguised sorts $Bool, Cfg \in S$. We write $T_{\Sigma,s}(Var)$ for the set of terms of sort s over a set Var of S-indexed variables and $T_{\Sigma,s}$ for the set of ground terms of sort s. We identify the $Bool$-sorted operations in Σ with a set Π of predicates.

Example 1. Consider the Maude definition of the IMP+ language. Then, Σ is the algebraic signature containing all the sorts and operations described in the previous section, including the `Bool` and `Cfg` sorts. The operation `eval : BExp Map -> Bool` has sort `Bool` and is thus identified with a predicate in the set Π. The sort `Cfg` has the constructor `<_,_,_,_> : Stmt Map Funs Ints -> Cfg`.

Definition 1 (Pattern). *A pattern is an expression of the form $(\exists X)\pi \wedge \phi$, where $X \subset Var$, $\pi \in T_{\Sigma,Cfg}(X)$ and ϕ is a FOL formula over the FOL signature (Σ, Π) with free variables in X.*

We often denote patterns by φ and write $\varphi \triangleq (\exists X)\pi \wedge \phi$ to emphasise its components: the quantified variables X, the *basic pattern* π, and ϕ, the *condition*. We let *FreeVars*(φ) denote the set of variables freely occurring in a pattern φ, defined as usual (i.e., not under the incidence of a quantifier). We often identify basic patterns π with $(\exists\emptyset)\pi \wedge true$, and *elementary patterns* $\pi \wedge \phi$ with $(\exists\emptyset)\pi \wedge \phi$.

Example 2. The left and right-hand sides of the rules defining the semantics of IMP+ are basic patterns, `< S, M, F, O > /\ eval(true,M)` is an elementary pattern, and $(\exists O)$ `< S, M, F, O > /\ eval(true,M)` is a pattern.

We now describe the semantics of patterns. We assume a model M of the algebraic signature Σ. In the case of the Maude specification of IMP+ the model M, M is the initial model induced by the specification's equations and axioms. For sorts $s \in S$ we write M_s for the interpretation (a.k.a. carrier set) of the sort s. We call *valuations* the functions $\rho : Var \to M$ that assign to variables in Var a value in M of a corresponding sort, and *configurations* the elements in M_{Cfg}.

Definition 2 (Pattern semantics). *Given a pattern $\varphi \triangleq (\exists X)\pi \wedge \phi$, $\gamma \in M_{Cfg}$ a configuration, and $\rho : Var \to M$ a valuation, the satisfaction relation $(\gamma, \rho) \models \varphi$ holds iff there exists a valuation ρ' with $\rho'|_{Var \setminus X} = \rho|_{Var \setminus X}$ such that $\gamma = \rho'(\pi)$ and $\rho' \models \phi$ (where the latter $\models$ denotes satisfaction in FOL, and $\rho_{|Var \setminus X}$ denotes the restriction of the valuation ρ to the set $Var \setminus X$).*

We let $[\![\varphi]\!]$ denote the set $\{\gamma \in M_{Cfg} \mid (\exists\rho : Var \to M)(\gamma, \rho) \models \varphi\}$. A formula φ is *valid in M*, denoted by $M \models \varphi$, if it is satisfied by all pairs (γ, ρ).

We now recall Reachability-Logic (RL) formulas, the transition systems that they induce, and their all-paths semantics [4] that we will be using in this paper.

Definition 3 (RL Formulas). *An RL formula is a pair of patterns $\varphi \Rightarrow \varphi'$.*

Examples of RL formulas were given in the introduction. The rules defining the semantics of IMP+ are also RL formulas (for the conditional rules, just assume that the expression following `if` is the condition of the rule's left-hand side).

Let $\mathcal{S}$ denote a fixed set of RL formulas, e.g., the semantics of a given language. We define the transition system defined by $\mathcal{S}$ together with some notions related to this transition system, and then the notion of validity for RL formulas.

Definition 4 (Transition System defined by $\mathcal{S}$). *The transition system defined by $\mathcal{S}$ is $(M_{Cfg}, \Rightarrow_\mathcal{S})$, where $\Rightarrow_\mathcal{S} = \{(\gamma, \gamma') \mid (\exists\varphi \Rightarrow \varphi' \in \mathcal{S})(\exists\rho)(\gamma, \rho) \models \varphi \wedge (\gamma', \rho) \models \varphi'\}$. We write $\gamma \Rightarrow_\mathcal{S} \gamma'$ for $(\gamma, \gamma') \in \Rightarrow_\mathcal{S}$. A state γ is terminal if there is no γ' such that $\gamma \Rightarrow_\mathcal{S} \gamma'$. A path is a sequence $\gamma_0 \cdots \gamma_n$ such that $\gamma_i \Rightarrow_\mathcal{S} \gamma_{i+1}$ for all $0 \le i \le n-1$. Such a path is complete if γ_n is terminal.*

An RL formula $\varphi \Rightarrow \varphi'$ is valid, written $\mathcal{S} \models \varphi \Rightarrow \varphi'$, if for all pairs (γ_0, ρ) such that $(\gamma_0, \rho) \models \varphi$, and all complete paths $\gamma_0 \Rightarrow_\mathcal{S} \cdots \Rightarrow_\mathcal{S} \gamma_n$, there exists $0 \le i \le n$ such that $(\gamma_i, \rho) \models \varphi'$.

Note that the validity of RL formulas is only determined by finite, complete paths. Infinite paths, induced by nonterminating programs, are not considered. Thus, termination is assumed: as a program logic, RL is a logic of partial correctness. We restrict our attention to RL formulas satisfying the following assumption:

Assumption 1 *RL formulas have the form $\pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r$ and satisfy FreeVars$(\pi_r) \subseteq$ FreeVars$(\pi_l) \cup Y$, FreeVars$(\phi_r) \subseteq$ FreeVars$(\pi_l) \cup$ FreeVars(π_r), and FreeVars$(\phi_l) \subseteq$ FreeVars(π_l).*

That is, the left-hand side is an elementary pattern, and the right hand side is a pattern, possibly with quantifiers. Such formulas are typically expressive enough for expressing language semantics (for this purpose, quantifiers are not even required)[2] and program properties. For program properties, existentially quantified variables in the right-hand side are useful to denote values computed by a given program, which are not known before the program computes them, such as s - the sum of natural numbers up to a given bound - in the formula (1).

3.2 Language-Parametric Symbolic Execution

We now briefly present symbolic execution, a well-known program analysis technique that consists in executing programs with symbolic input (e.g. a symbolic value x) instead of concrete input (e.g. 0). We reformulate the language-independent symbolic execution approach we already presented elsewhere [6], with some simplifications (e.g., unlike [6] we do not use coinduction). The approach consists in transforming the signature Σ and semantics $\mathcal{S}$ of a programming language so that, under reasonable restrictions, executing a program with the modified semantics amounts to executing the program symbolically.

Consider the signature Σ corresponding to a language definition. Let *Fol* be a new sort whose terms are all FOL formulas, including existential and universal quantifiers. Let *Id* and *IdSet* be new sorts denoting identifiers and sets of identifiers, with a union operation $\_,\_$. Let *Cfg*$^{\mathfrak{s}}$ be a new sort, with constructor $(\exists\_)\_\wedge\_ : IdSet \times Cfg \times Fol \to Cfg^{\mathfrak{s}}$. Thus, patterns $(\exists X)\pi \wedge \phi$ correspond to terms $(\exists X)\pi \wedge \phi$ of sort *Cfg*$^{\mathfrak{s}}$ in the enriched signature and reciprocally. Consider also the following set of RL formulas, called the *symbolic version of $\mathcal{S}$*:

$$\mathcal{S}^{\mathfrak{s}} \triangleq \{(\exists\mathcal{X})\pi_l \wedge \psi \Rightarrow (\exists\mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r)|\pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r \in \mathcal{S}\}$$

with ψ a new variable of sort *Fol*, and $\mathcal{X}$ a new variable of sort *IdSet*.

Example 3. The following conditional rule is part of the semantics $\mathcal{S}$ of IMP+:
`<if C then S1 else S2 endif ; S, M, F, O> => <S1 ; S, M, F, O>` if `eval(C,M)`
Written as an RL formula (with patterns in left and right-hand sides) it becomes[3]
`<if C then S1 else S2 endif ; S, M, F, O>` $\wedge$ `eval(C,M)=> <S1 ; S, M, F, O>`
The corresponding rule in $\mathcal{S}^{\mathfrak{s}}$ becomes an unconditional rule: $(\exists\mathcal{X})$ `<if C then S1 else S2 endif ; S, M, F, O>` $\wedge\psi$ `=>` $(\exists\mathcal{X})$ `<S1 ; S, M, F, O>` $\wedge$ $(\psi\wedge$ `eval(C,M))`.

---

[2] see, e.g., the languages defined in the $\mathbb{K}$ framework: http://k-framework.org.
[3] We liberally use a mixture of Maude and math notation for the sake of the example.

The interest of the above nontrivial construction is that, under reasonable assumptions, stated below, rewriting with the rules in $\mathcal{S}^\mathfrak{s}$ achieves a *simulation* of rewriting with the rules in $\mathcal{S}$, which is a result that we need for our approach.

Assumption 2 *There exists a* builtin subsignature $\Sigma^b \subsetneq \Sigma$. *The sorts and operations in* Σ^b *are* builtin, *while all others are* non-builtin. *The sort Cfg is not builtin. Non-builtin operation symbols may only be subject to a (possibly empty) set of* linear, regular, and non-collapsing *axioms.*

We recall that an axiom $u = v$ is linear if both u, v are linear (a term is linear if any variable occurs in it at most once); it is regular if both u, v have the same set of variables; and it is non-collapsing if both u, v have non-builtin sorts.

Example 4. For the IMP+ language specification we assume that the non-builtin sorts are Cfg, Stmt (for statements), and Funs (which map function identifiers to statements). Statements were declared to be associative with unity, whereas maps of identifiers to statements were taken to be associative and commutative with unity. All these axioms have the properties requested by Assumption 2.

In order to formulate the simulation result we now define the transition relation generated by the set of symbolic RL rules $\mathcal{S}^\mathfrak{s}$. It is essentially rewriting modulo the congruence $\cong$ on $T_\Sigma(Var)$ induced by the axioms in Assumption 2. Let $Var^b \subset Var$ be the set of variables of builtin sorts. We first need the following technical assumption, which does not restrict the generality of our approach:

Assumption 3 *For every* $\pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r \in \mathcal{S}$, $\pi_l \in T_{\Sigma \setminus \Sigma^b}(Var)$, π_l *is linear, and* $Y \subseteq Var^b$.

The assumption can always be made to hold by replacing in π_l all non-variable terms in Σ^b and all duplicated variables by fresh variables, and by equating in the condition ϕ_l the new variables to the terms that they replaced.

For the sake of complying with the definition of rewriting we need to extend the congruence $\cong$ to terms of sort $Cfg^\mathfrak{s}$ by $(\exists X)\pi_1 \wedge \phi \cong (\exists X)\pi_2 \wedge \phi$ iff $\pi_1 \cong \pi_2$.

Definition 5 (Relation $\Rightarrow_{\alpha^\mathfrak{s}}$). *For* $\alpha^\mathfrak{s} \triangleq (\exists \mathcal{X})\pi_l \wedge \psi \Rightarrow (\exists \mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r)$ $\in \mathcal{S}^\mathfrak{s}$ *we write* $(\exists X)\pi \wedge \phi \Rightarrow_{\alpha^\mathfrak{s}} (\exists X, Y)\pi' \wedge \phi'$ *whenever* $(\exists X)\pi \wedge \phi \ \alpha^\mathfrak{s}$ *is rewritten by* $\alpha^\mathfrak{s}$ *to* $(\exists X, Y)\pi' \wedge \phi'$, *i.e., there exists a substitution* σ' *on* $Var \cup \{\mathcal{X}, \psi\}$ *such that* $\sigma'((\exists \mathcal{X})\pi_l \wedge \psi) \cong (\exists X)\pi \wedge \phi$ *and* $\sigma'((\exists \mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r)) = (\exists X, Y)\pi' \wedge \phi'$.

Lemma 1 ($\Rightarrow_{\alpha^\mathfrak{s}}$ simulates $\Rightarrow_\alpha$). *For all* $\gamma, \gamma' \in M_{Cfg}$, *all patterns* φ *with* $FreeVars(\varphi) \subseteq Var^b$, *and all valuations* ρ, *if* $(\gamma, \rho) \models \varphi$ *and* $\gamma \Rightarrow_\alpha \gamma'$ *then there exists* φ' *with* $FreeVars(\varphi') \subseteq Var^b$ *such that* $\varphi \Rightarrow_{\alpha^\mathfrak{s}} \varphi'$ *and* $(\gamma', \rho) \models \varphi'$.

As a consequence, any concrete execution (following $\Rightarrow_\mathcal{S}$) such that the initial configuration satisfies a given initial pattern φ is simulated by a symbolic execution (following $\Rightarrow_{\mathcal{S}^\mathfrak{s}}$) starting in φ. We shall also use the following notion of *derivative*, which collects all the symbolic successors of a pattern by a rule:

Definition 6 (Derivatives). $\Delta_\alpha(\varphi) = \{\varphi' \mid \varphi \Rightarrow_{\alpha^\mathfrak{s}} \varphi'\}$ *for any* $\alpha \in \mathcal{S}$.

Since the symbolic successors are computed by rewriting, the derivative operation is computable and always returns a finite set of patterns.

4 Proving RL Formulas Incrementally

In this section we present an incremental method for proving RL formulas. We first state two technical conditions and prove that they are equivalent to RL formula validity. The equivalence works for so-called *terminal* formulas, whose right-hand side specifies a completed program; however, a generalisation to non-terminal formulas, required for incremental verification, is also given. Thus, RL formula verification amounts to checking the two above-mentioned conditions.

For this, we present a graph construction based on symbolic execution that, if it terminates successfully, ensures that the two conditions in question hold for a given RL formula. The graph construction is parameterised by a set of formulas that have already been proved valid (using the same method, or any other sound one). These formulas correspond to subprograms of the given program fragment that the current formula under proof specifies. The current formula, once proved, can then be used in proofs of formulas specifying larger program fragments. We consider a fixed set $\mathcal{S}$ or RL formulas and their transition relation $\Rightarrow_{\mathcal{S}}$. The first of the two following definitions says that all terminal configurations reachable from a given pattern "end up" as instances of a quantified basic pattern:

Definition 7 (Capturing All Terminal Configurations). *We say that a pattern $(\exists Y)\pi'$ captures all terminal configurations for a pattern φ if for all (γ, ρ) such that $(\gamma, \rho) \models \varphi$, and all complete paths $\gamma \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma'$, $(\gamma', \rho) \models (\exists Y)\pi'$.*

The second definition characterises FOL formulas that hold in a given quantified pattern, i.e., conditions satisfied by all configurations reachable from a given initial pattern whenever they "reach" the quantified pattern in question:

Definition 8 (Invariant at, Starting from). *We say that a FOL formula $(\exists Y)\phi'$ is invariant at a pattern $(\exists Y)\pi'$ starting from a pattern φ if for all (γ, ρ) such that $(\gamma, \rho) \models \varphi$, all paths $\gamma \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma'$, and all valuations ρ' with $\rho'|_{Var \setminus Y} = \rho|_{Var \setminus Y}$, if $\gamma' = \rho'(\pi')$, then $\rho' \models \phi'$.*

Note that the *same* values of the variables Y were used for satisfying π' and ϕ'.

Definition 9. *A basic pattern π' is terminal if for all valuations ρ, $\rho(\pi')$ is a terminal configuration. A rule $\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$ is terminal if π' is terminal.*

The following proposition characterises the validity of terminal RL formulas:

Proposition 1 (Equivalent Conditions for Terminal Formula Validity). *Consider a terminal formula $\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$. Then $\mathcal{S} \models \pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$ iff*

1. *$(\exists Y)\phi'$ is invariant at $(\exists Y)\pi'$ starting from $\pi \wedge \phi$, and*
2. *$(\exists Y)\pi'$ captures all terminal configurations for $\pi \wedge \phi$.*

Remark 1. The ($\Leftarrow$) implication in Proposition 1 is the important one for the soundness of our method. Its proof naturally follows from definitions. For the reverse implication, the following assumption is required: for all right-hand sides $\varphi_r \triangleq (\exists Y)\pi_r \wedge \phi_r$ of rules in $\mathcal{S}$, if $\rho(\pi_r) = \rho'(\pi_r)$ then $\rho|_{FreeVars(\pi_r)} = \rho'|_{FreeVars(\pi_r)}$.

The assumption does not restrict generality as it can always be made to hold, by replacing subterms of patterns by fresh variables (and adding equations to the condition) and by noting that the *Cfg* sort is interpreted syntactically in the model M. Then, $\pi_r \triangleq f(x_1, \ldots, x_n)$ where f is the constructor for the *Cfg* sort, and $\rho(f(x_1, \ldots, x_n)) = \rho'(f(x_1, \ldots, x_n))$ iff $\rho(x) = \rho'(x_i)$ for all variables x_i.

Remark 2. Proposition 1 works for terminal RL formulas. We shall need the following observation: assume that an RL formula of the following form $\langle P \ldots \rangle \wedge \phi \Rightarrow (\exists Y)\langle skip \ldots \rangle \wedge \phi'$ has been proved valid, where P is a program, *skip* denotes the empty program, and suspension dots denote the rest of the configurations (which depend on the programming language). Then, assuming a sequencing operation[4] denoted by semicolon, the following formula $\langle P; Q \ldots \rangle \wedge \phi \Rightarrow (\exists Y)\langle Q \ldots \rangle \wedge \phi'$ is also valid: if each terminal path executing P ended up in the empty program, then each path executing $P; Q$ still has Q to execute after having executed P. As shown later in this section, the validity of such "generalized" formulas enables us to incrementally use a proved-valid formula in the proofs of other formulas.

Proposition 1 is the basis for proving RL formulas, by checking the conditions (1) and (2). We now show how the conditions can be checked mechanically.

Symbolic graph construction. The graph-construction procedure in Figure 2 uses symbolic execution and is used to check the conditions (1),(2) in Prop. 1. Before we describe the procedure we introduce the components that it uses.

A partial order $<$ on $\mathcal{S}$. The procedure assumes a set of RL formulas $\mathcal{S}$, which consist of the semantical rules $\mathcal{S}_0$ of a programming language and a (possibly empty) set of RL formulas $\mathcal{G}$ that were already proved valid in an earlier step of our envisaged incremental verification method. Such formulas, sometimes called *circularities* in RL verification, specify subprograms of the program under verification, and are assumed here to have the form $\langle P; Q, \ldots \rangle \wedge \phi \Rightarrow (\exists Y)\langle Q, \ldots \rangle \wedge \phi'$ (cf. Remark 2). During symbolic execution, circularities can be symbolically applied "in competition with" rules in the semantics (e.g., when the program to be executed is $P; Q$, the symbolic version of the above rule can be applied, but the symbolic version of the semantical rule for the first instruction of P can be applied as well). We solve the conflict between semantical rules and circularities by giving priority to the latter.

We use the following notations. Let $lhs(\alpha)$ denote the left-hand side of a formula α. Let $\mathcal{G} < \mathcal{S}_0$ denote the fact that for every $g \in \mathcal{G}$ and $\alpha \in \mathcal{S}_0$, $g < \alpha$. Let $\mathcal{S}_0 \models \mathcal{G}$ denote $\mathcal{S}_0 \models g$, for all $g \in \mathcal{G}$, and $min(<)$ denote the minimal elements of $<$.

Assumption 4 *We assume a partial order relation $<$ on $\mathcal{S} \triangleq \mathcal{S}_0 \cup \mathcal{G}$ satisfying: $\mathcal{G} < \mathcal{S}_0$, $\mathcal{S}_0 \models \mathcal{G}$, and for all $\alpha' \in \mathcal{S}$ and pairs (γ, ρ), if $(\gamma, \rho) \models lhs(\alpha')$ then there exists a rule $\alpha \in min(<)$ such that $(\gamma, \rho) \models lhs(\alpha)$.*

---

[4] "Sequencing" and "empty" do not need to be actual statements of the programming language; they can just be artifacts required by the language's operational semantics.

0: $G = (N \triangleq \{\pi \wedge \phi\}, E \triangleq \emptyset)$, *Failure* $\leftarrow$ *false*, *New* $\leftarrow N$
1: **while** not *Failure* and *New* $\neq \emptyset$

  2: **choose** $\varphi \triangleq (\exists X_n)\pi_n \wedge \phi_n \in New$; *New* $\leftarrow New \setminus \{\varphi\}$
  3: **if** $match_{\cong}(\pi_n, \pi') = \emptyset$ **then**
      4: **if** $\bigvee_{\alpha \in min(<)} inclusion(\varphi, lhs(\alpha)) = true$ **then**
        5: **forall** $\alpha \in min(<)$, **forall** $\varphi' \in \Delta_\alpha(\varphi)$
          6: **if** $inclusion(\varphi', \varphi)$ **then** $E \leftarrow E \cup \{\varphi \xrightarrow{\alpha} (\pi \wedge \phi)\}$
          7: **else** *New* $\leftarrow New \cup \{\varphi'\}$; $E \leftarrow E \cup \{\varphi \xrightarrow{\alpha} \varphi'\}$ **endif**
        8: $N \leftarrow N \cup New$
      9: **else** *Failure* $\leftarrow$ *true* **endif**
10: **elseif** not $inclusion(\varphi, (\exists Y)\pi' \wedge \phi')$ **then** *Failure* $\leftarrow$ *true* **endif**.

Fig. 2. Graph construction. $match_{\cong}()$ is matching modulo the non-bultin axioms (cf. Section 3.2), and $inclusion()$ is the object of Definition 10.

This assumption is satisfied by taking as minimal elements of $<$ previously proved circularities, which gives them priority over rules in the semantics that can be applied in competition with them. The other rules in the semantics, which are not in competition with circularities, are not related by $<$ with other formulas and are thus minimal by definition (and valid, since $\alpha \in \mathcal{S}$ implies $\mathcal{S} \models \alpha$).

Inclusion between patterns. The graph-construction procedure uses a test of inclusion between patterns, which satisfies the following definition.

Definition 10 (Inclusion). *An inclusion test is a function that, given patterns* φ, φ', *returns true if for all pairs* (γ, ρ), *if* $(\gamma, \rho) \models \varphi$ *then* $(\gamma, \rho) \models \varphi'$.

The graph construction. We are now ready to present the procedure in Fig. 2. The procedure takes as input an RL formula $\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$ and a set $\mathcal{S}$ of RL formula with an order $<$ on $\mathcal{S}$ as discussed earlier in this section. It builds a graph (N, E) with N the set of nodes (initially, $\{\pi \wedge \phi\}$) and E the set of edges (initially empty). It uses two variables to control a while loop: a Boolean variable *Failure* (initially *false*) and a set of nodes *New* (initially equal to $N = \{\pi \wedge \phi\}$).

At each iteration of the while loop, a node $\varphi_n \triangleq (\exists X_n)\pi_n \wedge \phi_n$ is taken out from *New* (line 2) and checks whether there is a matcher modulo $\cong$ (cf. Section 3.2) of π' onto π_n (line 3). If this is the case, then π_n is an instance of the (terminal) basic pattern π', and the procedure goes to line 10 to check whether φ_n "as a whole" is included in $(\exists Y)\pi' \wedge \phi'$. If this is not the case, then, informally, this indicates a terminal path that does not satisfy the right-hand side of the formula under proof, i.e., of the fact that $(\exists Y)\phi'$ is not invariant at $(\exists Y)\pi'$, in contradiction with the first hypothesis of Prop. 1 that the procedure is checking; *Failure* is reported, which terminates the execution of the procedure. However, if the test at line 3 indicated that π_n is not an instance of the terminal pattern π', then another inclusion test is performed (line 4): whether there exists a minimal rule in $\mathcal{S}$ (i.e., a rule in the language's semantics, or a circularity already proved, as discussed earlier in this section) whose left-hand side includes φ_n. If this is not the case then, informally, this indicates a terminal configuration that is not an

instance of $(\exists Y)\pi'$, which contradicts the second hypothesis of Prop. 1, making the procedure terminate again with *Failure = true*.

If, however, the inclusion test at line 4 succeeds then all symbolic successors φ'_n of φ_n by minimal rules α w.r.t. $<$ are computed. Each of these patterns is tested for inclusion in the initial node $\pi\wedge\phi$. If inclusion holds then an edge is added from φ_n to the initial node, labelled by the rule that generated the symbolic successor in question. Otherwise, a new node φ'_n is created, and an edge from the current node φ_n to the new node, labelled by the rule that generated it, is created, and the while loop proceeds to the next iteration.

The graph-construction procedure does not terminate in general, since the verification of RL formulas is undecidable. However, if it does terminate with *Failure = false* then the two conditions equivalent to the validity of the procedure's input $\pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ hold, i.e., $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$, which is the desired conclusion. This is established by the results in the rest of this section.

The paths in the constructed graph simulate concrete execution paths whose transitions are given by rules from $\mathcal{S}_0$. This is formalised and used in the proof of the main theorem states that the hypotheses of Proposition 1, equivalent to RL formula validity, are checked by the graph-construction procedure.

Theorem 1. *If the procedure in Fig. 2 terminates with Failure = false on a terminal RL formula $\pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$, then $(\exists Y)\phi'$ is invariant at $(\exists Y)\pi'$ starting from $\pi\wedge\phi$, and $(\exists Y)\pi'$ captures all terminal configurations starting from $\pi\wedge\phi$.*

Theorem 1 uses the following (and last) assumptions on RL formulas:

Assumption 5 *All rules $\varphi_l \Rightarrow \varphi_r \in \mathcal{S}$ have the following properties:*

1. for all pairs (γ, ρ) such that $(\gamma, \rho) \models \varphi_l$ there exists γ' such that $(\gamma', \rho) \models \varphi_r$[5].
2. $[\![\varphi_l]\!] \cap [\![\varphi_r]\!] = \emptyset$

The first of the above assumptions says that if the left-hand side of a rule matches a configuration then there is nothing in the right-hand side preventing the application. This property is called *weak well-definedness* in [4] and is shown there to be a necessary condition for obtaining a sound proof system for RL. The second condition just says that the left and right-hand sides of rules cannot share instances - such rules could generate self-loops on instances, which are useless. We then obtain as a corollary the soundness of our RL formula proof method:

Corollary 1 (Soundness). *If procedure in Figure 2 terminates with Failure = false on a terminal RL formula $\pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ then $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$.*

Incremental verification. We are now ready to describe our incremental RL formula verification method. The method works in a setting where each formula has an associated *code* that it specifies, and that for a given RL formula f, $code(f)$ returns the given code. Considering the RL formulas (1) and (2) in Section 1, $code(1)$ is the `sum` program in Figure 1 and $code(2)$ is the `while` subprogram.

---

[5] this property is called weak well-definedness in [4].

The problem to be solved is: given two sets of formulas: $\mathcal{S}$ (the semantics of a language) and $\mathcal{G}$ (the specification of a given program and of some of its subprograms) prove *for all $g \in \mathcal{G}$, $\mathcal{S} \models g$* (for short, $\mathcal{S} \models \mathcal{G}$).

We use partial orders $<$ on $\mathcal{S}$ (initially empty) and $\sqsubset$ on $\mathcal{G}$, defined by $g_1 \sqsubset g_2$ whenever $code(g_1)$ is a strict subprogram of $code(g_2)$. Without restriction of generality we take the formulas in $\mathcal{G}$ to be terminal (which is natural: a piece of code is specified by stating what the code "does" when it terminates). The verification consists repeatedly applying the following steps while $\mathcal{G} \neq \emptyset$:

- choose $g \in \mathcal{G}$ minimal w.r.t. $\sqsubset$ and prove it, based on Corollary 1;
- remove g from $\mathcal{G}$, transform g into a non-terminal formula (cf. Remark 2) and add the resulting formula g' to $\mathcal{S}$;
- extend $<$ on the newly obtained set $\mathcal{S}$ so that g' is smaller than any formula in $\mathcal{S}$ that can be applied concurrently with g'.

Example 5. Consider the `sum` program in Figure 1. $\mathcal{S}$ consists of the semantical rules of IMP+, and $\mathcal{G}$ consists of formulas (1) and (2) in Section 1, with (2) $\sqsubset$ (1).

At the first iteration (2) is chosen. It is verified based on Corollary 1 (which builds the graph according to the procedure shown in Figure 2), then transformed into a nonterminal formula, removed from $\mathcal{G}$ and added to $\mathcal{S}$. The relation $<$ is extended so that the newly added formula is smaller than the semantical rule for the `while` instruction, since the two rules can be applied concurrently.

At the second (and final) iteration, (1) is verified. The graph-construction procedure exploits the fact that (2) is minimal in $\mathcal{S}$ and thus it will be applied instead of the semantical rule for `while`, producing a finite graph by avoiding an infinite loop unfolding, and allowing Corollary 1 to establish that (1) is valid.

5 Incrementally Verifying the KMP Algorithm

The KMP (Knuth-Morris-Pratt) algorithm is a linear-time string-matching algorithm. The algorithm optimises the naive search of a pattern P into a text T by using some additional information collected from the pattern.

For instance, let us consider $T = $ `ABADABCDA` and $P = $ `ABAC`. It can be easily observed that `ABAC` does not match `ABADABCDA` starting with the first position because there is a mismatch on the fourth position, namely `C` $\neq$ `D`. A naive algorithm, after having detected this, would restart the matching process of P at the second position of T (which fails immediately) then at the third one, where it woud first match an `A` before detecting another mismatch (between `B` and `D`). The KMP optimises this by comparing directly the `B` and `D`, as it "already knows" that they are both preceded by `A`, thereby saving one redundant comparison.

The overall effect is that the worst-case complexity of KMP is determined by the sum of the lengths of P and T, whereas that of a naive algorithm is determined by the product of the two lengths.

The KMP algorithm pre-processes the pattern P by computing a so-called *prefix function* π. Let P_j denote the subpattern of P up to a position j. For such position j, $\pi(j)$ equals the *length of the longest proper prefix of P_j, which is also a suffix of P_j*. In the case of a mismatch between the position i in T and the

```
/*C3: m >= 1 */                              /*Main program */
function computePrefix(){                    /*C6: m>=1 /\ n>=1 */
  k := 0;                                    q := 0;
  pi[1] := 0;                                i := 1;
  q := 2;                                    computePrefix();
  while (q <== m) do                         while (i <== n) do
  /*C2: 0 <= k /\ k < q /\ q <= m+1 /\       /C5: *1 <= m /\ 0 <=q <= m /\ 1 <= i<= n+1 /\
  (forall u:1..k)(p[u]=p[q-1-k+u]) /\          (forall u:1..q)(p[u]=t[i-1-q+u]) /\
  (forall u:1..q-1)(pi[u]=Pi(u)) /\            (forall u:1..m)(pi[u]=Pi(u)) /\
  (forall j)((j > k /\ j in Pi*(q-1))          (exists v)((forall u:v+1..i-1) Theta(u)<m /\
          implies p[j+1] != p[q]) /\                                  allOcc(Out,p,t,v))/\
   k in Pi*(q-1) /\ Pi(q)<=k+1  */            Theta(i)<=q+1 */
    while !(k <== 0) &&                       while !(q <== 0) && !(p[q ++ 1] === t[i]) do
        !(p[k ++ 1] === p[q]) do              /*C4: 1<=m /\ 0<=q<m /\ 1 <= i <= n/\
    /*C1: 0 <= k /\ k < q /\ q <= m /\          (forall u:1..q)(p[u]=t[i-1-q+u]) /\
    (forall u:1..k)(p[u]=p[q-1-k+u]) /\         (forall u:1..m)(pi[u]=Pi(u)) /\
    (forall u:1..q-1)(pi[u]=Pi(u)) /\           (exists v)((forall u:v+1..i-1)
    (forall j)((j > k /\ j in Pi*(q-1))              Theta(u)<m /\ allOcc(Out,p,t,v))/\
              implies p[j+1] != p[q]) /\        Theta(i)<=q+1 */
     k in Pi*(q-1) /\ Pi(q)<=k+1  */            q := pi[q]
     k := pi[k]                                endwhile
    endwhile                                  if (p[q ++ 1] === t[i]) then q := q ++ 1
    if (p[k ++ 1] === p[q]) then              else skip endif;
     k := k ++ 1                              if (q === m) then print (i -- m) ; q := p[q]
    else skip                                 else skip endif;
    endif                                     i := i ++ 1
    pi[q] := k;                              endwhile
    q := q ++ 1;                            /*allOcc(Out, p, t, n) */
  endwhile}
/*(forall u:1..m)(pi[u]=Pi(u)) */
```

Fig. 3. The KMP algorithm in IMP+: prefix function (left) and the main program (right). Grey-text annotations are syntactic sugar for RL formulas. `Pi`, `Theta`, `allOcc`, and `Pi*` denote the functions π, θ, and $allOcc$, and the set π^* respectively (cf. Def. 11).

position j in P, the algorithm proceeds with the comparison of the positions i and $\pi[j]$. This is why, in the above example, KMP direcly compared the B and D.

We prove that the KMP algorithm is correct, i.e., given a non-empty pattern P and a non-empty string T, the algorithm finds *all* the occurrences of P in T. We use the incremental method presented in Section 4 on an encoding of KMP in the IMP+ language formally defined in Maude (cf. Section 2).

The program is shown in Figure 3. Its specification uses the following notions:

Definition 11. – P_j *denotes the prefix of P up to (and including) j. P_0 is the empty string ϵ. If a string P' is a strict suffix of P we write $P' \sqsupset P$.*
  – *The prefix function for P is $\pi : \{1, \ldots, m\} \to \{0, \ldots, m-1\}$ defined by $\pi(i) = max\{j \mid 0 \le j < i \wedge P_j \sqsupset P_i\}$. We let $\pi^*(q) = \{\pi(q), \pi(\pi(q)), \ldots\}$.*
  – *Let T be a string of length n. We define $\theta : \{1, \ldots, n\} \to \{0, \ldots, m\}$ the function which, for a given $i \in \{1, \ldots, n\}$, returns the longest prefix of P which is a suffix of T_i: $\theta(i) = max\{j \mid 0 \le j \le m \wedge P_j \sqsupset T_i\}$.*
  – *Let T be a string of length n and Out a list. The function $allOcc(Out, P, T, i)$ returns true iff the list Out contains all the occurrences of P in $T[1..i]$.*

The grey-text annotations, written as pre/post conditions and invariants, are syntactical sugar for RL formulas. The annotations are numbered (C_1 to C_6)

according to the order in which the RL formulas are verified by our incremental method. So, for example, the annotation for the inner loop of the `computePrefix` function is the first to be verified, and corresponds to an RL formula for the form

$$\langle \text{while C do} \ldots \text{endwhile}, \ldots \rangle \wedge C \wedge C_1 \Rightarrow \langle \text{skip}, \ldots \rangle \wedge \neg C \wedge C_1$$

where `while C do ... endwhile` denotes the inner loop of `computePrefix`. Similarly, the specification of the KMP program is an RL formula of the form:

$$\langle \text{KMP}, \ldots, \text{.Ints} \rangle \wedge C_6 \Rightarrow \langle \text{skip}, \ldots, \text{Out} \rangle \wedge allOcc(\text{Out}, \ldots),$$

where `KMP` denotes the whole program, `.Ints` denotes an empty list of integers (cf. Section 2), `Out` is a list of integers denoting the program's output, and $allOcc(\text{Out}, \ldots)$ states that `Out` contains all positions of the pattern in the text.

The RL formulas corresponding to the annotations ($C_1 \ldots C_6$) were verified in the given order. Once a formula was verified, it was generalised (cf. Remark 2) and added to the rules denoting the semantics of IMP+ as new, prioritary rules. Each rule verification follows the construction of a graph (cf. procedure in Figure 2), performed by symbolic execution, implemented by rewriting as described in Section 3. For this purpose we have intensively use Maude's metalevel mechanisms in order to control the application of rewrite rules.

The main verification effort (besides coming up with the annotations $C_1 \ldots C_6$) went into the inclusion test between patterns that occurs in our graph-construction procedure. For this purpose we have used certain properties of the $\pi, \pi*$, and θ mathematical functions from [12], which we include in Maude as equations used for the purpose of simplification. Some elementary simplifications involving properties of integers and Booleans were performed via Maude's interface to the Z3 solver. Collectively, these properties can be seen as axioms that define the class of models in which the correctness of our KMP program holds.

Benefits of incremental verification. In earlier work [5] we attempted to verify KMP using a circular approach of the "all-or-nothing" variety. The main difficulty with such approaches is that, if verification fails, one is left with nothing: any of the formulas being (simultaneously) verified could be responsible for the failure. The consequence was that (as we realised afterwards by revisiting the problem) our earlier verification was incorrect. We found some versions of the annotations $C_1 \ldots C_6$, which, as RL formulas, would only hold under unrealistic assumptions about the problem-domain functions π, π^*, and θ.

We decided to redo the KMP verification incrementally, starting with smaller program fragments, and rigorously proving at each step the required facts about the problem domain. Our incremental approach was first a language-*dependent* one [12], as it was based on proving pre/post conditions of functions and loop invariants. Of course, not all languages have the same kinds of functions and loops; some lack such constructions altogether. The method proposed in this paper is (with some restrictions) both incremental and language-independent, is formally proved correct, and was instrumental in successfully proving the KMP program, this time, under valid assertions regarding the problem domain.

6   Conclusion, Related Work and Future Work

In this paper we propose an incremental method for proving a class of RL formulas useful in practical situations. Mainly, RL formula verification is reduced to checking two technical conditions: the first is an invariance property, while the second is related to the so-called capturing of terminal configurations. Formally, the conjunction of these conditions is shown to be equivalent to RL formula validity. We also present a graph construction procedure based on symbolic execution which, if it terminates successfully, ensures that these conditions hold for a given RL formula. The method is successfully applied on the nontrivial Knuth-Morris-Pratt algorithm for string matching, encoded in a simple imperative language. The syntax and the semantics of this language have been defined in Maude, whose reflective features were intensively used for implementation purposes.

Using the proposed approach RL formulas are proved in a systematic manner. One first proves formulas that specify sub-programs of the program under verification, and then exploits the newly proved formulas to (incrementally) prove other formulas that specify larger subprograms. By contrast, monolithic/circular approaches [1,2,3,4,13,6] attempt to prove all formulas at once, in no particular order. In case of failure, in a monolithic approach, *any* circularly dependent subset of formulas under proof might be responsible for the failure; whereas in an incremental approach, there is only one subset of formulas to consider (and to modify in order to progress in the proof): the formula currently under proof, together with some already proved valid formulas. Thus, an incremental method saves the user some effort in the trial-and-error process of program verification.

Related work. Besides the already mentioned work on RL we cite some approaches in program verification; an exhaustive list is outside the scope of this paper.

Some approaches are based on exploring the state-space of a program, e.g., [14], in which software model checking is combined with symbolic execution and abstraction techniques to overcome state-space explosion. Our approach has some similarities with the above: we also use symbolic execution to construct a graph, which is an abstraction of the reachable state space of a program.

Some verification tools (e.g., Why3 [15]) are based on deductive methods. These tools use the program specifications (i.e., pre/post-conditions, invariants) to generate proof obligations, which are then discharged to external provers (e.g., COQ, Z3, . . .). Similarly, our implementation uses a version of Maude which includes a connection to the Z3 SMT solver (used for simplifying conditions).

In the same spirit, compositional methods for the formal verification (e.g., [16]) shift the focus of verification from global to local level in order to reduce the complexity of the verification process.

Future work. One issue that needs to be addressed is the handling of domain-specific properties. Each program makes computations over a certain domain (e.g., arrays), and in order to prove a program, certain properties of the underlying domain are required (e.g., relations between selecting and storing elements in an array). Currently, these properties are stated as axioms in Maude, and we are planning to connect Maude to an inductive prover in order to interactively prove the axioms in questions as properties satisfied by more basic definitions.

References

1. Grigore Roşu and Andrei Ştefănescu. Towards a unified theory of operational and axiomatic semantics. In *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP'12)*, volume 7392 of *Lecture Notes in Computer Science*, pages 351–363. Springer, 2012.

2. Grigore Roşu and Andrei Ştefănescu. Checking reachability using matching logic. In *Proceedings of the 27th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'12)*, pages 555–574. ACM, 2012.

3. Grigore Roşu, Andrei Ştefănescu, Ştefan Ciobâcă, and Brandon M. Moore. One-path reachability logic. In *Proceedings of the 28th Symposium on Logic in Computer Science (LICS'13)*, pages 358–367. IEEE, June 2013.

4. Andrei Ştefănescu, Ştefan Ciobâcă, Radu Mereuţă, Brandon M. Moore, Traian Florin Şerbănuţă, and Grigore Roşu. All-path reachability logic. In *Proceedings of the Joint 25th International Conference on Rewriting Techniques and Applications and 12th International Conference on Typed Lambda Calculi and Applications (RTA-TLCA'14)*, volume 8560 of *LNCS*, pages 425–440. Springer, July 2014.

5. Andrei Arusoaie, Dorel Lucanu, and Vlad Rusu. A Generic Framework for Symbolic Execution: Theory and Applications. Research Report RR-8189, Inria, September 2015.

6. Andrei Arusoaie, Dorel Lucanu, and Vlad Rusu. A Generic Framework for Symbolic Execution. Research Report RR-8189, Inria, September 2015. Available at `https://hal.inria.fr/hal-00766220`.

7. Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn Talcott. *All about Maude – a high-performance logical framework: how to specify, program and verify systems in Rewriting Logic*. Springer-Verlag, 2007.

8. Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS'08*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.

9. Grigore Roşu, Chucky Ellison, and Wolfram Schulte. Matching logic: An alternative to Hoare/Floyd logic. In Michael Johnson and Dusko Pavlovic, editors, *Proceedings of the 13th International Conference on Algebraic Methodology And Software Technology (AMAST '10)*, volume 6486 of *Lecture Notes in Computer Science*, pages 142–162, 2010.

10. Grigore Roşu and Andrei Ştefănescu. Matching Logic: A New Program Verification Approach (NIER Track). In *ICSE'11: Proceedings of the 30th International Conference on Software Engineering*, pages 868–871. ACM, 2011.

11. Grigore Roşu. Matching logic — extended abstract. In *Proceedings of the 26th International Conference on Rewriting Techniques and Applications (RTA'15)*, volume 36 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5–21, Dagstuhl, Germany, July 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

12. Verification of the KMP algorithm. `https://fmse.info.uaic.ro/imgs/kmp.pdf`.

13. Dorel Lucanu, Vlad Rusu, Andrei Arusoaie, and David Nowak. Verifying reachability-logic properties on rewriting-logic specifications. In Narciso Martí-Oliet, Peter Csaba Ölveczky, and Carolyn L. Talcott, editors, *Logic, Rewriting, and Concurrency - Essays dedicated to José Meseguer on the Occasion of His 65th Birthday*, volume 9200 of *Lecture Notes in Computer Science*, pages 451–474. Springer, 2015.

14. Willem Visser, Klaus Havelund, Guillaume P. Brat, Seungjoon Park, and Flavio Lerda. Model checking programs. *Autom. Softw. Eng.*, 10(2):203–232, 2003.

15. Jean-Christophe Filliâtre and Andrei Paskevich. Why3 — where programs meet provers. In Matthias Felleisen and Philippa Gardner, editors, *Proceedings of the 22nd European Symposium on Programming*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128. Springer, March 2013.

16. Willem P. de Roever, Frank S. de Boer, Ulrich Hannemann, Jozef Hooman, Yassine Lakhnech, Mannes Poel, and Job Zwiers. *Concurrency Verification: Introduction to Compositional and Noncompositional Methods*, volume 54 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.

17. Vlad Rusu, Dorel Lucanu, Traian-Florin Serbanuta, Andrei Arusoaie, Andrei Stefanescu, and Grigore Rosu. Language definitions as rewrite theories. *J. Log. Algebr. Meth. Program.*, 85(1):98–120, 2016.

Appendix: Proofs

Results from Section 3

This section is dedicated to proving Lemma 1 regarding the simulation result between the relations $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ and $\Rightarrow_{\mathcal{S}}$, where $\Rightarrow_{\mathcal{S}}$ is the transition relation induced by a set of RL formulas $\mathcal{S}$ (Definition 4) and $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ is the relation induced by rewriting with a the set of formulas $\mathcal{S}^{\mathfrak{s}}$ (Definition 5). Proving Lemma 1 is actually quite a challenge; even though it is background material regarding symbolic execution we include full details here in order to make this document self-contained.

The proof includes three main steps:

1. establishing a mutual simulation between $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ and another relation denoted by $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$, in which the details about rewriting occuring in $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ are spelled out;
2. proving the simulation of the relation $\Rightarrow_{\mathcal{S}}$ by the relation $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ in two steps:
   – first, when the relations as defined by RL formulas without quantifiers;
   – then, for RL formulas with quantified variables in the right-hand side.

A relation $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ and its mutual simulation with $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ The following definition introduces thr relation $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ on patterns. It is a version of the rewriting-based relation $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ where details about rewriting (matching with a substitution, applying the substitution to the right-hand side of a rule) are spelled out.

Definition 12 (relation $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$). *Let φ be a pattern with $\varphi \triangleq (\exists X)\pi \wedge \phi$ and $\alpha \triangleq \varphi_l \Rightarrow \varphi_r$ be a RL formula with $\varphi_l \triangleq \pi_l \wedge \phi_l$, $\varphi_r \triangleq (\exists Y)\pi_r \wedge \phi_r$, and $(X \cup \mathit{FreeVars}(\varphi)) \cap (Y \cup \mathit{FreeVars}(\varphi_l, \varphi_r)) = \emptyset$. We write $\varphi \Rightarrow_{\alpha}^{\mathfrak{s}} \varphi'$, with $\varphi' \triangleq (\exists X, Y)\pi' \wedge \phi'$, whenever there exists a matcher $\sigma \in \mathit{match}_{\cong}(\pi, \pi_l)$ such that $\pi' = \sigma'(\pi_r)$ and $\phi' = \phi \wedge \sigma'(\phi_l \wedge \phi_r)$ where $\sigma' = \sigma \cup \mathit{Id}|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)}$.*

The following lemma establishes a mutual simulation between $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ and $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$. Remember that patterns $(\exists X)\pi \wedge \phi$ can equivalently be seen as terms of sort $\mathit{Cfg}^{\mathfrak{s}}$, and that we extended the congruence $\cong$ from terms of sort Cfg to to terms of sort $\mathit{Cfg}^{\mathfrak{s}}$, as follows: $(\exists X)\pi_1 \wedge \phi \cong (\exists X)\pi_2 \wedge \phi$ iff $\pi_1 \cong \pi_2$.

For convenience we also recall here the definition of the relation $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$ (Def. 5): for $\alpha^{\mathfrak{s}} \triangleq (\exists \mathcal{X})\pi_l \wedge \psi \Rightarrow (\exists \mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r) \in \mathcal{S}^{\mathfrak{s}}$ we write $(\exists X)\pi \wedge \phi \Rightarrow_{\alpha^{\mathfrak{s}}} (\exists X, Y)\pi' \wedge \phi'$ whenever $(\exists X)\pi \wedge \phi$ is rewritten by $\alpha^{\mathfrak{s}}$ to $(\exists X, Y)\pi' \wedge \phi'$, i.e., there exists a substitution σ'' such that $\sigma''((\exists \mathcal{X})\pi_l \wedge \psi) \cong (\exists X)\pi \wedge \phi$ and $(\exists X, Y)\pi' \wedge \phi' = \sigma''((\exists \mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r))$.

Lemma 2. $(\exists X)\pi \wedge \phi \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} (\exists X, Y)\pi' \wedge \phi'$ *iff* $(\exists X)\pi \wedge \phi \Rightarrow_{\mathcal{S}^{\mathfrak{s}}} (\exists X, Y)\pi' \wedge \phi'$.

Proof. ($\Rightarrow$) Assume $(\exists X)\pi \wedge \phi \Rightarrow_{\mathcal{S}}^{\mathfrak{s}} (\exists X, Y)\pi' \wedge \phi'$, thus, there exists $\alpha \triangleq \pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r \in \mathcal{S}$, $\sigma : \mathit{FreeVars}(\pi_l) \rightarrow T_{\Sigma}(\mathit{FreeVars}(\pi)) \in \mathit{match}_{\cong}(\pi, \pi_l)$, and $\sigma' \triangleq \sigma \cup \mathit{Id}|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)}$ such that $\pi' = \sigma'(\pi_r)$ and $\phi' = \phi \wedge \sigma'(\phi_l \wedge \phi_r)$.

Consider the rule $\alpha^{\mathfrak{s}} \triangleq (\exists \mathcal{X})\pi_l \wedge \psi \Rightarrow (\exists \mathcal{X}, Y)\pi_r \wedge (\psi \wedge \phi_l \wedge \phi_r) \in \mathcal{S}^{\mathfrak{s}}$ and the substitution $\sigma'' \triangleq \sigma' \cup (\mathcal{X} \leftarrow X) \cup (\psi \leftarrow \phi)$. We have:

1. $\sigma''((\exists\mathcal{X})\pi_l\wedge\psi) \cong (\exists X)\pi\wedge\phi$; (remember how we extended the congruence $\cong$ to terms in $Cfg^{\mathfrak{s}}$ by letting $(\exists X)\pi_1\wedge\phi \cong (\exists X)\pi_2\wedge\phi$ if and only if $\pi_1 \cong \pi_2$);
2. $\sigma''((\exists\mathcal{X},Y)\pi_r\wedge(\psi \wedge \phi_l \wedge \phi_r)) = (\exists X,Y)\sigma'(\pi_r)\wedge(\phi \wedge \sigma'(\phi_l \wedge \phi_r)) = (\exists X,Y)\pi'\wedge\phi'$.

Thus, $(\exists X)\pi\wedge\phi \Rightarrow_{\alpha^{\mathfrak{s}}} (\exists X,Y)\pi'\wedge\phi'$, which proves the ($\Rightarrow$) implication.

($\Leftarrow$) Assume $(\exists X)\pi\wedge\phi \Rightarrow_{\mathcal{S}^{\mathfrak{s}}} (\exists X,Y)\pi'\wedge\phi'$. Thus, there exist $\alpha^{\mathfrak{s}} \triangleq (\exists\mathcal{X})\pi_l\wedge\psi \Rightarrow (\exists\mathcal{X},Y)\pi_r\wedge(\psi \wedge \phi_l \wedge \phi_r) \in \mathcal{S}^{\mathfrak{s}}$ and a substitution σ'', mapping $FreeVars((\exists\mathcal{X})\pi_l\wedge\psi)$ to terms over $FreeVars((\exists X)\pi\wedge\phi)$, such that

1. $\sigma''((\exists\mathcal{X})\pi_l\wedge\psi) \cong (\exists X)\pi\wedge\phi$, where the congruence $\cong$ over $Cfg^{\mathfrak{s}}$ satisfies $(\exists X)\pi_1\wedge\phi \cong (\exists X)\pi_2\wedge\phi$ if and only if $\pi_1 \cong \pi_2$. *Due to this property*, σ'' has the form $\sigma' \cup (\mathcal{X} \leftarrow X) \cup (\psi \leftarrow \phi)$ with $\sigma' = \sigma \cup Id|_{Var\setminus FreeVars(\pi_l)}$ and $\sigma : FreeVars(\pi_l) \rightarrow T_\Sigma(FreeVars(\pi)) \in match_{\cong}(\pi, \pi_l)$. We obtain $\sigma'(\pi_l) \cong \pi$ and $\sigma''(\psi) = \phi$ and and $\sigma''(\mathcal{X}) = X$;
2. $\sigma''((\exists\mathcal{X},Y)\pi_r\wedge(\psi \wedge \phi_l \wedge \phi_r)) = (\exists X,Y)\pi'\wedge\phi'$, thus, we obtain $\sigma''(\pi_r) = \sigma'(\pi_r) = \pi'$ and also $\sigma''(\psi \wedge \phi_l \wedge \phi_r) = \phi \wedge \sigma'(\phi_l \wedge \phi_r) = \phi'$

We have thus obtained $\pi' = \sigma'(\pi_r)$ and $\phi' = \phi \wedge \sigma'(\phi_l \wedge \phi_r)$ for some σ' s.t. $\sigma'(\pi_l) \cong \pi$, which implies $(\exists X)\pi\wedge\phi \Rightarrow_\alpha^{\mathfrak{s}} (\exists X,Y)\pi'\wedge\phi'$. Thus proves the ($\Leftarrow$) implication and the lemma. $\qquad\square$

Simulation of $\Rightarrow_{\mathcal{S}}$ by $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$. Now that we have established mutual simulation of $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ and $\Rightarrow_{\mathcal{S}^{\mathfrak{s}}}$, there remains to establish the simulation of $\Rightarrow_{\mathcal{S}}$ by $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$. We will do this in two steps. In the first step we shall consider relations $\Rightarrow_{\mathcal{S}}$ and $\Rightarrow_{\mathcal{S}}^{\mathfrak{s}}$ defined by RL formulas $\varphi \Rightarrow \varphi'$ in which existential quantifiers in the right-hand side have been replaced by fresh variables. Thus, we have formulas satisfying $FreeVars(\varphi') \not\subseteq FreeVars(\varphi)$, for which the definition of validity becomes:

Definition 13. *An* RL *formula* $\varphi \Rightarrow \varphi'$ *with* $FreeVars(\varphi') \not\subseteq FreeVars(\varphi)$ *is valid w.r.t. a set* $\mathcal{S}$ *of* RL *formulas, written* $\mathcal{S} \models \varphi \Rightarrow \varphi'$*, if for all pairs* (γ_0, ρ) *such that* $(\gamma_0, \rho) \models \varphi$*, and all complete paths* $\gamma_0 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n$*, there exists* $0 \le i \le n$ *and a valuation* ρ' *with* $\rho'|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$ *such that* $(\gamma_i, \rho') \models \varphi'$*.*

Remark 3. The new valuation ρ' is required in order to avoid the undesired capturing of additional variables (in $FreeVars(\varphi')\setminus FreeVars(\varphi)$) by the valuation ρ; however, for variables of φ the two valuations coincide. Note that this notion of validity is obtained from the standard one (Definition 4) applied to $\varphi \Rightarrow (\exists Y)\varphi'$.

We shall be using the following result, which holds thanks to Assumption 2. Its proof can also be found in [17] but we give here a more compact proof.

Definition 14. *For a valuation* $\rho : Var \rightarrow M$ *and a substitution* $\sigma : X \rightarrow T_\Sigma(Y)$*, we write* $\rho \prec \sigma$ *if* $\rho|_X = (\rho \circ \sigma)|_X$*.*

Lemma 3. *Let* $\sigma_1 : X \rightarrow T_\Sigma(Y)$ *and* $\sigma_2 : Z \rightarrow T_\Sigma(X)$*. If* $\rho \prec \sigma_1$ *and* $\rho \prec \sigma_2$ *then* $\rho \prec \sigma_1 \circ \sigma_2$*.*

Proof. Let $z \in Z$ be chosen arbitrarily. We have to show that $\rho(z) = \rho(\sigma_1(\sigma_2(z)))$.

Now, $t_2 \triangleq \sigma_2(z)$ is a term over $T_\Sigma(X)$, and $t_1 \triangleq \sigma_1(\sigma_2(z))$ is a term over $T_\Sigma(Y)$, obtained by replacing in t_2 each $x \in \textit{FreeVars}(t_2)$ by $\sigma_1(x) \in T_\Sigma(Y)$.

From $\rho \prec \sigma_1$ we know that $\rho(x) = \rho(\sigma_1(x))$ for all $x \in \textit{FreeVars}(t_2)$. Thus, $\rho(t_2)$, obtained by replacing in t_2 function symbols f by their ρ-interpretation f_ρ and variables in $\textit{FreeVars}(t_2)$ by their ρ-valuation $\rho(x)$, is also equal to the value obtained by function symbols by their ρ-interpretation *and variables* $x \in$ $\textit{FreeVars}(t_2)$ *by* $\rho(\sigma_1(x))$. But the latter value is exactly $\rho(\sigma_1(t_2)) = \rho(\sigma_1(\sigma_2(z)))$.

Thus, $\rho(t_2) = \rho(\sigma_1(\sigma_2(z)))$, and from $\rho \prec \sigma_2$ we obtain $\rho(z) = \rho((\sigma_2(z)) = \rho(t_2)$. The conclusion follows by transitivity of equality. $\qquad\square$

Remark 4. The domains and codomains of substitutions σ_1, σ_2 in Lemma 3 were chosen so that the composition $\sigma_1 \circ \sigma_2 : Z \to T_\Sigma(Y)$ is well defined. The lemma can be generalized to for substitutions $\sigma_1 : X \to T_\Sigma(Y)$ and $\sigma_2 : Z \to T_\Sigma(X')$ with $X \neq X'$ as follows: consider the substitution σ_1' with a domain $X \cup X'$, which extends σ_1 extended as the identity on variables in $X' \setminus X$. By slight notation abuse we can define $\sigma_1 \circ \sigma_2$ as the (properly defined) composition $\sigma_1' \circ \sigma_2$. Now, from $\rho \prec \sigma_1$ we trivially obtain $\rho \prec \sigma_1'$, and, together with the hypothesis $\rho \prec \sigma_2$ Lemma 3 then gives us $\rho \prec \sigma_1' \circ \sigma_2$, i.e., $\rho \prec \sigma_1 \circ \sigma_2$ since we defined $\sigma_1 \circ \sigma_2 \triangleq \sigma_1' \circ \sigma_2$. Thus, with the above generalized definition for the composition of (otherwise non-composable) substitutions, Lemma 3 holds. This is used below.

Lemma 4 (Unification by Matching). *For non-variable terms $t \in T_\Sigma(\textit{Var}^b)$, linear terms $t' \in T_{\Sigma \setminus \Sigma^b}(\textit{Var})$, and all valuations $\rho : \textit{Var} \to M$ such that $\rho(t) = \rho(t')$, there exists a substitution σ such that $t =_A \sigma(t')$ and $\rho \prec \sigma$.*

Proof. Let $\mathcal{D}$ be a model of the builtin subtheory Σ^b, satisfying any builtin axioms that the signature may have. Let A be the set of non-builtin axioms, which by an assumption in the paper are known to be linear, regular and non-collapsing. We consider the model M to be $T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}/A$, i.e., equivalence classes modulo A of terms in which the only subterms of a builtin sort are constants in $\mathcal{D}$.

The equality $\rho(t) = \rho(t')$ implies that for all $\hat{t} \in \rho(t)$ there exists $\hat{t}' \in \rho(t')$ such that $\hat{t} =_A \hat{t}'$. We fix arbitrarily such terms $\hat{t}, \hat{t}' \in T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$. From $\hat{t}' \in \rho(t')$ we get that $\hat{t}'$ is obtained from $t' \in T_{(\Sigma \setminus \Sigma^b)}(\textit{Var})$ by substituting variables $x_1 \dots x_n \in \textit{FreeVars}(t')$ with representatives $\widehat{\rho(x_1)} \dots \widehat{\rho(x_n)}$ chosen in, respectively, $\rho(x_1) \dots \rho(x_n)$, i.e., $t' = c[x_1 \dots x_n]$ and $\hat{t}' = c[\widehat{\rho(x_1)} \dots \widehat{\rho(x_n)}]$ for some context c. The relation $\hat{t} =_A \hat{t}'$ is obtained using a finite number of axioms in A.

1. First we consider the case when no axioms are applied: $\hat{t} = \hat{t}'$, i.e., the two terms are syntactically equal. We prove by induction on positions (strings over natural numbers $\mathbb{N}$)) that ($\spadesuit$) *any position ω in t' is also a position in t, and if t'_ω is a non-variable position then so is t_ω and the top function symbols of t'_ω and t_ω coincide.*

   (a) in the base case ω is the empty string which is obviously a position of t. Since t, t' are not variables they have some top symbols f resp. g. From $\hat{t} = \hat{t}'$ we get $f = g$ which proves the base case.

(b) for the inductive step: let ω be a position in t' of length $k+1$. Thus, there is a position ω' of length k, where $t'_{\omega'}$ has the form $f(\tau_1, \ldots \tau_q)$ where $q > 0$ and f is a non-builtin function symbol. By induction hypothesis the position ω' is also a position of t. Now, $t_{\omega'}$ cannot be a variable, because if this was the case, in $\hat{t}$ the subterm $\hat{t}_{\omega'}$ would have a builtin sort, which cannot be equal to the non-builtin $\hat{t}'_{\omega'}$. Thus, $t_{\omega'}$ is of the form $g(\tau'_1, \ldots \tau'_r)$ for some $r \geq 0$. Now, the syntactical equality $\hat{t} = \hat{t}'$ implies $\hat{t}_{\omega'} = \hat{t}'_{\omega}$, where $\hat{t}$ has top symbol g and $\hat{t}$ has top symbol f: we obtain $f = g$ and $r = q > 0$. In particular the position ω of length $k+1$ is also a position of t. Thus, t_ω exists, and so does $\hat{t}_\omega$, and $\hat{t} = \hat{t}'$ also implies $\hat{t}_\omega = \hat{t}'_\omega$. Using the same reasoning as in the base case above we obtain that if t'_ω is not a variable then so is t_ω and the the top function symbols of t'_ω and t_ω coincide, which proves the inductive step and ($\spadesuit$).

Since $t' = c[x_1, \ldots x_n]$ we obtain in particular that the positions $\omega(x_1), \omega(x_n)$ of the variables $x_1, \ldots x_n$ in t' are also positions in t. We build the substitution σ by mapping x_i to $t_{\omega(x_i)}$ for $i = 1, \ldots, n$ - it is a substitution, since t' is linear- and we obtain $\sigma(t') = t$. Moreover, from $\hat{t}' = c[\widehat{\rho(x_1)} \ldots \widehat{\rho(x_n)}]$ we obtain $\widehat{\rho(x_i)} = \hat{t}'_{\omega(x_i)} = \hat{t}_{\omega(x_i)} \in \rho(t_{\omega(x_i)})$ and then $\rho(x_i) = [\widehat{\rho(x_i)}]_A = [\hat{t}_{\omega(x_i)}]_A = \rho(t_{\omega(x_i)}) = (\rho(\sigma((x_i)) = (\rho \circ \sigma)(x_i)$ which proves $\rho \prec \sigma$.

2. Next, we consider the case in which only one axiom, say, $u = v$ is involved once in establishing $\hat{t} =_A \hat{t}'$. By our assumption of the axioms A, u and v are linear terms that only contain non-data operations, and have the same set of free variables, say, $\{y_1, \ldots y_m\}$. We assume without restriction of generality that $y_i \notin vars(t, t')$ for $i = 1, \ldots, m$. Thus, there exists a substitution $\mu : \{y_1, \ldots, y_n\} \to T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$ and a common position ω of $\hat{t}$ and $\hat{t}'$ such that $\mu(u) = \hat{t}_\omega$, $\mu(v) = \hat{t}'_\omega$, and $\hat{t}[\mu(u)]_\omega = \hat{t}'[\mu(u)]_\omega$.

  – using $\hat{t}_\omega = \mu(u)$ together with $\hat{t}_\omega \in T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$ and the fact that that u is linear and has no builtin subtems other than variables allows us to apply the same reasoning as in Item 1, and obtain $\sigma_u : FreeVars(u) \to T_\Sigma(FreeVars(t_\omega))$ such that $\sigma_u(u) = t_\omega$. Next, we note that our lemma holds or not independently of the value of ρ in $Var \setminus FreeVars(t, t')$. Thus, we can assume $\rho(y) = \rho(\sigma_u(y))$ for all $y \in FreeVars(u)$, i.e., $\rho \prec \sigma_u$.

  – using $\mu(v) = \hat{t}'_\omega$ and the fact that t'_ω is linear and $\mu(v) \in T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$, we apply again the same reasoning as in Item 1 and obtain $\sigma_v : FreeVars(t'_\omega) \to T_\Sigma(FreeVars(v))$ such that $\sigma_v(t'_\omega) = v$ and $\rho \prec \sigma_v$.

  – from $\hat{t}[\mu(u)]_\omega = \hat{t}'[\mu(u)]_\omega$, by using once more the reasoning in Item 1, we obtain that for each position ω' of $t'[u]_\omega$ *that is not under* ω, ω' is also a position of $t[u]_\omega$ and if ω' is not a variable position of t' then the top symbols of $(t'[u]_\omega)_{\omega'}$ and $(t[u]_\omega)_{\omega'}$ coincide. We get a substitution $\sigma' : FreeVars(t') \setminus FreeVars(t'_\omega) \to T_\Sigma(FreeVars(t))$, such that $\rho \prec \sigma'$.

  – then, we build $\sigma' \circ \sigma_u : (FreeVars(t') \setminus FreeVars(t'_\omega)) \uplus FreeVars(u) \to T_\Sigma(FreeVars(t))$. By the properties above we have $(\sigma' \circ \sigma_u)(t'[u]_\omega) = t$.

  – next, $t'[u]_\omega =_A t'[v]_\omega = \sigma_v(t')$, thus, $((\sigma' \circ \sigma_u) \circ \sigma_v)(t') =_A t$.

  – finally, let $\sigma \triangleq ((\sigma' \circ \sigma_u) \circ \sigma_v)$; we have obtained $\sigma(t') =_A t$ above, and $\rho \prec \sigma$ follows from $\rho \prec \sigma_v$, $\rho \prec \sigma_u$, $\rho \prec \sigma'$ and Remark 4.

3. There remains to consider the general case when several axioms are involved in the relation $\hat{t} =_A \hat{t}'$. Thus, there are axioms $a_1, \ldots, a_p$ $(p > 1)$ and terms in $T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$: $\hat{t}_0 \triangleq \hat{t}, \hat{t}_1, \ldots, \hat{t}_p \triangleq \hat{t}'$ such that $\hat{t} \triangleq \hat{t}_0 =_{a_1} \hat{t}_1 \cdots =_{a_p} \hat{t}_p \triangleq \hat{t}'$. Now, for each of the *ground* terms $\hat{t}_1, \ldots, \hat{t}_{p-1}$ in the sequence there exists a corresponding term *with variables* $t_1, \ldots, t_{p-1}$, such that t_i is obtained from $\hat{t}_i$ by substituting constants in $\mathcal{D}$ with *fresh* variables in Var^b. Thus, $t_1 \ldots t_{p-1}$ are linear and belong to $T_{(\Sigma \setminus \Sigma^b)}(Var^b)$ (since the corresponding ground terms $\hat{t}_i$ are in $T_{(\Sigma \setminus \Sigma^b)(\mathcal{D})}$). Remember also that $t' \in T_{(\Sigma \setminus \Sigma^b)}(Var)$ is linear by hypothesis. Thus, we can repeatedly apply the reasoning at Item 2 and obtain substitutions $\sigma_i : FreeVars(t_i) \to T_\Sigma(FreeVars(t_{i-1}))$ such that $\sigma_i(t_i) =_A t_{i-1}$ and $\rho \prec \sigma_i$, for $i = 1, \ldots, p$. With $\sigma = \sigma_p \circ \sigma_{p-1} \cdots \circ \sigma_1$ we get $\sigma(t') =_A t$ and (by Remark 4) $\rho \prec \sigma$. This concludes the lemmas's proof. $\square$

We shall also be using the following remarks, which directly follow from definitions in the paper.

Remark 5. If $\mathcal{S} \models \alpha$ with $\alpha \triangleq \varphi \Rightarrow \varphi'$ then for all pairs (γ_0, ρ) such that $(\gamma_0, \rho) \models \varphi$, all complete paths $\gamma_0 \Rightarrow_\mathcal{S} \cdots \Rightarrow_\mathcal{S} \gamma_n$, we know there exists $i \in \{0, \ldots, n\}$ such that $(\gamma_i, \rho') \models \varphi'$ for some ρ' with $\rho'|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$. Then $\gamma_0 \Rightarrow_\alpha \gamma_i$.

Remark 6. From Assumptions 1 and 3 it follows that for rules $\pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r$, $FreeVars(\pi_r \wedge \phi_r) \setminus FreeVars(\pi_l \wedge \phi_l) \subseteq Y \subseteq Var^b$.

The next remark regards the definition of the relation $\Rightarrow_\mathcal{S}^\mathfrak{s}$ restricted to unquantified patterns and generated for unquantified RL formulas (but with additional variables in their right-hand side) as discussed above. It is essentially Definition 12 in which the sets X, Y of quantified variables are both empty.

Remark 7. Let φ be a pattern with $\varphi \triangleq \pi \wedge \phi$ and $\alpha \triangleq \varphi_l \Rightarrow \varphi_r$ be a RL formula with $\varphi_l \triangleq \pi_l \wedge \phi_l$, $\varphi_r \triangleq \pi_r \wedge \phi_r$, and $FreeVars(\varphi) \cap FreeVars(\varphi_l, \varphi_r) = \emptyset$. We have $\varphi \Rightarrow_\alpha^\mathfrak{s} \varphi'$, with $\varphi' \triangleq \pi' \wedge \phi'$, whenever there exists a matcher $\sigma \in match_\cong(\pi, \pi_l)$ such that $\pi' = \sigma'(\pi_r)$ and $\phi' = \phi \wedge \sigma'(\phi_l \wedge \phi_r)$ where $\sigma' = \sigma \cup Id|_{Var \setminus FreeVars(\pi_l)}$.

Lemma 5 (restricted $\Rightarrow_\mathcal{S}^\mathfrak{s}$ simulates $\Rightarrow_\mathcal{S}$). *For all $\gamma, \gamma' \in M_{Cfg}$, pattern $\varphi \triangleq \pi \wedge \phi$ with $FreeVars(\varphi) \subseteq Var^b$, and valuation ρ, if $(\gamma, \rho) \models \varphi$ and $\gamma \Rightarrow_\alpha \gamma'$ then there is $\varphi' \triangleq \pi' \wedge \phi'$ with $FreeVars(\varphi') \subseteq Var^b$ such that $\varphi \Rightarrow_\alpha^\mathfrak{s} \varphi'$ and $(\gamma', \rho') \models \varphi'$, for some valuation ρ' such that $\rho'|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$.*

Proof. From $\gamma \Rightarrow_\alpha \gamma'$ we obtain the rule $\alpha \triangleq \varphi_l \Rightarrow \varphi_r \in \mathcal{S}$ and $\varphi_l \triangleq \pi_l \wedge \phi_l$, $\varphi_r \triangleq \pi_r \wedge \phi_r$, and a valuation μ such that $(\gamma, \mu) \models \pi_l \wedge \phi_l$ and $(\gamma', \mu) \models \pi_r \wedge \phi_r$. Since the rules are defined up to the names of their free variables, we can assume $FreeVars(\varphi) \cap FreeVars(\varphi_l, \varphi_r) = \emptyset$. Let then ρ'' be any valuation such that $\rho''|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$, $\rho''|_{FreeVars(\varphi_l, \varphi_r)} = \mu|_{FreeVars(\varphi_l, \varphi_r)}$. We thus have

1. $(\gamma, \rho'') \models \pi \wedge \phi$, hence, (i) $\gamma = \rho''(\pi)$ and (iv) $\rho'' \models \phi$;
2. $(\gamma, \rho'') \models \pi_l \wedge \phi_l$, hence, (ii) $\gamma = \rho''(\pi_l)$ and (v) $\rho'' \models \phi_l$;
3. $(\gamma', \rho'') \models \pi_r \wedge \phi_r$, hence, (iii) $\gamma' = \rho''(\pi_r)$ and (vi) $\rho'' \models \phi_r$.

From (i) and (ii) we obtain $\rho''(\pi) = \rho''(\pi_l)$ and, using Lemma 4 (unification by matching) we obtain $\sigma : \mathit{FreeVars}(\pi_l) \to T_\Sigma(\mathit{FreeVars}(\pi)) \in \mathit{match}_\cong(\pi, \pi_l)$ such that $\rho'' \prec \sigma$, that is, there exists a valuation $\eta : \mathit{Var} \to M$ such that $\rho''|_{\mathit{FreeVars}(\pi_l)} = (\eta \circ \sigma)|_{\mathit{FreeVars}(\pi_l)}$. Let $\sigma' \triangleq \sigma \cup Id|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)}$. Then consider a valuation η' s.t. $\eta'|_{\mathit{FreeVars}(\pi_l)} \triangleq \eta|_{\mathit{FreeVars}(\pi_l)}$ and (vii) $\eta'|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)} \triangleq \rho''|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)}$. We have (viii) $\rho'' = \eta' \circ \sigma'$.

Let $\pi' \triangleq \sigma'(\pi_r)$, $\phi' \triangleq \phi \wedge \sigma'(\phi_l \wedge \phi_r)$, $\varphi' \triangleq \pi' \wedge \phi'$. Using Remark 7 we obtain $\varphi \Rightarrow_S^{\mathfrak{s}} \varphi'$. Moreover, $\mathit{FreeVars}(\varphi') \subseteq \mathit{Var}^b$ since $\mathit{FreeVars}(\varphi') = \mathit{FreeVars}(\sigma'(\pi_r)) \cup \mathit{FreeVars}(\phi) \cup \mathit{FreeVars}(\sigma'(\phi_l)) \cup \mathit{FreeVars}(\sigma'(\phi_l))$, and σ' maps $\mathit{FreeVars}(\pi_l)$ to terms over $\mathit{FreeVars}(\pi) \subseteq \mathit{Var}^b$ and each of the sets $\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l)$, $\mathit{FreeVars}(\phi_r) \setminus \mathit{FreeVars}(\pi_l)$, which are subsets of $\mathit{FreeVars}(\varphi_r) \setminus \mathit{FreeVars}(\varphi_l) \subseteq \mathit{Var}^b$, to the identity. Note that Remark 6 was used in the above reasoning.

There remains to find ρ' with $\rho'|_{\mathit{FreeVars}(\varphi)} = \rho|_{\mathit{FreeVars}(\varphi)}$ such that $(\gamma', \rho') \models \varphi'$. We choose $\rho' \triangleq \rho''$, which does satisfy $\rho''|_{\mathit{FreeVars}(\varphi)} = \rho|_{\mathit{FreeVars}(\varphi)}$ by construction. We have:

1. $\mathit{FreeVars}(\pi_r) = \mathit{FreeVars}(\pi_l) \cup (\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l))$. Since σ' maps $\mathit{FreeVars}(\pi_l)$ to terms over $\mathit{FreeVars}(\pi)$, and is the identity over variables not in $\mathit{FreeVars}(\pi_l)$ (in particular, over $\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l)$) we get $\mathit{FreeVars}(\sigma'(\pi_r)) \subseteq \mathit{FreeVars}(\pi) \cup (\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l)) \subseteq \mathit{Var} \setminus \mathit{FreeVars}(\pi_l)$. Using (vii) we obtain $\rho''|_{\mathit{FreeVars}(\sigma'(\pi_r))} = \eta'|_{\mathit{FreeVars}(\sigma'(\pi_r))}$, thus, $\rho''(\sigma'(\pi_r)) = \eta'(\sigma'(\pi_r)) = (\eta' \circ \sigma')(\pi_r) = \rho''(\pi_r)$ (the last equality, using (viii)). But from (iii) we have $\gamma' = \rho''(\pi_r)$, thus, $\boxed{\gamma' = \rho''(\sigma'(\pi_r))}$;

2. from (iv) : $\boxed{\rho'' \models \phi}$;

3. from (vii,viii) we get $\rho''|_{\mathit{FreeVars}(\pi)} = (\eta' \circ \sigma')|_{\mathit{FreeVars}(\pi)} = \eta'|_{\mathit{FreeVars}(\pi)}$. By Assumption 1, $\mathit{FreeVars}(\phi_l) \subseteq \mathit{FreeVars}(\pi_l)$, thus, $\mathit{FreeVars}(\sigma'(\phi_l)) \subseteq \mathit{FreeVars}(\sigma'(\pi_l)) \subseteq \mathit{FreeVars}(\pi)$. We obtain $\rho''(\sigma'(\phi_l)) = \eta'(\sigma'(\phi_l)) = (\eta' \circ \sigma')(\phi_l) = \rho''(\phi_l)$. But by (v), $\rho''(\phi_l) = \mathit{true}$ hence, $\boxed{\rho'' \models \sigma'(\phi_l)}$;

4. from (viii) we know $\rho''|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)} = \eta'|_{\mathit{Var} \setminus \mathit{FreeVars}(\pi_l)}$. Next, Assumption 1 ensures $\mathit{FreeVars}(\phi_r) \subseteq \mathit{FreeVars}(\pi_l) \cup (\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l))$ and then $\mathit{FreeVars}(\sigma'(\phi_r)) \subseteq \mathit{FreeVars}(\pi) \cup (\mathit{FreeVars}(\pi_r) \setminus \mathit{FreeVars}(\pi_l)) \subseteq \mathit{Var} \setminus \mathit{FreeVars}(\pi_l)$. We then obtain $\rho''|_{\mathit{FreeVars}(\sigma'(\phi_r))} = \eta'|_{\mathit{FreeVars}(\sigma'(\phi_r))}$, thus, $\rho''(\sigma'(\phi_r)) = \eta'(\sigma'(\phi_r)) = (\eta' \circ \sigma')(\phi_r) = \rho''(\phi_r)$ (the last equality, using (viii)). But by (vi), $\rho''(\phi_r) = \mathit{true}$, hence, $\boxed{\rho'' \models \sigma'(\phi_r)}$.

The boxed conclusions of items 1-4 above imply $(\gamma', \rho'') \models \sigma'(\pi_r) \wedge \phi \wedge \sigma'(\varphi_l \wedge \varphi_r)$, i.e., $(\gamma', \rho'') \models \varphi'$, which concludes the proof of the lemma. $\square$

There remains to prove that general version $\Rightarrow_S^{\mathfrak{s}}$, over quantified patterns and generated by rules with quanatified variables in right-hand sides simulates $\Rightarrow_S$.

Lemma 6 (general $\Rightarrow_S^{\mathfrak{s}}$ simulates $\Rightarrow_S$). *For all $\gamma, \gamma' \in M_{Cfg}$, pattern $\varphi \triangleq (\exists X)\pi \wedge \phi$ with $X \cup \mathit{FreeVars}(\varphi) \subseteq \mathit{Var}^b$, and valuation ρ, if $(\gamma, \rho) \models \varphi$ and $\gamma \Rightarrow_\alpha \gamma'$ then there is $\varphi' \triangleq (\exists X, Y)\pi' \wedge \phi'$ with $X \cup Y \cup \mathit{FreeVars}(\varphi') \subseteq \mathit{Var}^b$ such that $\varphi \Rightarrow_\alpha^{\mathfrak{s}} \varphi'$ and $(\gamma', \rho) \models \varphi'$.*

Proof. We take $\varphi' \triangleq (\exists X, Y)(\sigma'(\pi_r) \wedge \phi \wedge \sigma'(\phi_l \wedge \phi_r)$ and by Def. 12, $\varphi \Rightarrow_\alpha^{\mathfrak{s}} \varphi'$ with $\alpha \triangleq \pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r \in \mathcal{S}$.

We first prove $FreeVars(\varphi') \subseteq FreeVars(\varphi)$. We have $\varphi \triangleq (\exists X)\pi \wedge \phi$ and, since $\varphi' = (\exists X, Y)(\sigma'(\pi_r) \wedge \phi \wedge \sigma'(\phi_l \wedge \phi_r))$ we get $FreeVars(\varphi') = (FreeVars(\sigma'(\pi_r)) \cup FreeVars(\phi) \cup FreeVars(\sigma'(\phi_l)) \cup FreeVars(\sigma'(\phi_r))) \setminus (X \cup Y)$.

We have $FreeVars(\sigma'(\pi_r)) = FreeVars(\pi) \cup (FreeVars(\pi_r) \setminus FreeVars(\pi_l))$, and then $FreeVars(\sigma'(\pi_r)) \setminus (X \cup Y) = (FreeVars(\pi) \setminus X) \cup ((FreeVars(\pi_r) \setminus Y) \setminus FreeVars(\pi_l))$. But $(FreeVars(\pi) \setminus X) \subseteq FreeVars(\varphi)$ and by Assumption 1, $(FreeVars(\pi_r) \setminus Y) \setminus FreeVars(\pi_l) = \emptyset$. Hence, $FreeVars(\sigma'(\pi_r)) \subseteq FreeVars(\varphi)$.

Then, $FreeVars(\phi) \setminus (X \cup Y) = FreeVars(\phi) \setminus X \subseteq FreeVars(\varphi)$.

Next, $FreeVars(\sigma'(\phi_l)) \subseteq FreeVars(\sigma'(\pi_l))$ using Assumption 1, and then we obtain $FreeVars(\sigma'(\phi_l)) \subseteq FreeVars(\pi)$, and then $FreeVars(\sigma'(\phi_l)) \setminus (X \cup Y) = FreeVars(\sigma'(\phi_l)) \setminus X = FreeVars(\pi) \setminus X \subseteq FreeVars(\varphi)$.

Finally, by Assumption 1 we have $FreeVars(\phi_r) \subseteq FreeVars(\pi_l) \cup Y$, thus, $FreeVars(\sigma'(\phi_r)) \subseteq FreeVars(\pi) \cup Y$. We then obtain $FreeVars(\sigma'(\phi_r)) \setminus (X \cup Y) = FreeVars(\sigma'(\phi_r)) \setminus Y \subseteq FreeVars(\pi) \subseteq FreeVars(\varphi)$.

The proof of $FreeVars(\varphi') \subseteq FreeVars(\varphi)$ is now complete.

Since we already obtained $\varphi \Rightarrow_\alpha^{\mathfrak{s}} \varphi'$ there only remains to prove $(\gamma', \rho) \models \varphi'$ and $X \cup Y \cup FreeVars(\varphi') \subseteq Var^b$.

We now prove $(\gamma', \rho) \models \varphi'$. We have the hypothesis $\gamma \Rightarrow_\alpha \gamma'$, thus, there exist valuations ρ, ρ' such that $\rho'_{Var \setminus Y} = \rho_{Var \setminus Y}$, $(\gamma, \rho) \models \pi_l \wedge \phi_l$, and $(\gamma', \rho') \models \pi_r \wedge \phi_r$. Since $FreeVars(\pi_l \wedge \phi_l) \subseteq (Var \setminus Y)$ we also have $(\gamma, \rho') \models \pi_l \wedge \phi_l$. Thus, using the unquantified version of the rule α, i.e., $\alpha' \triangleq \pi_l \wedge \phi_l \Rightarrow \pi_r \wedge \phi_r$, we obtain $\gamma \Rightarrow_{\alpha'} \gamma'$.

We note that $FreeVars(\pi \wedge \phi) = FreeVars(\varphi) \cup X \subseteq Var^b$. We can then apply Lemma 5 and obtain $\pi'' \wedge \phi''$ such that $\pi \wedge \phi \Rightarrow_{\alpha'}^{\mathfrak{s}} \pi'' \wedge \phi''$. Using the "unquantified" definition $\Rightarrow_{\alpha'}^{\mathfrak{s}}$ we obtain that $\pi'' \wedge \phi''$ coincides with $\pi' \wedge \phi'$. Lemma 5 also gives us a valuation ρ'', with $\rho''|_{FreeVars(\pi \wedge \phi)} = \rho'|_{FreeVars(\pi \wedge \phi)}$, such that $(\gamma', \rho'') \models \pi' \wedge \phi'$. Thus, using the definition of valuation of quantified patterns, we also obtain $(\gamma', \rho'') \models (\exists X, Y)\pi' \wedge \phi'$, i.e., $(\gamma', \rho'') \models \varphi'$.

We have obtained above $\rho''|_{FreeVars(\pi \wedge \phi)} = \rho'|_{FreeVars(\pi \wedge \phi)}$, hence, we also have $\rho''|_{FreeVars(\pi \wedge \phi) \setminus X} = \rho'|_{FreeVars(\pi \wedge \phi) \setminus X}$, that is $\rho''|_{FreeVars(\varphi)} = \rho'|_{FreeVars(\varphi)}$. Moreover, we know $\rho'_{Var \setminus Y} = \rho_{Var \setminus Y}$, thus, using our lemma's hypothesis that $FreeVars(\varphi)$ is disjoint from Y, we obtain $\rho'|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$. Thus, $\rho''|_{FreeVars(\varphi)} = \rho|_{FreeVars(\varphi)}$, and using $FreeVars(\varphi') \subseteq FreeVars(\varphi)$ (established at the beginning of this proof) and $(\gamma', \rho'') \models \varphi'$, we obtain $(\gamma', \rho) \models \varphi'$.

There only remains to be proved that $X \cup Y \cup FreeVars(\varphi') \subseteq Var^b$. We have $X \cup Y \cup FreeVars(\varphi') = FreeVars(\pi' \wedge \phi')$ and Lemma 5 tells us that $FreeVars(\pi' \wedge \phi') \subseteq Var^b$. This concludes the proof of our lemma. $\square$

Lemma 1 is then a corollary of Lemmas 6 and 2, which concludes the proofs of results from Section 3.

Proofs of Results from Section 4

Proposition 1 (Equivalent Conditions for Terminal Formula Validity)
Consider a terminal formula $\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$. Then $\mathcal{S} \models \pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$ iff

1. $(\exists Y)\phi'$ is invariant at $(\exists Y)\pi'$ starting from $\pi\wedge\phi$, and
2. $(\exists Y)\pi'$ captures all terminal configurations for $\pi\wedge\phi$.

Proof. ($\Leftarrow$) This implication is a simple consequence of Definitions 7 and 8. To prove $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ we consider an arbitrary pair (γ,ρ) such that $(\gamma,\rho) \models \varphi$ and an arbitrary complete path $\gamma \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma'$. Since $(\exists Y)\pi'$ captures all terminal configurations for $\pi\wedge\phi$, there exists ρ' with $\rho'|_{Var\setminus Y} = \rho|_{Var\setminus Y}$ such that $\gamma' = \rho'(\pi')$.

But since $(\exists Y)\phi'$ is invariant at $(\exists Y)\pi'$ starting from $\pi\wedge\phi$, we obtain $\rho' \models \phi'$.

Thus, for an arbitrary pair (γ,ρ) such that $(\gamma,\rho) \models \varphi$ and an arbitrary complete path $\gamma \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma'$, we have obtained ρ' with $\rho'|_{Var\setminus Y} = \rho|_{Var\setminus Y}$ such that $(\gamma',\rho') \models \pi'\wedge\phi'$, which is $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\forall Y)\pi'\wedge\phi'$ according to Definition 13.

($\Rightarrow$) We have to show that $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ implies (i) $(\exists Y)\phi'$ is invariant at $(\exists Y)\pi'$ starting from $\pi\wedge\phi$ and (ii) $(\exists Y)\pi'$ captures all terminal configurations for $\pi\wedge\phi$.

We first prove (i). Consider an arbitrary path $\gamma_0 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n$ and valuations ρ,ρ' such that $(\gamma_0,\rho) \models \pi\wedge\phi$, $\rho'|_{Var\setminus Y} = \rho|_{Var\setminus Y}$, and $\gamma_n = \rho'(\pi')$.

Since π' is terminal, γ_n is terminal and thus the path $\gamma_0 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n$ is complete. From $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ we obtain that there exists $0 \leq i \leq n$ and a valuation ρ'' with $\rho''|_{Var\setminus Y} = \rho|_{Var\setminus Y}$ such that $(\gamma_i,\rho'') \models \pi'\wedge\phi'$. Thus, $\gamma_i = \rho''(\pi')$ is terminal as well, and thus $\gamma_i = \gamma_n$, since there cannot be two terminal configurations on one path. Thus, $\gamma_n = \rho''(\pi')$ and $\rho'' \models \phi'$. We now show $\rho' \models \varphi'$.

By Assumption 1, $FreeVars(\phi') \subseteq FreeVars(\pi) \cup FreeVars(\pi')$. Consider then any variable $x \in FreeVars(\phi')$.

- if $x \in FreeVars(\pi)$: from $\rho''|_{FreeVars(\pi\wedge\phi)} = \rho|_{FreeVars(\pi\wedge\phi)}$, $\rho'|_{FreeVars(\pi\wedge\phi)} = \rho|_{FreeVars(\pi\wedge\phi)}$ we obtain $\rho''|_{FreeVars(\pi)} = \rho|_{FreeVars(\pi)} = \rho'|_{FreeVars(\pi)}$, meaning that $\rho''(x) = \rho'(x)$;
- if $x \in FreeVars(\pi')$: By Remark 1, from $\rho''(\pi') = \rho'(\pi')(= \gamma_n)$ we obtain $\rho'|_{FreeVars(\pi')} = \rho''|_{FreeVars(\pi')}$ and thus $\rho''(x) = \rho'(x)$.

Thus, $\rho''|_{FreeVars(\phi')} = \rho'|_{FreeVars(\phi')}$, and from $\rho'' \models \phi'$ we obtain $\rho' \models \phi'$.

Recapitulating, for the arbitrary path $\gamma_0 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n$ and valuations ρ,ρ' such that $(\gamma_0,\rho) \models \pi\wedge\phi$, $\rho'|_{Var\setminus Y} = \rho|_{Var\setminus Y}$, and $\gamma_n = \rho'(\pi')$, $\rho' \models \phi'$ holds. This is just $(\exists Y)\phi'$ *is invariant at* $(\exists Y)\pi'$ *starting from* $\pi\wedge\phi$, which proves (i).

There remains to prove (ii). Consider now an arbitrary *terminal* path $\gamma_0 \Rightarrow_{\mathcal{S}} \cdots \Rightarrow_{\mathcal{S}} \gamma_n$ and a valuation ρ such that $(\gamma_0,\rho) \models \pi\wedge\phi$. Since $\mathcal{S} \models \pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$ we obtain $(\gamma_n,\rho) \models (\exists Y)\pi'\wedge\phi'$. In particular, $(\gamma_n,\rho) \models (\exists Y)\pi'$, i.e., $(\exists Y)\pi'$ captures all terminal configurations for $\pi\wedge\phi$, which proves (ii), the ($\Rightarrow$) implication, and the proposition. $\square$

Lemma 7 (Simulation by Graph). *Consider any complete path* $\tau = \gamma_0 \Rightarrow_{\alpha_1} \cdots \Rightarrow_{\alpha_n} \gamma_n$ *with* $\alpha_1,\ldots,\alpha_n \in \mathcal{S}_0$, *s.t.* $(\gamma_0,\rho) \models \pi\wedge\phi$. *Then, there is a path in the*

graph constructed by the procedure in Figure 2 which simulates a subsequence of τ, i.e., there exist $k \geq 0$, a subsequence ($0 = i_0 < \cdots < i_k = n$), and a path $(\pi \wedge \phi =)\varphi_0 \xrightarrow{\alpha_{i_1}} \cdots \xrightarrow{\alpha_{i_k}} \varphi_n$ in the graph s.t. $(\gamma_{i_j}, \rho) \models \varphi_{i_j}$ for $j = 0 \ldots k$.

Proof. We show how to inductively construct the sequence of indices ($0 = i_0 < \cdots < i_k = n$) and the corresponding path in the graph.

The first index is (by definition) $i_0 = 0$. In this case the path in the graph reduces to the sole node $\pi_{i_0} \wedge \phi_{i_0} = \pi_0 \wedge \phi_0 = \pi \wedge \phi$, and the valuation ρ together with $\gamma_{i_0} = \gamma_0$ obviously satisfies $(\gamma_{i_0}, \rho) \models \pi_{i_0} \wedge \phi_{i_0}$.

Assume now that we have built the subsequence up to some index $0 \leq i_m \leq n$. Thus, we have built the sequence ($0 = i_0 < \cdots < i_m$) and the path $(\pi \wedge \phi =)\varphi_0 \xrightarrow{\alpha_{i_1}} \cdots \xrightarrow{\alpha_{i_m}} \varphi_m$ satisfying the conclusions of the lemma. If $i_m = n$ the conclusion of the lemma holds directly so we can assume $i_m < n$ in the rest of the proof.

We show how to extend the sequence of indices, and the path in the graph for satisfying the lemma's conclusions.

We know that $\varphi_{i_m} \triangleq (\exists Z)\pi_{i_m} \wedge \phi_{i_m}$ is a node in the graph and that $(\gamma_{i_m}, \rho) \models (\exists Z)\pi_{i_m} \wedge \phi_{i_m}$.

Consider the configuration γ_{i_m} on the sequence τ. Since $i_m < n$ the configuration γ_{i_m} has a successor on τ i.e., there is a rule $\alpha_{i_m} \in \mathcal{S}$ such that $(\gamma_{i_m}, \rho) \models lhs(\alpha_{i_m})$. By Assumption 4 on the relation $<$, there exists a rule $\alpha \triangleq \pi_l \wedge \phi_l \Rightarrow (\exists Y)\pi_r \wedge \phi_r \in min(<)$ such that $(\gamma_{i_m}, \rho) \models \pi_l \wedge \phi_l$. Since $\mathcal{S} = \mathcal{S}_0 \cup \mathcal{G}$, we distinguish two cases:

- $\alpha \in \mathcal{G}$, and we obtain $\mathcal{S}_0 \models \alpha$ (since, cf. Assumption 4, $\mathcal{S}_0 \models \mathcal{G}$), or
- $\alpha \in \mathcal{S}_0$, and thus, $\mathcal{S}_0 \models \alpha$.

Next, using the definition of RL formula validity, on the (complete) path $\gamma_{i_m} \cdots \gamma_n$ (which is a nonempty suffix of τ) there exists an index, say, $i_{m+1} \leq n$, and ρ' such that $(\gamma_{i_{m+1}}, \rho') \models \pi_r \wedge \phi_r$. Moreover, $i_{m+1} > i_m$ since $\gamma_{i_m} \in [\![\pi_l \wedge \phi_l]\!]$, which by Assumption 5.2 is disjunct from $[\![\pi_r \wedge \phi_r]\!]$ that contains $\gamma_{i_{m+1}}$.

By Remark 5 we have $\gamma_{i_m} \Rightarrow_\alpha \gamma_{i_{m+1}}$, thus, using Lemma 1 and Definition 6 of derivatives, $(\exists Z)\pi_{i_m} \wedge \phi_{i_m} \Rightarrow_{\alpha^s} \Delta_\alpha((\exists Z)\pi_{i_m} \wedge \phi_{i_m})$, and $(\gamma_{i_{m+1}}, \rho) \models \Delta_\alpha((\exists Z)\pi_{i_m} \wedge \phi_{i_m})$. We take i_{m+1} to be the next element of the sequence ($0 = i_0 < \cdots < i_m$), and extend the path $\pi_0 \wedge \phi_0 \xrightarrow{\alpha_{i_1}} \cdots \xrightarrow{\alpha_{i_m}} (\exists Z)\pi_{i_m} \wedge \phi_{i_m}$ with the transition $(\exists Z)\pi_{i_m} \wedge \phi_{i_m} \xrightarrow{\alpha} \varphi_{i_{m+1}}$, where $\varphi_{i_{m+1}} \triangleq \Delta_\alpha((\exists Z)\pi_{i_m} \wedge \phi_{i_m})$ if $inclusion(\Delta_\alpha((\exists Z)\pi_{i_m} \wedge \phi_{i_m}), \pi \wedge \phi) = false$, and $\varphi_{i_{m+1}} \triangleq \pi \wedge \phi$ otherwise. From $(\gamma_{i_{m+1}}, \rho) \models \Delta_\alpha((\exists Z)\pi_{i_m} \wedge \phi_{i_m})$ we obtain that $(\gamma_{i_{m+1}}, \rho) \models \varphi_{i_{m+1}}$ in both cases above.

Thus, we have obtained the next index i_{m+1} in the sequence ($0 = i_0 < \cdots < i_k = n$) of indices in τ, the next node $\varphi_{i_{m+1}}$ in the path $(\pi \wedge \phi =)\varphi_0 \xrightarrow{\alpha_{i_1}} \cdots \xrightarrow{\alpha_{i_n}} \varphi_n$ in the graph satisfying all the lemma's conclusions. This completes the inductive construction of the elements whose existence is stated by the lemma. $\square$

We now prove the main result of the paper (Theorem 1). Since the theorem consists in two implications we have naturally decomposed it into two lemmas.

Lemma 8. *Consider a terminal* RL *formula* $\pi \wedge \phi \Rightarrow (\exists Y)\pi' \wedge \phi'$. *If the procedure in Figure 2 terminates with Failure $=$ false, then $(\exists Y)\phi'$ is an invariant at $(\exists Y)\pi'$ starting from $\pi \wedge \phi$.*

Proof. To prove that $(\exists Y)\phi'$ is an invariant at $(\exists Y)\pi'$ starting from $\pi \wedge \phi$ we show that, for any path $\tau = \gamma_0 \Rightarrow_{\alpha_1} \cdots \Rightarrow_{\alpha_n} \gamma_n$, any valuation ρ such that $(\gamma_0, \rho) \models \pi \wedge \phi$, and any valuation ρ' with $\rho'|_{Var \setminus Y} = \rho|_{Var \setminus Y}$ and $\gamma_n = \rho'(\pi')$, it holds that $\rho' \models \phi'$.

Consider then an arbitrary path τ and valuations ρ, ρ' satisfying the above properties; we need to show $\rho' \models \phi'$. Since π' is terminal we know that $\gamma_n = \rho'(\pi')$ is terminal as well.

Next, by Lemma 7, there exists $k \geq 0$, a subsequence $(0 = i_0 < \cdots < i_k = n)$ of indices in τ, a path $(\pi \wedge \phi =)\varphi_0 \xrightarrow{\alpha_{i_1}} \cdots \xrightarrow{\alpha_{i_k}} (\exists U)\pi_n \wedge \phi_n$ in the graph, with, in particular, $(\gamma_n, \rho) \models (\exists U)\pi_n \wedge \phi_n$.

Thus, $\gamma_n = \rho''(\pi_n)$ for some valuation ρ'' such that $\rho''|_{Var \setminus U} = \rho|_{Var \setminus U}$, in particular, and $\rho''|_{Free Vars((\exists U)\pi_n \wedge \phi_n)} = \rho|_{Free Vars((\exists U)\pi_n \wedge \phi_n)}$.

At the beginning of the proof of Lemma 6 we establish that the sets of free variables of patterns along a symbolic paths cannot increase. Thus, we also have $Free Vars((\exists U)\pi_n \wedge \phi_n) \subseteq Free Vars(\pi \wedge \phi)$, and $\rho''|_{Free Vars(\pi \wedge \phi)} = \rho|_{Free Vars(\pi \wedge \phi)}$.

We let $\varrho : Var \to M$ be a valuation s.t.:

- $\varrho|_{Free Vars(\pi') \setminus Free Vars(\pi \wedge \phi)} = \rho'|_{Free Vars(\pi') \setminus Free Vars(\pi \wedge \phi)}$,
- $\varrho|_{Free Vars(\pi \wedge \phi)} = \rho''|_{Free Vars(\pi \wedge \phi)} = \rho'|_{Free Vars(\pi \wedge \phi)}$, and
- $\varrho|_{Free Vars(\pi_n) \setminus Free Vars(\pi \wedge \phi)} = \rho''|_{Free Vars(\pi_n) \setminus Free Vars(\pi \wedge \phi)}$.

then $\varrho(\pi_n) = \rho''(\pi_n) = \gamma_n = \rho'(\pi') = \varrho(\pi')$. From $\varrho(\pi_n) = \varrho(\pi')$ and Lemma 4 (unification by matching) we obtain $match_\cong(\pi_n, \pi') \neq \emptyset$. Thus, in our procedure, after $(\exists U)\pi_n \wedge \phi_n$ is selected from *New*, the test **if** $match_\cong(\pi_n, \pi') = \emptyset$ at line 3 leads us into the **else** branch, i.e., line 8. And since we assumed the procedure terminated with *Failure = false* it must be that $inclusion((\exists U)\pi_n \wedge \phi_n, (\exists Y)\pi' \wedge \phi')$ on line 10 evaluates to *true*. Thus, by Definition 10 of inclusion test, $(\gamma_n, \rho) \models (\exists Y)\pi' \wedge \phi'$, thus, there exists ρ'' with $\rho''|_{Var \setminus Y} = \rho|_{Var \setminus Y}$ such that $(\gamma_n, \rho'') \models \pi' \wedge \phi'$, in particular, $\rho'' \models \phi'$.

There remains to prove that $\rho''|_{Free Vars(\phi')} = \rho'|_{Free Vars(\phi')}$. By Assumption 1, $Free Vars(\phi') \subseteq Free Vars(\pi) \cup Free Vars(\pi')$. Let $x \in Free Vars(\pi) \cup Free Vars(\pi')$; we prove $\rho'(x) = \rho''(x)$.

The first case is $x \in Free Vars(\pi')$. From $\gamma_n = \rho'(\pi') = \rho''(\pi')$, by Assumption 1 we know that $\rho'(x) = \rho''(x)$, which settles this case.

In the second case, $x \in Free Vars(\pi)$. Since variables in Y are quantified, we can assume that $Free Vars(\pi) \cap Y = \emptyset$. On $Var \setminus Y$ valuations ρ, ρ', ρ'' coincide. Thus, $\rho'(x) = \rho''(x)$, which settles this case as well. From $\rho''|_{Free Vars(\phi')} = \rho'|_{Free Vars(\phi')}$ and $\rho'' \models \phi'$ we obtain our conclusion: $\rho' \models \phi'$. $\qquad \square$

This proves the first implication in Theorem 1. We now deal with the second one.

Lemma 9. *Consider a terminal* RL *formula* $\pi\wedge\phi \Rightarrow (\exists Y)\pi'\wedge\phi'$. *If the procedure Figure 2) terminates with Failure = false, then* $(\exists Y)\pi'$ *captures all terminal configurations starting from* $\pi\wedge\phi$.

Proof. We need to show that for all (γ_0,ρ) such that $(\gamma,\rho) \models \pi\wedge\phi$, and all complete paths $\tau = \gamma_0 \Rightarrow_{\alpha_1} \cdots \Rightarrow_{\alpha_n} \gamma_n$, $(\gamma_n,\rho) \models (\exists Y)\pi'$.

By Lemma 7, there exists $k \geq 0$, a subsequence $(0 = i_0 < \cdots < i_k = n)$ of indices in τ, a path $(\pi\wedge\phi =)\pi_0\wedge\phi_0 \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_n} (\exists Z)\pi_n\wedge\phi_n$ in the graph, such that $(\gamma_n,\rho) \models (\exists Z)\pi_n\wedge\phi_n$.

We first prove by contradiction that $match_\cong(\pi_n,\pi') \neq \emptyset$. Thus, assume $match_\cong(\pi_n,\pi') = \emptyset$. Then, when after choosing the node $(\exists Z)\pi_n\wedge\phi_n$ in *New* the procedure enters the **then** branch after the test at line 3, and proceeds with performing the test at line 4. Since we assumed the procedure terminates with *Failure = false*, that test is passed, i.e., $inclusion((\exists Z)\pi_n\wedge\phi_n, lhs(\alpha)) = true$ for some $\alpha \triangleq \pi_l\wedge\phi_l \Rightarrow (\exists W)\pi_r\wedge\phi_r \in min(<)$. Since $(\gamma_n,\rho) \models (\exists Z)\pi_n\wedge\phi_n$, by Definition 10 of inclusion test $(\gamma_n,\rho) \models \pi_l\wedge\phi_l$.

But by Assumption 5, all rules in $\mathcal{S}$ are weakly well-defined, thus, there exist γ_{n+1} such that $(\gamma_{n+1},\rho) \models (\exists W)\pi_r\wedge\phi_r$, thus, $\gamma_n \Rightarrow_\alpha \gamma_{n+1}$, i.e., γ_n is not terminal, which contradicts the hypothesis that $\tau = \gamma_0 \Rightarrow_{\alpha_1} \cdots \Rightarrow_{\alpha_n} \gamma_n$ is a complete path. Thus, $match_\cong(\pi_n,\pi') \neq \emptyset$.

Hence, the test $match_\cong(\pi_n,\pi') \neq \emptyset$ at line 3 leads into the **else** branch at line 10. And since the procedure terminates with *Failure = false*, we obtain $inclusion((\exists Z)\pi_n\wedge\phi_n, (\exists Y)\pi'\wedge\phi') = true$, which means, by Definition 10 of the inclusion test, that $(\gamma_n,\rho) \models (\exists Y)\pi'\wedge\phi'$, in particular $(\gamma_n,\rho) \models (\exists Y)\pi'$. $\square$

Theorem 1 is now a corrolary of Lemmas 8 and 9, which concludes this appendix.