ALL-PATH REACHABILITY LOGIC

ANDREI ȘTEFĂNESCU [a,*], STEFAN CIOBÂCĂ [b], RADU MEREUTA [b,c], BRANDON M. MOORE [c,*], GRIGORE ROȘU [d,c], AND TRAIAN FLORIN ȘERBĂNUȚĂ [c,e]

[a] Galois, Inc., USA
   e-mail address: andrei@galois.com

[b] Alexandru Ioan Cuza University, Romania
   e-mail address: {stefan.ciobaca,radu.mereuta}@info.uaic.ro

[c] Runtime Verification, Inc., USA
   e-mail address: brandon.moore@runtimeverification.com

[d] University of Illinois at Urbana-Champaign, USA
   e-mail address: grosu@illinois.edu

[e] University of Bucharest, Romania
   e-mail address: traian.serbanuta@fmi.unibuc.ro

Abstract. This paper presents a language-independent proof system for reachability properties of programs written in non-deterministic (e.g., concurrent) languages, referred to as *all-path reachability logic*. It derives partial-correctness properties with all-path semantics (a state satisfying a given precondition reaches states satisfying a given postcondition on all terminating execution paths). The proof system takes as axioms any unconditional operational semantics, and is sound (partially correct) and (relatively) complete, independent of the object language. The soundness has also been mechanized in Coq. This approach is implemented in a tool for semantics-based verification as part of the 𝕂 framework (http://kframework.org).

1. Introduction

Operational semantics are easy to define and understand. Giving a language an operational semantics can be regarded as "implementing" a formal interpreter. Operational semantics require little formal training, scale up well and, being executable, can be tested. Thus, operational semantics are typically used as trusted reference models for the defined languages. Despite these advantages, operational semantics are rarely used directly for program verification (i.e., verifying properties of a given program, rather than performing meta-reasoning about a given language), because such proofs tend to be low-level and tedious, as they involve formalizing and working directly with the corresponding transition system. Hoare or dynamic logics allow higher level reasoning at the cost of (re)defining the language as a set of abstract proof rules, which are harder to understand and trust. The state-of-the-art in mechanical

---

* Most of the present work was done while still being PhD students at the University of Illinois at Urbana-Champaign.

program verification is to develop and prove such language-specific proof systems sound w.r.t to a trusted operational semantics [1–3], but that needs to be done for each language separately.

Defining more semantics for the same language and proving the soundness of one semantics in terms of another are highly uneconomical tasks where real programming languages are concerned, often taking several years to complete. Ideally, we would like to have only one semantics for a language, together with a generic theory and a set of generic tools and techniques allowing us to get all the benefits of any other semantics without paying the price of defining other semantics. Recent work [4–7] shows this *is possible*, by proposing a *language-independent proof system* which derives program properties directly from an operational semantics, *at the same proof granularity and compositionality* as a language-specific axiomatic semantics. Specifically, it introduces *(one-path) reachability rules*, which generalize both operational semantics reduction rules and Hoare triples, and gives a proof system which derives new reachability rules (program properties) from a set of given reachability rules (the language semantics).

However, the existing proof system has a major limitation: it only derives reachability rules with a *one-path* semantics, that is, it guarantees a program property holds on one but not necessarily all execution paths, which suffices for deterministic languages but not for non-deterministic (concurrent) languages. Here we remove this limitation, proposing the first generic all-path reachability proof system for program verification.

Using *matching logic* [8] as a configuration specification formalism (Section 2), where a pattern φ specifies all program configurations that match it, we first introduce the novel notion of an *all-path reachability rule* $\varphi \Rightarrow^\forall \varphi'$ (Section 3), where φ and φ' are matching logic patterns. Rule $\varphi \Rightarrow^\forall \varphi'$ is valid iff any program configuration satisfying φ reaches, on any complete execution path, some configuration satisfying φ'. This subsumes Hoare-style partial correctness in non-deterministic languages. We then present a proof system for deriving an all-path reachability rule $\varphi \Rightarrow^\forall \varphi'$ from a set $\mathcal{S}$ of semantics rules (Section 4). $\mathcal{S}$ consists of reduction rules $\varphi_l \Rightarrow^\exists \varphi_r$, where φ_l and φ_r are simple patterns as encountered in operational semantics (Section 5), which can be non-deterministic. The proof system derives more general sequents "$\mathcal{S}, \mathcal{A} \vdash_\mathcal{C} \varphi \Rightarrow^\forall \varphi'$", with $\mathcal{A}$ and $\mathcal{C}$ two sets of reachability rules. Intuitively, $\mathcal{A}$'s rules (*axioms*) are already established valid, and thus can be immediately used. Those in $\mathcal{C}$ (*circularities*) are only claimed valid, and can be used only after taking execution steps based on the rules in $\mathcal{S}$ or $\mathcal{A}$. The most important proof rules are

Step :
$$\models \varphi \rightarrow \bigvee_{\varphi_l \Rightarrow^\exists \varphi_r \,\in\, \mathcal{S}} \exists FreeVars(\varphi_l)\ \varphi_l$$
$$\frac{\models \exists c\ (\varphi[c/\Box] \wedge \varphi_l[c/\Box]) \wedge \varphi_r \rightarrow \varphi' \qquad \text{for each}\ \ \varphi_l \Rightarrow^\exists \varphi_r\ \in\ \mathcal{S}}{\mathcal{S}, \mathcal{A}\ \vdash_\mathcal{C} \varphi \Rightarrow^\forall \varphi'}$$

Circularity :
$$\frac{\mathcal{S}, \mathcal{A}\ \vdash_{\mathcal{C}\cup\{\varphi\Rightarrow^\forall\varphi'\}} \varphi \Rightarrow^\forall \varphi'}{\mathcal{S}, \mathcal{A}\ \vdash_\mathcal{C} \varphi \Rightarrow^\forall \varphi'}$$

STEP is the key proof rule which deals with non-determinism: it derives a sequent where φ reaches φ' in one step on all paths. The first premise ensures that any configuration satisfying φ has successors, the second that all successors satisfy φ' ($\Box$ is the configuration placeholder). CIRCULARITY adds the current goal to $\mathcal{C}$ at any point in a proof, and generalizes the various

language-specific axiomatic semantics invariant rules in a language-independent way (this form was introduced in [4]).

We illustrate on examples how our proof system enables state exploration (similar to symbolic model-checking), and verification of program properties (Section 5). We show that our proof system is sound (Section 7) and relatively complete (Section 8). We describe our implementation of the proof system as part of the $\mathbb{K}$ framework [9] (Section 6).

Contributions. This paper makes the following specific contributions:

(1) A language-independent proof system for deriving all-path reachability properties, with (complete) proofs of its soundness and relative completeness; the soundness result has also been mechanized in Coq, to serve as a foundation for certifiable verification.
(2) An implementation of it as part of the $\mathbb{K}$ framework (`http://kframework.org`).

An earlier, shorter version of this paper appeared at RTA-TCLA 2014 [10]. The main differences are the inclusion of proofs for soundness (Section 7), and relative completeness (Section 8).

2. Topmost Matching Logic

Matching logic [11,12] is a first-order logic (FOL) variant designed for specifying and reasoning about structure by means of patterns and pattern matching. In this section, we briefly recall topmost matching logic [8], a subset of the full matching logic theory described in [11,12]. We employ topmost matching logic to specify and reason about arbitrary program configurations. For simplicity, we use "matching logic" instead of "topmost matching logic" in this paper. A matching logic formula, called a *pattern*, is a first-order logic (FOL) formula with special predicates, called basic patterns. A *basic pattern* is a configuration term with variables. Intuitively, a pattern specifies both structural and logical constraints: a configuration satisfies the pattern iff it matches the structure (basic patterns) and satisfies the constraints.

Matching logic is parametric in a signature and a model of configurations, making it a prime candidate for expressing state properties in a language-independent verification framework. The configuration signature can be as simple as that of IMP (Fig. 2), or as complex as that of the C language [13,14] (with more than 100 semantic components).

We use basic concepts from multi-sorted first-order logic. Given a *signature* Σ which specifies the sorts and arities of the function symbols (constructors or operators) used in configurations, let $T_\Sigma(\mathit{Var})$ denote the *free Σ-algebra* of terms with variables in Var. $T_{\Sigma,s}(\mathit{Var})$ is the set of Σ-terms of sort s. A valuation $\rho\colon \mathit{Var}\to\mathcal{T}$ with $\mathcal{T}$ a Σ-algebra extends uniquely to a (homonymous) Σ-*algebra morphism* $\rho\colon T_\Sigma(\mathit{Var})\to\mathcal{T}$. Many mathematical structures needed for language semantics have been defined as Σ-algebras, including: boolean algebras, natural/integer/rational numbers, lists, sets, bags (or multisets), maps (e.g., for states, heaps), trees, queues, stacks, etc.

Let us fix the following: (1) an algebraic signature Σ, associated to some desired configuration syntax, with a distinguished sort *Cfg*, (2) a sort-wise infinite set *Var* of variables, and (3) a Σ-algebra $\mathcal{T}$, the *configuration model*, which may but need not be a term algebra. As usual, $\mathcal{T}_{Cfg}$ denotes the elements of $\mathcal{T}$ of sort *Cfg*, called *configurations*.

Definition 2.1 [8]**.** A matching logic formula, or a **pattern**, is a first-order logic (FOL) formula which additionally allows terms in $T_{\Sigma,Cfg}(\mathit{Var})$, called **basic patterns**, as predicates. A pattern is **structureless** if it contains no basic patterns.

We define satisfaction $(\gamma, \rho) \models \varphi$ over configurations $\gamma \in \mathcal{T}_{Cfg}$, valuations $\rho \colon Var \to \mathcal{T}$ and patterns φ as follows (among the FOL constructs, we only show $\exists$):

$(\gamma, \rho) \models \exists X \; \varphi$ iff $(\gamma, \rho') \models \varphi$ for some $\rho' \colon Var \to \mathcal{T}$ with $\rho'(y) = \rho(y)$ for all $y \in Var \backslash X$

$(\gamma, \rho) \models \pi$ iff $\gamma = \rho(\pi)$ where $\pi \in T_{\Sigma, Cfg}(Var)$

We write $\models \varphi$ when $(\gamma, \rho) \models \varphi$ for all $\gamma \in \mathcal{T}_{Cfg}$ and all $\rho \colon Var \to \mathcal{T}$.

A basic pattern π is satisfied by all the configurations γ that *match* it; in $(\gamma, \rho) \models \pi$ the ρ can be thought of as the "witness" of the matching, and can be further constrained in a pattern. For instance, the pattern from Section 5

$$\langle \mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1} \; || \; \mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1}, \; \mathtt{x} \mapsto n \rangle \wedge (n = 0 \vee n = 1)$$

is matched by the configurations with code "$\mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1} \; || \; \mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1}$" and state mapping program variable $\mathtt{x}$ into integer n with n being either 0 or 1. We use *italic* for mathematical variables in *Var* and $\mathtt{typewriter}$ for program variables (program variables are represented in matching logic as constants of sort *PVar*, see Section 5).

The "topmost" matching logic variant that we use in this paper is a fragment of the general matching logic approach proposed in [11, 12], where terms are regarded as predicates only at the top level, where they have the sort *Cfg*. The general matching logic approach in [11, 12] allows the distinctive terms-as-predicates view as well as the mix of logical and structural constraints at all levels, for any terms of any sorts. As an example, the pattern above may be equivalently written as

$$\langle \mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1} \; || \; \mathtt{x}\!:=\!\mathtt{x}\!+\!\mathtt{1}, \; \mathtt{x} \mapsto n \wedge (n = 0 \vee n = 1) \rangle$$

thus localizing the logical constraints to the actual place in the configuration where they matter. We remind the reader that in this paper we limit ourselves to "topmost" matching logic, and that we take the freedom to drop the "topmost" and just call it "matching logic".

Next, we recall how matching logic formulae can be translated into FOL formulae, so that its satisfaction becomes FOL satisfaction in the model of configurations, $\mathcal{T}$. Then, we can use conventional theorem provers or proof assistants for pattern reasoning.

Definition 2.2 [8]**.** Let $\square$ be a fresh *Cfg* variable. For a pattern φ, let $\varphi^{\square}$ be the FOL formula formed from φ by replacing basic patterns $\pi \in T_{\Sigma, Cfg}(Var)$ with equalities $\square = \pi$. If $\rho \colon Var \to \mathcal{T}$ and $\gamma \in \mathcal{T}_{Cfg}$ then let the valuation $\rho^{\gamma} \colon Var \cup \{\square\} \to \mathcal{T}$ be such that $\rho^{\gamma}(x) = \rho(x)$ for $x \in Var$ and $\rho^{\gamma}(\square) = \gamma$.

With the notation in Definition 2.2, $(\gamma, \rho) \models \varphi$ iff $\rho^{\gamma} \models \varphi^{\square}$, and $\models \varphi$ iff $\mathcal{T} \models \varphi^{\square}$. Thus, matching logic is a methodological fragment of the FOL theory of $\mathcal{T}$. We drop $\square$ from $\varphi^{\square}$ when it is clear in context that we mean the FOL formula instead of the matching logic pattern. It is often technically convenient to eliminate $\square$ from φ, by replacing $\square$ with a *Cfg* variable c and using $\varphi[c/\square]$ instead of φ. We use the FOL representation in the STEP proof rule in Fig. 1, and to establish relative completeness in Section 8.

3. Specifying Reachability

In this section we define one-path and all-path reachability. We begin by recalling some matching logic reachability [6] notions that we need for specifying reachability.

Definition 3.1 [6]. A (one-path) **reachability rule** is a pair $\varphi \Rightarrow^\exists \varphi'$, where φ and φ' are patterns (which can have free variables). Rule $\varphi \Rightarrow^\exists \varphi'$ is **weakly well-defined** iff for any $\gamma \in \mathcal{T}_{Cfg}$ and $\rho : Var \to \mathcal{T}$ with $(\gamma, \rho) \models \varphi$, there exists $\gamma' \in \mathcal{T}_{Cfg}$ with $(\gamma', \rho) \models \varphi'$. A **reachability system** is a set of reachability rules. Reachability system $\mathcal{S}$ is **weakly well-defined** iff each rule is weakly well-defined. $\mathcal{S}$ induces a **transition system** $(\mathcal{T}, \Rightarrow^{\mathcal{T}}_{\mathcal{S}})$ on the configuration model: $\gamma \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma'$ for $\gamma, \gamma' \in \mathcal{T}_{Cfg}$ iff there is some rule $\varphi \Rightarrow^\exists \varphi'$ in $\mathcal{S}$ and some valuation $\rho : Var \to \mathcal{T}$ with $(\gamma, \rho) \models \varphi$ and $(\gamma', \rho) \models \varphi'$. A $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$**-path** is a finite sequence $\gamma_0 \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma_1 \Rightarrow^{\mathcal{T}}_{\mathcal{S}} ... \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma_n$ with $\gamma_0,...,\gamma_n \in \mathcal{T}_{Cfg}$. A $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$-path is **complete** iff it is not a strict prefix of any other $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$-path.

We assume an operational semantics is a set of (unconditional) reduction rules "$l \Rightarrow^\exists r$ if b", where $l, r \in T_{\Sigma, Cfg}(Var)$ are program configurations with variables and $b \in T_{\Sigma, Bool}(Var)$ is a condition constraining the variables of l, r. Styles of operational semantics using only such (unconditional) rules include evaluation contexts [15], the chemical abstract machine [16] and $\mathbb{K}$ [9] (see Section 5 for an evaluation contexts semantics). Several large languages have been given semantics in such styles, including C [13] (more than 3000 rules) and R5RS Scheme [17]. The reachability proof system works with any set of rules of this form, being agnostic to the particular style of semantics.

Such a rule "$l \Rightarrow^\exists r$ if b" states that a ground configuration γ which is an instance of l and satisfies the condition b reduces to an instance γ' of r. Matching logic can express terms with constraints: $l \wedge b$ is satisfied by exactly the γ above. Thus, we can regard such a semantics as a particular weakly well-defined reachability system $\mathcal{S}$ with rules of the form "$l \wedge b \Rightarrow^\exists r$". The weakly well-defined condition on $\mathcal{S}$ guarantees that if γ matches the left-hand-side of a rule in $\mathcal{S}$, then the respective rule induces an outgoing transition from γ. The transition system induced by $\mathcal{S}$ describes precisely the behavior of any program in any given state. In [4–6] we show that reachability rules capture one-path reachability properties and Hoare triples for deterministic languages.

Formally, let us fix an operational semantics given as a reachability system $\mathcal{S}$. Then, we can specify reachability in the transition system induced by $\mathcal{S}$, both all-path and one-path, as follows:

Definition 3.2. An **all-path reachability rule** is a pair $\varphi \Rightarrow^\forall \varphi'$ of patterns φ and φ'.

An all-path reachability rule $\varphi \Rightarrow^\forall \varphi'$ is satisfied, $\mathcal{S} \models \varphi \Rightarrow^\forall \varphi'$, iff for all complete $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$-paths τ starting with $\gamma \in \mathcal{T}_{Cfg}$ and for all $\rho : Var \to \mathcal{T}$ such that $(\gamma, \rho) \models \varphi$, there exists some $\gamma' \in \tau$ such that $(\gamma', \rho) \models \varphi'$.

A one-path reachability rule $\varphi \Rightarrow^\exists \varphi'$ is satisfied, $\mathcal{S} \models \varphi \Rightarrow^\exists \varphi'$, iff for all $\gamma \in \mathcal{T}_{Cfg}$ and $\rho : Var \to \mathcal{T}$ such that $(\gamma, \rho) \models \varphi$, there is either a $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$-path from γ to some γ' such that $(\gamma', \rho) \models \varphi'$, or there is a diverging execution $\gamma \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma_1 \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma_2 \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \cdots$ from γ.

The racing increment example in Section 5 can be specified by

$$\langle \texttt{x:=x+1} \mathbin{||} \texttt{x:=x+1}, \texttt{x} \mapsto m \rangle \Rightarrow^\forall \exists n \, (\langle \texttt{skip}, \texttt{x} \mapsto n \rangle \wedge (n = m +_{Int} 1 \vee n = m +_{Int} 2)$$

which states that every terminating execution reaches a state where execution of both threads is complete and the value of $\texttt{x}$ has increased by 1 or 2 (this code has a race).

A Hoare triple describes the resulting state after execution finishes, so it corresponds to a reachability rule where the right side contains no remaining code. However, all-path reachability rules are strictly more expressive than Hoare triples, as they can also specify intermediate configurations (the code in the right-hand-side need not be empty). Reachability

Step :

$$\frac{\models \varphi \to \bigvee_{\varphi_l \Rightarrow^\exists \varphi_r \ \in \ \mathcal{S}} \exists \mathit{FreeVars}(\varphi_l) \ \varphi_l \qquad \models \exists c \ (\varphi[c/\Box] \wedge \varphi_l[c/\Box]) \wedge \varphi_r \to \varphi' \qquad \text{for each} \ \ \varphi_l \Rightarrow^\exists \varphi_r \ \in \ \mathcal{S}}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'}$$

Axiom :

$$\frac{\varphi \Rightarrow^\forall \varphi' \ \in \ \mathcal{A}}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'}$$

Reflexivity :

$$\frac{\cdot}{\mathcal{S}, \mathcal{A} \vdash \varphi \Rightarrow^\forall \varphi}$$

Transitivity :

$$\frac{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow^\forall \varphi_2 \qquad \mathcal{S}, \mathcal{A} \cup \mathcal{C} \vdash \varphi_2 \Rightarrow^\forall \varphi_3}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow^\forall \varphi_3}$$

Case Analysis :

$$\frac{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow^\forall \varphi \qquad \mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_2 \Rightarrow^\forall \varphi}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1 \vee \varphi_2 \Rightarrow^\forall \varphi}$$

Abstraction :

$$\frac{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi' \qquad X \cap \mathit{FreeVars}(\varphi') = \emptyset}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \exists X \ \varphi \Rightarrow^\forall \varphi'}$$

Consequence :

$$\frac{\models \varphi_1 \to \varphi_1' \qquad \mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1' \Rightarrow^\forall \varphi_2' \qquad \models \varphi_2' \to \varphi_2}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi_1 \Rightarrow^\forall \varphi_2}$$

Circularity :

$$\frac{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C} \cup \{\varphi \Rightarrow^\forall \varphi'\}} \varphi \Rightarrow^\forall \varphi'}{\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'}$$

Figure 1: Proof system for reachability. We make the standard assumption that the free variables of $\varphi_l \Rightarrow^\exists \varphi_r$ in the STEP proof rule are fresh (e.g., disjoint from those of $\varphi \Rightarrow^\forall \varphi'$).

rules provide a unified representation for both language semantics and program specifications: $\varphi \Rightarrow^\exists \varphi'$ for semantics and $\varphi \Rightarrow^\forall \varphi'$ for all-path reachability specifications. Note that, like Hoare triples, reachability rules can only specify properties of complete paths (that is, terminating execution paths). One can use existing Hoare logic techniques to break reasoning about a non-terminating program into reasoning about its terminating components.

4. Reachability Proof System

Fig. 1 shows our novel proof system for all-path reachability. The target language is given as a weakly well-defined reachability system $\mathcal{S}$. The soundness result (Thm. 7.1) guarantees that $\mathcal{S} \models \varphi \Rightarrow^\forall \varphi'$ if $\mathcal{S} \vdash \varphi \Rightarrow^\forall \varphi'$ is derivable. Note that the proof system derives more general sequents of the form $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'$, where $\mathcal{A}$ and $\mathcal{C}$ are sets of reachability rules. The rules in $\mathcal{A}$ are called *axioms* and rules in $\mathcal{C}$ are called *circularities*. If either $\mathcal{A}$ or $\mathcal{C}$ does not appear in a sequent, it means the respective set is empty: $\mathcal{S} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'$ is a shorthand for $\mathcal{S}, \emptyset \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'$, and $\mathcal{S}, \mathcal{A} \vdash \varphi \Rightarrow^\forall \varphi'$ is a shorthand for $\mathcal{S}, \mathcal{A} \vdash_{\emptyset} \varphi \Rightarrow^\forall \varphi'$, and $\mathcal{S} \vdash \varphi \Rightarrow^\forall \varphi'$ is a shorthand for $\mathcal{S}, \emptyset \vdash_{\emptyset} \varphi \Rightarrow^\forall \varphi'$. Initially, both $\mathcal{A}$ and $\mathcal{C}$ are empty. Note that "$\to$" in STEP and CONSEQUENCE denotes logical implication.

The intuition is that the reachability rules in $\mathcal{A}$ can be assumed valid, while those in $\mathcal{C}$ have been postulated but not yet justified. After making progress from φ (at least one derivation by STEP or by AXIOM with the rules in $\mathcal{A}$), the rules in $\mathcal{C}$ become (coinductively) valid (can be used in derivations by AXIOM). During the proof, circularities can be added to $\mathcal{C}$ via CIRCULARITY, flushed into $\mathcal{A}$ by TRANSITIVITY, and used via AXIOM. The desired semantics for sequent $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^\forall \varphi'$ (read "$\mathcal{S}$ with axioms $\mathcal{A}$ and circularities $\mathcal{C}$ proves $\varphi \Rightarrow^\forall \varphi'$") is: $\varphi \Rightarrow^\forall \varphi'$ holds if the rules in $\mathcal{A}$ hold and those in $\mathcal{C}$ hold after taking at least one step from φ in the transition system $(\Rightarrow^\mathcal{T}_{\mathcal{S}}, \mathcal{T})$, and if $\mathcal{C} \neq \emptyset$ then φ reaches φ' after at least one step on all complete paths. As a consequence of this definition, any rule $\varphi \Rightarrow^\forall \varphi'$

derived by CIRCULARITY has the property that φ reaches φ' after at lest one step, due to CIRCULARITY having a prerequisite $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C} \cup \{\varphi \Rightarrow^\forall \varphi'\}} \varphi \Rightarrow^\forall \varphi'$ (with a non-empty set of circularities). We next discuss the proof rules.

STEP derives a sequent where φ reaches φ' in one step on all paths. The first premise ensures any configuration matching φ matches the left-hand-side φ_l of some rule in $\mathcal{S}$ and thus, as $\mathcal{S}$ is weakly well-defined, can take a step. Formally, if $(\gamma, \rho) \models \varphi$, then there exists some rule $\varphi_l \Rightarrow^\exists \varphi_r \in \mathcal{S}$ and some valuation ρ' of the free variables of φ_l such that $(\gamma, \rho') \models \varphi_l$, and thus γ has at least one $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-successor generated by the rule $\varphi_l \Rightarrow^\exists \varphi_r$. The second premise ensures that each $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-successor of a configuration matching φ matches φ'. Formally, if $\gamma \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'$ and γ matches φ then there is some rule $\varphi_l \Rightarrow^\exists \varphi_r \in \mathcal{S}$ and $\rho : Var \to \mathcal{T}$ such that $(\gamma, \rho) \models \varphi \wedge \varphi_l$ and $(\gamma', \rho) \models \varphi_r$; then the second premise implies γ' matches φ'.

Designing a proof rule for deriving an execution step along all paths is non-trivial. For instance, one might expect STEP to require as many premises as there are transitions going out of φ, as is the case for the examples presented later in this paper. However, that is not possible, as the number of successors of a configuration matching φ may be unbounded even if each matching configuration has a finite branching factor in the transition system. STEP avoids this issue by requiring only one premise for each rule by which some configuration φ can take a step, even if that rule can be used to derive multiple transitions. To illustrate this situation, consider a language defined by $\mathcal{S} \equiv \{\langle n_1 \rangle \wedge n_1 >_{Int} n_2 \Rightarrow^\exists \langle n_2 \rangle\}$, with n_1 and n_2 non-negative integer variables. A configuration in this language is a singleton with a non-negative integer. Intuitively, a positive integer transits into a strictly smaller non-negative integer, in a non-deterministic way. The branching factor of a non-negative integer is its value. Then $\mathcal{S} \models \langle m \rangle \Rightarrow^\forall \langle 0 \rangle$. Deriving it reduces (by CIRCULARITY and other proof rules) to deriving $\langle m_1 \rangle \wedge m_1 >_{Int} 0 \Rightarrow^\forall \exists m_2 (\langle m_2 \rangle \wedge m_1 >_{Int} m_2)$. The left-hand-side is matched by any positive integer, and thus its branching factor is infinity. Deriving this rule with STEP requires only two premises, $\models (\langle m_1 \rangle \wedge m_1 >_{Int} 0) \to \exists n_1 n_2 (\langle n_1 \rangle \wedge n_1 >_{Int} n_2)$ and $\models \exists c\, (c = \langle m_1 \rangle \wedge m_1 >_{Int} 0 \wedge c = \langle n_1 \rangle \wedge n_1 >_{Int} n_2) \wedge \langle n_2 \rangle \to \exists m_2 (\langle m_2 \rangle \wedge m_1 >_{Int} m_2)$. A similar situation arises in real life for languages with thread pools of arbitrary size.

AXIOM applies a trusted rule. REFLEXIVITY and TRANSITIVITY capture the corresponding closure properties of the reachability relation. REFLEXIVITY requires $\mathcal{C}$ to be empty to ensure that all-path rules derived with non-empty $\mathcal{C}$ take at least one step. TRANSITIVITY enables the circularities as axioms for the second premise, since if $\mathcal{C}$ is not empty, the first premise is guaranteed to take at least one step. CONSEQUENCE, CASE ANALYSIS and ABSTRACTION are adapted from Hoare logic. Ignoring circularities, these seven rules discussed so far constitute formal infrastructure for symbolic execution.

CIRCULARITY has a coinductive nature, allowing us to make new circularity claims. We typically make such claims for code with repetitive behaviors, such as loops, recursive functions, jumps, etc. If there is a derivation of the claim using itself as a circularity, then the claim holds. This would obviously be unsound if the new assumption was available immediately, but requiring progress (taking at least one step in the transition system $(\mathcal{T}, \Rightarrow_{\mathcal{S}}^{\mathcal{T}})$) before circularities can be used ensures that only diverging executions can correspond to endless invocation of a circularity.

One important aspect of concurrent program verification, which we do not address in this paper, is proof compositionality. Our focus here is limited to establishing a sound and complete language-independent proof system for all-path reachability rules, to serve as a foundation for further results and applications, and to discuss our current implementation of it. We only mention that we have already studied proof compositionality for earlier one-path

IMP language syntax			IMP evaluation contexts syntax	
$PVar$	$::=$	program variables	$Context$	$::=$ ■
Exp	$::=$	$PVar$ | Int | Exp **op** Exp		| $\langle Context,\ State \rangle$
$Stmt$	$::=$	**skip** | $PVar := Exp$		| $Context$ **op** Exp | Int **op** $Context$
	|	$Stmt; Stmt$ | $Stmt \parallel Stmt$		| $PVar := Context$ | $Context; Stmt$
	|	**if**$(Exp)\ Stmt$ **else** $Stmt$		| $Context \parallel Stmt$ | $Stmt \parallel Context$
	|	**while**$(Exp)\ Stmt$		| **if**$(Context)\ Stmt$ **else** $Stmt$

IMP operational semantics			
lookup	$\langle C,\ \sigma \rangle[x] \Rightarrow^{\exists} \langle C,\ \sigma \rangle[\sigma(x)]$	**op**	i_1 **op** $i_2 \Rightarrow^{\exists} i_1\ op_{Int}\ i_2$
asgn	$\langle C,\ \sigma \rangle[x := i] \Rightarrow^{\exists} \langle C,\ \sigma[x \leftarrow i]\rangle[\texttt{skip}]$	**seq**	$\texttt{skip}; s \Rightarrow^{\exists} s$
cond$_1$	$\texttt{if}(i)\ s_1\ \texttt{else}\ s_2 \Rightarrow^{\exists} s_1$ if $i \neq 0$	**cond$_2$**	$\texttt{if}(0)\ s_1\ \texttt{else}\ s_2 \Rightarrow^{\exists} s_2$
while	$\texttt{while}(e)\ s \Rightarrow^{\exists} \texttt{if}(e)\ s; \texttt{while}(e)\ s\ \texttt{else skip}$	**finish**	$\texttt{skip} \parallel \texttt{skip} \Rightarrow^{\exists} \texttt{skip}$

Figure 2: IMP language syntax and operational semantics based on evaluation contexts.

variants of reachability logic [5], showing that there is a mechanical way to translate any Hoare logic proof derivation into a reachability proof of similar size and structure, but based entirely on the operational semantics of the language. The overall conclusion of our previous study, which we believe will carry over to all-path reachability, was that compositional reasoning can be achieved methodologically using our proof system, by proving and then using appropriate reachability rules as lemmas. However, note that this works only for theoretically well-behaved languages which enjoy a compositional semantics. For example, a language whose semantics assumes a bounded heap size, or which has constructs whose semantics involve the entire program, e.g., call/cc, will lack compositionality.

5. Verifying Programs

In this section we show a few examples of using our proof system to verify programs based on an operational semantics. In a nutshell, the proof system enables generic symbolic execution combined with circular reasoning. Symbolic execution is achieved by rewriting modulo domain reasoning.

First, we introduce a simple parallel imperative language, IMP. Fig. 2 shows its syntax and an operational semantics based on evaluation contexts [15] (we choose evaluation contexts for presentation purposes only). IMP has only integer expressions. When used as conditions of `if` and `while`, zero means false and any non-zero integer means true (like in C). Expressions are formed with integer constants, program variables, and conventional arithmetic constructs. Arithmetic operations are generically described as **op**. IMP statements are assignment, `if`, `while`, sequential composition and parallel composition. IMP has shared memory parallelism without explicit synchronization. The examples use the parallel construct only at the top-level of the programs. The second example shows how to achieve synchronization using the existing language constructs.

The program configurations of IMP are pairs $\langle \texttt{code},\ \sigma \rangle$, where `code` is a program fragment and σ is a state term mapping program variables into integers. As usual, we assume appropriate definitions for the integer and map domains available, together with associated operations like arithmetic operations ($i_1\ op_{Int}\ i_2$, etc.) on the integers and lookup ($\sigma(x)$) and update ($\sigma[x \leftarrow i]$) on the maps. We also assume a context domain with a plugging operation ($C[t]$) that composes a context and term back into a configuration. A configuration context consists of a code context and a state. The definition in Fig. 2 consists of eight
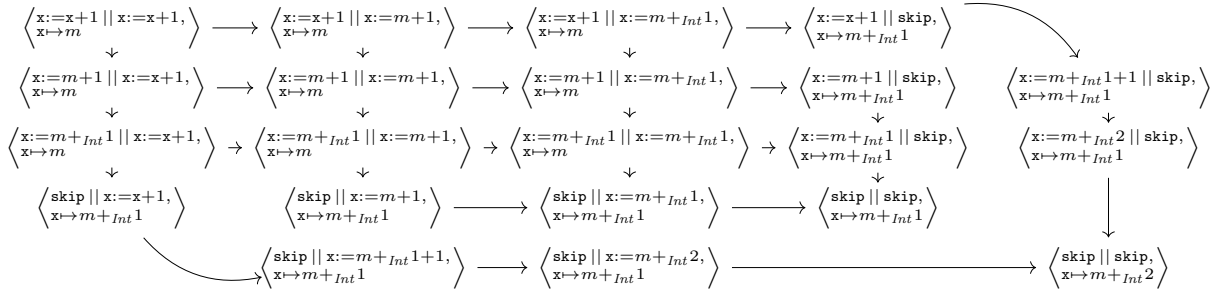
$$\langle \substack{\texttt{x:=x+1 || x:=x+1,} \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=x+1 || x:=m+1,} \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=x+1 || x:=}m+_{Int}1, \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=x+1 || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle$$

$$\langle \substack{\texttt{x:=m+1 || x:=x+1,} \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=m+1 || x:=m+1,} \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=m+1 || x:=}m+_{Int}1, \\ \texttt{x}\mapsto m} \rangle \longrightarrow \langle \substack{\texttt{x:=m+1 || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle \qquad \langle \substack{\texttt{x:=}m+_{Int}1+1 \texttt{ || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle$$

$$\langle \substack{\texttt{x:=}m+_{Int}1 \texttt{ || x:=x+1,} \\ \texttt{x}\mapsto m} \rangle \to \langle \substack{\texttt{x:=}m+_{Int}1 \texttt{ || x:=m+1,} \\ \texttt{x}\mapsto m} \rangle \to \langle \substack{\texttt{x:=}m+_{Int}1 \texttt{ || x:=}m+_{Int}1, \\ \texttt{x}\mapsto m} \rangle \to \langle \substack{\texttt{x:=}m+_{Int}1 \texttt{ || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle \qquad \langle \substack{\texttt{x:=}m+_{Int}2 \texttt{ || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle$$

$$\langle \substack{\texttt{skip || x:=x+1,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle \qquad \langle \substack{\texttt{skip || x:=m+1,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle \longrightarrow \langle \substack{\texttt{skip || x:=}m+_{Int}1, \\ \texttt{x}\mapsto m+_{Int}1} \rangle \longrightarrow \langle \substack{\texttt{skip || skip,} \\ \texttt{x}\mapsto m+_{Int}1} \rangle$$

$$\langle \substack{\texttt{skip || x:=}m+_{Int}1+1, \\ \texttt{x}\mapsto m+_{Int}1} \rangle \longrightarrow \langle \substack{\texttt{skip || x:=}m+_{Int}2, \\ \texttt{x}\mapsto m+_{Int}1} \rangle \longrightarrow \langle \substack{\texttt{skip || skip,} \\ \texttt{x}\mapsto m+_{Int}2} \rangle$$

Figure 3: State space of the racing increment example

reduction rules between program configurations, which make use of first-order variables: x is a variable of sort $PVar$; e is a variable of sort Exp; s, s_1, s_2 are variables of sort $Stmt$; i, i_1, i_2 are variables of sort Int; σ is a variable of sort $State$; C is a variable of sort $Context$. A rule reduces a configuration by splitting it into a context and a redex, rewriting the redex and possibly the context, and then plugging the resulting term into the resulting context. As an abbreviation, a context is not mentioned if not used; e.g., the rule **op** is in full $\langle C, \ \sigma \rangle [i_1 \ \textbf{op} \ i_2] \Rightarrow^\exists \langle C, \ \sigma \rangle [i_1 \ op_{Int} \ i_2]$. For example, configuration $\langle \texttt{x} := (2 + 5) - 4, \ \sigma \rangle$ reduces to $\langle \texttt{x} := 7 - 4, \ \sigma \rangle$ by applying the $\textbf{op}_+$ rule with $C \equiv \texttt{x} := \blacksquare - 4$, $\sigma \equiv \sigma$, $i_1 \equiv 2$ and $i_2 \equiv 5$. We can regard the operational semantics of IMP above as a set of reduction rules of the form "$l \Rightarrow^\exists r$ if b", where l and r are program configurations with variables constrained by boolean condition b. As discussed in Section 3, our proof system works with any rules of this form.

Next, we illustrate the proof system on a few examples. The first example shows that our proof system enables exhaustive state exploration, similar to symbolic model-checking but based on the operational semantics. Although humans prefer to avoid such explicit proofs and instead methodologically use abstraction or compositional reasoning whenever possible (and such methodologies are not excluded by our proof system), a complete proof system must nevertheless support them. The code $\texttt{x} := \texttt{x} + 1 \ || \ \texttt{x} := \texttt{x} + 1$ exhibits a race on $\texttt{x}$: the value of $\texttt{x}$ increases by 1 when both reads happen before either write, and by 2 otherwise. The all-path rule that captures this behavior is

$$\langle \texttt{x:=x+1 || x:=x+1}, \ \texttt{x} \mapsto m \rangle \Rightarrow^\forall \exists n \ (\langle \texttt{skip}, \ \texttt{x} \mapsto n \rangle \land (n = m +_{Int} 1 \lor n = m +_{Int} 2)$$

We show that the program has exactly these behaviors by deriving this rule in the proof system. Call the right-hand-side pattern G. The proof contains subproofs of $c \Rightarrow^\forall G$ for every reachable configuration c, tabulated in Fig. 3. The subproofs for c matching G use REFLEXIVITY and CONSEQUENCE, while the rest use TRANSITIVITY, STEP, and CASE ANALYSIS to reduce to the proofs for the next configurations. For example, the proof fragment below shows how $\langle \texttt{x} := m + 1 \, || \, \texttt{x} := \texttt{x} + 1, \ \texttt{x} \mapsto m \rangle \Rightarrow^\forall G$ reduces to $\langle \texttt{x} := m +_{Int} 1 \, || \, \texttt{x} := \texttt{x} + 1, \ \texttt{x} \mapsto m \rangle \Rightarrow^\forall G$ and $\langle \texttt{x} := m + 1 \, || \, \texttt{x} := m + 1, \ \texttt{x} \mapsto m \rangle \Rightarrow^\forall G$:

$$\text{STEP} \ \cfrac{ \cfrac{\ldots}{\langle \substack{\texttt{x:=m+1 || x:=x+1,} \\ \texttt{x}\mapsto m}\rangle \Rightarrow^\forall \langle \substack{\texttt{x:=}m+_{Int}1\texttt{ || x:=x+1,}\\ \texttt{x}\mapsto m}\rangle \lor \langle \substack{\texttt{x:=m+1 || x:=m+1,}\\ \texttt{x}\mapsto m}\rangle} \qquad \cfrac{\cfrac{\ldots}{\langle \substack{\texttt{x:=}m+_{Int}1\texttt{ || x:=x+1,}\\ \texttt{x}\mapsto m}\rangle \Rightarrow^\forall G} \quad \cfrac{\ldots}{\langle \substack{\texttt{x:=m+1 || x:=m+1,}\\ \texttt{x}\mapsto m}\rangle \Rightarrow^\forall G}}{\langle \substack{\texttt{x:=}m+_{Int}1\texttt{ || x:=x+1,}\\ \texttt{x}\mapsto m}\rangle \lor \langle \substack{\texttt{x:=m+1 || x:=m+1,}\\ \texttt{x}\mapsto m}\rangle \Rightarrow^\forall G} \text{CA}}{\langle \texttt{x} := m + 1 \, || \, \texttt{x} := \texttt{x} + 1, \ \texttt{x} \mapsto m \rangle \Rightarrow^\forall G} \text{TRANS}$$

For the rule hypotheses of STEP above, note that all rules but **lookup** and **op$_+$** make the overlap condition $\exists c \left(\left\langle {\mathtt{x}:=m+1 \,||\, \mathtt{x}:=\mathtt{x}+1,} \atop {\mathtt{x}\mapsto m} \right\rangle [c/\square] \wedge \varphi_l[c/\square] \right)$ unsatisfiable, and only one choice of free variables works for the **lookup** and **op$_+$** rules. For **lookup**, φ_l is $\langle C, \sigma \rangle[x]$ and the overlap condition is only satisfiable if the logical variables C, σ and x are equal to $(\mathtt{x} := m + 1 \,||\, \mathtt{x} := \blacksquare + 1)$, $(\mathtt{x} \mapsto m)$, and $\mathtt{x}$, respectively. Under this assignment, the pattern $\varphi_r = \langle C, \sigma \rangle[\sigma(x)]$ is equivalent to $\langle \mathtt{x} := m + 1 \,||\, \mathtt{x} := m + 1, \ \mathtt{x} \mapsto m \rangle$, the right branch of the disjunction. The **op$_+$** rule is handled similarly. The assignment for **lookup** can also witness the existential in the progress hypothesis of STEP. Subproofs for other states in Fig. 3 can be constructed similarly.

    The next two examples use loops and thus need to state and prove invariants. As discussed in [4], CIRCULARITY generalizes the various language-specific invariant proof rules encountered in Hoare logics. One application is reducing a proof of $\varphi \Rightarrow^\forall \varphi'$ to proving $\varphi_{\mathrm{inv}} \Rightarrow^\forall \varphi_{\mathrm{inv}} \vee \varphi'$ for some pattern invariant φ_{inv}. We first show $\models \varphi \rightarrow \varphi_{\mathrm{inv}}$, and use CONSEQUENCE to change the goal to $\varphi_{\mathrm{inv}} \Rightarrow^\forall \varphi'$. This is claimed as a circularity, and then proved by transitivity with $\varphi_{\mathrm{inv}} \vee \varphi'$. The second hypothesis $\{\varphi_{\mathrm{inv}} \Rightarrow^\forall \varphi'\} \vdash \varphi_{\mathrm{inv}} \vee \varphi' \Rightarrow^\forall \varphi'$ is proved by CASE ANALYSIS, AXIOM, and REFLEXIVITY.

```
f0 = 1;               f1 = 1;
turn = 1;             turn = 0;
while (f1 && turn)    while (f0 && (1 - turn))
  skip                  skip
x = x + 1;            x = x + 1;
f0 = 0;              f1 = 0;
```

Figure 4: Peterson's algorithm (threads $\mathtt{T_0}$ and $\mathtt{T_1}$)

    Next, we can use Peterson's algorithm for mutual exclusion to eliminate the race as shown in Fig. 4. The all-path rule $\varphi \Rightarrow^\forall \varphi'$ that captures the new behavior is

$$\langle \mathtt{T_0} \,||\, \mathtt{T_1}, \ (\mathtt{f0} \mapsto 0, \ \mathtt{f1} \mapsto 0, \ \mathtt{x} \mapsto N) \rangle$$
$$\Rightarrow^\forall \ \exists t \, \langle \mathtt{skip}, \ (\mathtt{f0} \mapsto 0, \ \mathtt{f1} \mapsto 0, \ \mathtt{x} \mapsto N +_{Int} 2, \ \mathtt{turn} \mapsto t) \rangle$$

Similarly to the unsynchronized example, the proof contains subproofs of $c \Rightarrow^\forall \varphi'$ for every configuration c reachable from φ. The main difference is that CIRCULARITY is used with each of these rules $c \Rightarrow^\forall \varphi'$ with one of the two threads of c in the `while` loop (these rules capture the invariants). Thus, when we reach a configuration c visited before, we use the rule added by CIRCULARITY to complete the proof.

    The final example is the program $\mathtt{SUM} \equiv$ "`s := 0; LOOP`" (where `LOOP` stands for "`while (n>0) (s := s+n; n := n-1)`"), which computes in `s` the sum of the numbers from 1 up to `n`. The all-path reachability rule $\varphi \Rightarrow^\forall \varphi'$ capturing this behavior is

$$\langle \mathtt{SUM}, \ (\mathtt{s} \mapsto s, \ \mathtt{n} \mapsto n) \rangle \wedge n \geq_{Int} 0 \ \Rightarrow^\forall \ \langle \mathtt{skip}, \ (\mathtt{s} \mapsto n *_{Int} (n +_{Int} 1)/_{Int} 2, \ \mathtt{n} \mapsto 0) \rangle$$

We derive the above rule in our proof system by using CIRCULARITY with the invariant rule

$$\exists n'(\langle \mathtt{LOOP}, \ (\mathtt{s} \mapsto (n -_{Int} n') *_{Int} (n +_{Int} n' +_{Int} 1)/_{Int} 2, \ \mathtt{n} \mapsto n') \rangle \wedge n' \geq_{Int} 0) \Rightarrow^\forall \varphi'$$

Previous work [4–7] presented a proof system able to derive similar rules, but which hold along *some* execution path, requiring a separate proof that the program is deterministic.

6. IMPLEMENTATION

Here we briefly discuss our prototype implementation of the proof system in Fig. 1 in $\mathbb{K}$ [9]. We choose $\mathbb{K}$ because it is a modular semantic language design framework, it is used for teaching programming languages at several universities, and there are several languages

defined in it including C [13], PHP [18], Python, and Java. For brevity, we do not present $\mathbb{K}$ here. We refer the reader to http://kframework.org for language definitions, a tutorial, and our prototype. As discussed in Section 3, we simply view a $\mathbb{K}$ semantics as a set of reachability rules of the form "$l \wedge b \Rightarrow^\exists r$".

The prototype is implemented in Java, and uses Z3 [19] for domain reasoning. It takes an operational semantics and uses it to perform concrete or symbolic execution. At its core, it performs *narrowing* of a conjunctive pattern with reachability rules between conjunctive patterns, where a conjunctive pattern is a pattern of the form $\exists X(\pi \wedge \psi)$, with X a set of variables, π a basic pattern (program configurations with variables), and ψ a structureless formula. Narrowing is necessary when a conjunctive pattern is too abstract to match the left-hand side of any rule, but is unifiable with the left-hand sides of some rules. For instance, consider the IMP code fragment "if (b) then x = 1; else x = 0;". This code does not match the left-hand sides of either of the two rules giving semantics to if (similar to **cond**$_1$ and **cond**$_2$ in Fig. 2), but it is unifiable with the left-hand sides of both rules. Intuitively, if we use the rules of the semantics, taking steps of rewriting on a ground configuration yields concrete execution, while taking steps of narrowing on a conjunctive pattern yields symbolic execution. In our practical evaluation, we found that conjunctive patterns tend to suffice to specify both the rules for operational semantics and program specifications.

For each step of narrowing, the $\mathbb{K}$ engine uses *unification modulo theories.* In our implementation, we distinguish a number of mathematical theories (e.g. booleans, integers, sequences, sets, maps, etc) which the underlying SMT solver can reason about. Specifically, when unifying a conjunctive pattern $\exists X(\pi \wedge \psi)$ with the left-hand side of a rule $\exists X_l(\pi_l \wedge \psi_l)$ (we assume $X \cap X_l = \emptyset$), the $\mathbb{K}$ engine begins with the syntactic unification of the basic patterns π and π_l. Upon encountering corresponding subterms (π' in π and π'_l in π_l) which are both terms of one of the theories above, it records an equality $\pi' = \pi'_l$ rather than decomposing the subterms further (if one is in a theory, and the other one is in a different theory or is not in any theory, the unification fails). If this stage of unification is successful, we end up with a conjunction ψ_u of constraints, some having a variable in one side and some with both sides in one of the theories. Satisfiability of $\exists X \cup X_l(\psi \wedge \psi_u \wedge \psi_l)$ is then checked by the SMT solver. If it is satisfiable, then narrowing takes a step from $\exists X(\pi \wedge \psi)$ to $\exists X \cup X_l \cup X_r(\pi_r \wedge \psi \wedge \psi_u \wedge \psi_l \wedge \psi_r)$, where $\exists X_r(\pi_r \wedge \psi_r)$ is the right-hand side of the rule. Intuitively, "collecting" the constraints $\psi_u \wedge \psi_l \wedge \psi_r$ is similar to collecting the path constraint in traditional symbolic execution (but is done in a language-generic manner). For instance, in the if case above, narrowing with the two semantics rules results in collecting the constraints $b =$ true and $b =$ false.

The $\mathbb{K}$ engine accepts a set of user provided rules to prove together, which capture the behavior of the code being verified. Typically, these rules specify the behavior of recursive functions and while loops. For each rule, the $\mathbb{K}$ engine searches starting from the left-hand side for formulae which imply the right-hand side, starting with $\mathcal{S}$ the semantics and $\mathcal{C}$ all the rules it attempts to prove. By a derived proof rule called *Set Circularity*, this suffices to show that each rule is valid. As an optimization, AXIOM is given priority over STEP (use specifications rather than stepping into the code).

Most work goes into implementing the STEP proof rule, and in particular calculating how $\rho \models \exists c \, (\varphi[c/\Box] \wedge \varphi_l[c/\Box])$ can be satisfied. This holds when $\rho^\gamma \models \varphi$ and $\rho^\gamma \models \varphi_l$, which can be checked with unification modulo theories. To use STEP in an automated way, the $\mathbb{K}$ tool constructs φ' for a given φ as a disjunction of $\varphi_r \wedge \psi_u \wedge \psi \wedge \psi_l$ over each rule $\varphi_l \Rightarrow^\exists \varphi_r \in \mathcal{S}$ and each way ψ_u of unifying φ with φ_l. As discussed in Section 4, in general

this disjunction may not be finite, but it is sufficient for the examples that we considered. The CONSEQUENCE proof rule also requires unification modulo theories, to check validity of the implication hypothesis $\models \varphi_1 \rightarrow \varphi'_1$. The main difference from STEP is that the free variables of φ' become universality quantified when sending the query to the SMT solver. The implementation of the other proof rules is straight-forward.

7. SOUNDNESS

Here we discuss the soundness of our proof system. Unlike the similar results for Hoare logics and dynamic logics, which are separately proved for each language taking into account the particularities of that language, we prove soundness once and for all languages.

Soundness states that a syntactically derivable sequent holds semantically. Because of the utmost importance of the result below, we have also mechanized its proof. Our complete Coq formalization can be found at `http://fsl.cs.illinois.edu/rl`.

Theorem 7.1 (Soundness). *If $\mathcal{S} \vdash \varphi \Rightarrow^{\forall} \varphi'$ then $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$.*

Proof. Due to CIRCULARITY and TRANSITIVITY the sets $\mathcal{A}$ and $\mathcal{C}$ may not remain empty, so we need to prove a more general statement (Lemma 7.2 below), including hypotheses making some semantic assumptions about the rules in $\mathcal{A}$ and $\mathcal{C}$.

To express the different assumptions needed about the available rules in $\mathcal{A}$ and the still-unavailable rules in $\mathcal{C}$ we define a more specific satisfaction relation. Let $\delta \in \{+, *\}$ be a flag and let $n \in \mathbb{N}$ be a natural number. We define the new satisfaction relation $\mathcal{S} \models^{\delta}_n \varphi \Rightarrow^{\forall} \varphi'$ by restricting the paths in the definition of $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$ to length at most n, and requiring that φ be reached after at least one step if $\delta = +$. Formally, we define $\mathcal{S} \models^{\delta}_n \varphi \Rightarrow^{\forall} \varphi'$ to hold iff for any complete path $\tau = \gamma_1...\gamma_k$ of length $k \leq n$ and for any ρ such that $(\gamma_1, \rho) \models \varphi$, there exists $i \in \{1, ..., k\}$ such that $(\gamma_i, \rho) \models \varphi'$ and also that this i is not 1 (i.e. γ_1 makes progress) if $\delta = +$.

Observe that $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$ iff $\mathcal{S} \models^{*}_n \varphi \Rightarrow^{\forall} \varphi'$ for all $n \in \mathbb{N}$. The forward implication holds because a complete path $\tau = \gamma_1...\gamma_k$ of length $k \leq n$ is a complete path simpliciter. The reverse implication holds because given any complete path $\tau = \gamma_1...\gamma_k$ there are choices (such as k) of $n \in \mathbb{N}$ with $k \leq n$. Theorem 7.1 then follows from Lemma 7.2 by noting that the hypotheses about $\mathcal{A}$ and $\mathcal{C}$ are vacuously true when those sets are empty. $\square$

Lemma 7.2. *If $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^{\forall} \varphi'$ then for any n, $\mathcal{S} \models^{+}_n \mathcal{A}$ and $\mathcal{S} \models^{+}_{n-1} \mathcal{C}$ imply that $\mathcal{S} \models^{*}_n \varphi \Rightarrow^{\forall} \varphi'$, and also that $\mathcal{S} \models^{+}_n \varphi \Rightarrow^{\forall} \varphi'$ if $\mathcal{C}$ is not empty.*

Proof. To state the conclusion of the lemma more concisely we define flag Δ_C to be $*$ if C is an empty set and $+$ otherwise. The conclusion is $\mathcal{S} \models^{\Delta_C}_n \varphi \Rightarrow^{\forall} \varphi'$. The proof proceeds by induction on the proof tree showing $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \varphi \Rightarrow^{\forall} \varphi'$ (keeping induction hypotheses universally quantified over n), using an inner induction over n in the CIRCULARITY case.
CIRCULARITY: The induction hypothesis states that for any $m \in \mathbb{N}$.

$$\text{if } \mathcal{S} \models^{+}_m \mathcal{A} \text{ and } \mathcal{S} \models^{+}_{m-1} \mathcal{C} \cup \{\varphi \Rightarrow^{\forall} \varphi'\} \text{ then } \mathcal{S} \models^{+}_m \varphi \Rightarrow^{\forall} \varphi'. \tag{7.1}$$

We will strengthen $\Delta_{\mathcal{C}}$ to $+$ and prove that for any $n \in \mathbb{N}$, if $\mathcal{S} \models^{+}_n \mathcal{A}$ and $\mathcal{S} \models^{+}_{n-1} \mathcal{C}$ then $\mathcal{S} \models^{+}_n \varphi \Rightarrow^{\forall} \varphi'$. Proceed by induction on n.
(1) if $n = 0$, the conclusion $\mathcal{S} \models^{+}_0 \varphi \Rightarrow^{\forall} \varphi'$ of the implication is vacuously true.

(2) for $n > 0$, we have an inner induction hypothesis that $\mathcal{S} \models^+_{n-1} \mathcal{A}$ and $\mathcal{S} \models^+_{n-2} \mathcal{C}$ together imply $\mathcal{S} \models^+_{n-1} \varphi \Rightarrow^\forall \varphi'$, and must prove that $\mathcal{S} \models^+_n \mathcal{A}$ and $\mathcal{S} \models^+_{n-1} \mathcal{C}$ together imply $\mathcal{S} \models^+_n \varphi \Rightarrow^\forall \varphi'$, Instantiating the inner hypothesis with the stronger assumptions about $\mathcal{A}$ and $\mathcal{C}$ gives $\mathcal{S} \models^+_{n-1} \varphi \Rightarrow^\forall \varphi'$. With other assumptions this gives $\mathcal{S} \models^+_{n-1} \mathcal{C} \cup \{\varphi \Rightarrow^\forall \varphi'\}$. Now we apply the outer induction hypothesis, Equation (7.1), with $m = n$ to conclude $\mathcal{S} \models^+_n \varphi \Rightarrow^\forall \varphi'$.

This was the only case that made essential use of the quantification over n in the conclusion. For the remaining cases, we assume that $n \in \mathbb{N}$ is already fixed, and assumptions $\mathcal{S} \models^+_n \mathcal{A}$ and $\models^+_{n-1} \mathcal{C}$ are given in addition to the assumptions and inductive hypotheses coming from the specific proof rule. Inductive hypotheses with the same $\mathcal{A}$ and $\mathcal{C}$ as the conclusion will be stated in the form $\mathcal{S} \models^{\Delta c}_n \varphi_a \Rightarrow^\forall \varphi_b$, already specialized to n and applied to the semantic assumptions on $\mathcal{A}$ and $\mathcal{C}$. Furthermore, the remaining cases all prove the conclusion $\mathcal{S} \models^+_n \varphi \Rightarrow^\forall \varphi'$ directly from the definition, so we also assume that a complete path $\tau = \gamma_1, ..., \gamma_k$ with $k \leq n$ and a ρ with $(\gamma_1, \rho) \models \varphi$ are fixed. This leaves the obligation to find an $i \in \{1, ..., k\}$ such that $(\gamma_i, \rho) \models \varphi'$ and that $i > 1$ if $\mathcal{C}$ is nonempty.

TRANSITIVITY: We have inductive hypotheses that $\mathcal{S} \models^{\Delta c}_n \varphi \to \varphi_2$, and for any $m \in \mathbb{N}$, $\mathcal{S} \models^+_m \mathcal{A} \cup \mathcal{C}$ implies $\mathcal{S} \models^*_m \varphi_2 \to \varphi'$. Applying the first inductive hypothesis to τ and ρ, we receive an i such that $(\gamma_i, \rho) \models \varphi_2$. Now we make separate cases on whether $i = 1$.

- When $i = 1$, $\mathcal{C}$ must be empty, so our assumption $\mathcal{S} \models^+_n \mathcal{A}$ suffices to instantiate the second inductive hypothesis at $m = n$ to obtain $\mathcal{S} \models^*_n \varphi_2 \to \varphi'$. As $(\gamma_1, \rho) \models \varphi_2$ we can apply this to τ and ρ to obtain a j such that $(\gamma_j, \rho) \models \varphi'$, concluding this case.
- Otherwise, the suffix of τ beginning with γ_i has length strictly less than n, so it suffices to instantiate the second inductive hypothesis at $m = n - 1$. Our assumptions on $\mathcal{C}$ and $\mathcal{A}$ suffice to conclude $\mathcal{S} \models^+_{n-1} \mathcal{A} \cup \mathcal{C}$. We obtain $\mathcal{S} \models^*_{n-1} \varphi_2 \Rightarrow^\forall \varphi'$. Applying this to ρ and the complete path $\gamma_i, ..., \gamma_k$ yields j such that $(\gamma_{i+j-1}, \rho) \models \varphi'$. As $i + j - 1 > 1$, we can conclude this case with $i + j - 1$ whether or not $\mathcal{C}$ is empty.

STEP: This rule has assumptions $\models \varphi \to \bigvee_{\varphi_l \Rightarrow^\exists \varphi_r \, \in \, \mathcal{S}} \exists FreeVars(\varphi_l)\varphi_l$, and for any $\varphi_l \Rightarrow^\exists \varphi_r \in \mathcal{S} \models \exists c \, (\varphi[c/\square] \wedge \varphi_l[c/\square]) \wedge \varphi_r \to \varphi'$. By the first assumption γ_1 is not stuck, so complete path τ has a second entry γ_2. By the definition of a path, $\gamma_1 \Rightarrow^{\mathcal{T}}_{\mathcal{S}} \gamma_2$. By the definition of $\Rightarrow^{\mathcal{T}}_{\mathcal{S}}$, there exists $\varphi_l \Rightarrow^\exists \varphi_r \in S$ and valuation ρ' such that $(\gamma_1, \rho') \models \varphi_l$ and $(\gamma_2, \rho') \models \varphi_r$. Let $X = FreeVars(\varphi, \varphi')$ and $Y = FreeVars(\varphi_l, \varphi_r)$. We assume without loss of generality that $X \cap Y = \emptyset$. Fix any ρ'' which agrees with ρ on X and with ρ' on Y. With this choice $(\gamma_1, \rho'') \models \varphi$, $(\gamma_1, \rho'') \models \varphi_l$, and $(\gamma_2, \rho'') \models \varphi_r$. By the definition of satisfaction and the FOL translation, then also $(\gamma_2, \rho'') \models \exists c.(\varphi[c/\square] \wedge \varphi_l[c/\square])$ and finally $(\gamma_2, \rho'') \models \exists c.(\varphi[c/\square] \wedge \varphi_l[c/\square]) \wedge \varphi_r$. By the second assumption this implies $(\gamma_2, \rho'') \models \varphi'$. As ρ'' agrees with ρ on X, $(\gamma_2, \rho) \models \varphi'$ (and $2 > 1$), so taking $i = 2$ concludes this case.

REFLEXIVITY: This rule requires $\mathcal{C}$ be empty, and $\varphi = \varphi'$, so $i = 1$ suffices by $(\gamma_1, \rho) \models \varphi$.

AXIOM: This rule has assumption $\varphi \Rightarrow^\forall \varphi' \in \mathcal{A}$. By the assumption on $\mathcal{A}$, $\mathcal{S} \models^+_n \varphi \Rightarrow^\forall \varphi'$. Applying this to τ and ρ gives an $i > 1$ with $(\gamma_i, \rho) \models \varphi'$, concluding this case.

CONSEQUENCE: We have inductive hypothesis that $\mathcal{S} \models^{\Delta c}_n \varphi_1 \Rightarrow^\forall \varphi'_1$, and assumptions $\models \varphi \to \varphi_1$ and $\models \varphi'_1 \to \varphi'$. By the first implication $(\gamma_1, \rho) \models \varphi_1$ as well, so we can apply the inductive hypothesis with τ and ρ to obtain some i such that $(\gamma_i, \rho) \models \varphi'_1$ and $i > 1$ if $\mathcal{C}$ is nonempty. By the second implication, $(\gamma_i, \rho) \models \varphi'$ as well, so this i concludes this case.

CASE ANALYSIS: This rule requires φ have the form $\varphi_1 \vee \varphi_2$, and we have inductive hypotheses $\mathcal{S} \models^{\Delta c}_n \varphi_1 \Rightarrow^\forall \varphi'$ and $\mathcal{S} \models^{\Delta c}_n \varphi_2 \Rightarrow^\forall \varphi'$. Because $(\gamma_1, \rho) \models \varphi$, either $(\gamma_1, \rho) \models \varphi_1$

or $(\gamma_1, \rho) \models \varphi_2$. In either case we can use the respective inductive hypotheses with τ and ρ to obtain an i such that $(\gamma_i, \rho) \models \varphi'$ and $i > 1$ if $\mathcal{C}$ is nonempty, concluding this case.

ABSTRACTION: This rule requires φ have the form $\exists X \varphi_1$ with $X \notin \mathit{FreeVars}(\varphi')$, and we have the inductive hypothesis $\mathcal{S} \models_n^{\Delta c} \varphi_1 \Rightarrow^\forall \varphi'$. Because $(\gamma_1, \rho) \models \exists X \varphi_1$ there exists a ρ' which agrees with ρ except on X such that $(\gamma_1, \rho') \models \varphi_1$. Using the inductive hypothesis with τ and this ρ' gives an i such that $(\gamma_i, \rho') \models \varphi'$, and $i > 1$ if $\mathcal{C}$ is nonempty. As $X \notin \mathit{FreeVars}(\varphi')$, then also $(\gamma_i, \rho) \models \varphi'$, so this i concludes this case. $\qquad \square$

8. RELATIVE COMPLETENESS

Here we show relative completeness: any valid all-path reachability property of any program in any language with an operational semantics given as a reachability system $\mathcal{S}$ is derivable with the proof system in Fig. 1 from $\mathcal{S}$. As with Hoare and dynamic logics, "relative" means we assume an oracle capable of establishing validity in the first-order theory of the state, which here is the configuration model $\mathcal{T}$. Unlike the similar results for Hoare logics, which are separately proved for each language taking into account the particularities of that language, we prove relative completeness once and for all languages. An immediate consequence of relative completeness is that CIRCULARITY is sufficient to derive any repetitive behavior occurring in any program written in any language, and that STEP is also sufficient to derive any non-deterministic behavior!

We establish the relative completeness under the following assumptions:

Framework:

The semantics reachability system $\mathcal{S}$ is
— finite;
The configuration signature Σ has
— a sort $\mathbb{N}$;
— constant symbols 0 and 1 of $\mathbb{N}$;
— binary operation symbols $+$ and $\times$ on $\mathbb{N}$;
— an operation symbol $\alpha : \mathit{Cfg} \to \mathbb{N}$.
The configuration model $\mathcal{T}$ interprets
— $\mathbb{N}$ as the natural numbers;
— constant and operation symbols on $\mathbb{N}$ as corresponding operations;
— $\alpha : \mathit{Cfg} \to \mathbb{N}$ as an injective function.

The assumption that $\mathcal{S}$ is finite ensures STEP has a finite number of prerequisites. The assumption that the model $\mathcal{T}$ includes natural numbers with addition and multiplication is a standard assumption (also made by Hoare and dynamic logic completeness results) which allows the definition of Gödel's β predicate. The assumption that the model $\mathcal{T}$ includes some injective function $\alpha : \mathcal{T}_{\mathit{Cfg}} \to \mathbb{N}$ (that is, the set of configurations $\mathcal{T}_{\mathit{Cfg}}$ is countable) allows the encoding of a sequence of configurations into a sequence of natural numbers. We expect the operational semantics of any reasonable language to satisfy these conditions. Formally, we have the following

Theorem 8.1 (Relative Completeness). *For any semantics $\mathcal{S}$ satisfying the assumptions above, if $\mathcal{S} \models \varphi \Rightarrow^\forall \varphi'$ then $\mathcal{S} \vdash \varphi \Rightarrow^\forall \varphi'$.*

We present an informal proof sketch before going into the formal details, Our proof relies on the fact that pattern reasoning in first-order matching logic reduces to FOL reasoning in

$$step(c,\ c') \equiv \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\mu)\ (\varphi_l[c/\square] \wedge \varphi_r[c'/\square])$$

$$coreach(\varphi) \equiv \forall n \forall c_0...c_n \left(\square = c_0 \rightarrow \bigwedge_{0 \leq i < n} step(c_i,\ c_{i+1}) \rightarrow \neg \exists c_{n+1}\ step(c_n,\ c_{n+1}) \rightarrow \bigvee_{0 \leq i \leq n} \varphi[c_i/\square] \right)$$

Figure 5: FOL encoding of one step transition relation and all-path reachability.

the model $\mathcal{T}$. A key component of the proof is defining the $coreach(\varphi)$ predicate in plain FOL. This predicate states that every complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path τ starting at c includes some configuration satisfying φ. We express $coreach(\varphi)$ using the auxiliary predicate $step(c,\ c')$ that encodes the one step transition relation $(\Rightarrow_{\mathcal{S}}^{\mathcal{T}})$. Fig. 5 shows both definitions. As it is, $coreach(\varphi)$ is not a proper FOL formula, as it quantifies over a sequence of configurations. This is addressed using the injective function α to encode universal quantification over a sequence of configurations into universal quantification over a sequence of integers, which is in turn encoded into quantification over two integer variables using Gödel's β predicate (encoding shown in Fig. 6).

Next, using the definition above we encode the semantic validity of an all-path reachability rule as FOL validity: $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$ iff $\models \varphi \rightarrow coreach(\varphi')$. Therefore, the theorem follows by CONSEQUENCE from the sequent $\mathcal{S} \vdash coreach(\varphi') \Rightarrow^{\forall} \varphi'$. We derive this sequent by using CIRCULARITY to add the rule to the set of circularities, then by using STEP to derive one $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-step, and then by using TRANSITIVITY and AXIOM with the rule itself to derive the remaining $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-steps (circularities can be used after TRANSITIVITY). The formal derivation uses all eight proof rules.

Also recall that, as discussed in Section 2, matching logic is a methodological fragment of the FOL theory of the model $\mathcal{T}$. For technical convenience, in this section we work with the FOL translations $\varphi^{\square}$ instead of the matching logic formulae φ. We mention that in all the formulae used in this section, $\square$ only occurs in the context $\square = t$, thus we stay inside the methodological fragment. For the duration of the proof, we let $c, c', c_0, ..., c_n$ be distinct variables of sort Cfg which do not appear free in the rules in $\mathcal{S}$). We also let $\gamma, \gamma', \gamma_0, ..., \gamma_n$ range over (not necessarily distinct) configurations in the model $\mathcal{T}$, that is, over elements in $\mathcal{T}_{Cfg}$, and let ρ, ρ' range over valuations $Var \rightarrow \mathcal{T}$.

8.1. **Encoding Transition System Operations in FOL.** Fig. 5 shows the definition of the one step transition relation $(\Rightarrow_{\mathcal{S}}^{\mathcal{T}})$ and of the configurations that reach φ on all and complete paths. The former is a (proper) FOL formula, while the later is not, as it quantifies over a sequence of configuration. In Section 8.2 we use Gödel's β predicate to define $\overline{coreach}(\varphi)$, a FOL formula equivalent to $coreach(\varphi)$.

First, we establish the following general purpose lemma

Lemma 8.2. $(\rho(c), \rho) \models \varphi^{\square}$ iff $\rho \models \varphi^{\square}[c/\square]$.

Proof. With the notation in Definition 2.2, $(\rho(c), \rho) \models \varphi^{\square}$ iff $\rho^{\rho(c)} \models \varphi^{\square}$. Notice that if a valuation agrees on two variables, then it satisfies a formula iff it satisfies the formula obtained by substituting one of the two variables for the other. In particular, since $\rho^{\rho(c)}(\square) = \rho^{\rho(c)}(c)$, it follows that $\rho^{\rho(c)} \models \varphi^{\square}$ iff $\rho^{\rho(c)} \models \varphi^{\square}[c/\square]$. We notice that $\square$ does not occur in $\varphi^{\square}[c/\square]$, thus $\rho^{\rho(c)} \models \varphi^{\square}[c/\square]$ iff $\rho \models \varphi^{\square}[c/\square]$, and we are done. $\square$

The following lemma states that $step(c,\ c')$ actually has the semantic properties its name suggests.

Lemma 8.3. $\rho \models step(c,\ c')$ iff $\rho(c) \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \rho(c')$.

Proof. Assume $\rho \models step(c,\ c')$. Then, by the definition of $step(c,\ c')$, there exists some rule $\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}$ such that $\rho \models \exists Free Vars(\mu)\ (\varphi_l[c/\square] \wedge \varphi_r[c'/\square])$. Further, since c and c' do not occur in μ, there exists some ρ' which agrees with ρ on c and c' such that $\rho' \models \varphi_l[c/\square]$ and $\rho' \models \varphi_r[c'/\square]$. By Lemma 8.2, $\rho' \models \varphi_l[c/\square]$ iff $(\rho'(c), \rho') \models \varphi_l$ and $\rho' \models \varphi_r[c'/\square]$ iff $(\rho'(c'), \rho') \models \varphi_r$, so $(\rho'(c), \rho') \models \varphi_l$ and $(\rho'(c'), \rho') \models \varphi_r$. Since ρ and ρ' agree on c and c', it follows that $(\rho(c), \rho') \models \varphi_l$ and $(\rho(c'), \rho') \models \varphi_r$. By Definition 3.1, we conclude $\rho(c) \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \rho(c')$.

Conversely, assume $\rho(c) \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \rho(c')$. Then, by Definition 3.1, there exist some rule $\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}$ and some ρ' for which $(\rho(c), \rho') \models \varphi_l$ and $(\rho(c'), \rho') \models \varphi_r$. Further, since c and c' do not occur in μ, we can choose ρ' to agree with ρ on c and c'. Hence, $(\rho'(c), \rho') \models \varphi_l$ and $(\rho'(c'), \rho') \models \varphi_r$. By Lemma 8.2, $(\rho'(c), \rho') \models \varphi_l$ iff $\rho' \models \varphi_l[c/\square]$ and $(\rho'(c'), \rho') \models \varphi_r$ iff $\rho' \models \varphi_r[c'/\square]$, so $\rho' \models \varphi_l[c/\square]$ and $\rho' \models \varphi_r[c'/\square]$. Since the free variables occurring in $\varphi_l[c/\square] \wedge \varphi_r[c'/\square]$ are $Free Vars(\mu) \cup \{c, c'\}$ and ρ and ρ' agree on c and c', it follows that $\rho \models \exists Free Vars(\mu)\ (\varphi_l[c/\square] \wedge \varphi_r[c'/\square])$. By the definition of $step(c,\ c')$, we conclude $\rho \models step(c,\ c')$. ☐

The following lemma introduces a formula encoding a complete path of fixed length.

Lemma 8.4. $\rho \models \bigwedge_{0 \leq i < n} step(c_i,\ c_{i+1}) \wedge \not\exists c_{n+1}\ step(c_n,\ c_{n+1})$ iff $\rho(c_0), ..., \rho(c_{n+1})$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path.

Proof. By Lemma 8.3, we have that $\rho(c_i) \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \rho(c_{i+1})$ iff $\rho' \models step(c_i,\ c_{i+1})$, for each $0 \leq i < n$. Further, $\rho(c_0), ..., \rho(c_{n+1})$ is complete, iff there does not exist γ such that $\rho(c_n) \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma$. Again, by Lemma 8.3, that is iff $\rho \models \not\exists c_{n+1}\ step(c_n,\ c_{n+1})$. We conclude that $\rho \models \bigwedge_{0 \leq i < n} step(c_i,\ c_{i+1}) \wedge \not\exists c_{n+1}\ step(c_n,\ c_{n+1})$ iff $\rho(c_0), ..., \rho(c_{n+1})$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path, and we are done. ☐

The following lemma states that $coreach(\varphi)$ actually has the semantic properties its name suggests.

Lemma 8.5. $(\gamma, \rho) \models coreach(\varphi)$ iff for all complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-paths τ starting with γ it is the case that $(\gamma', \rho) \models \varphi$ for some $\gamma' \in \tau$.

Proof. First we prove the direct implication. Assume $(\gamma, \rho) \models coreach(\varphi)$, and let $\tau \equiv \gamma_0, ..., \gamma_n$ be a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path starting with γ. Then let ρ' agree with ρ on $Free Vars(\varphi)$ such that $\rho'(n) = n$ and $\rho'(c_i) = \gamma_i$ for each $0 \leq i \leq n$. According to the definition of $coreach(\varphi)$, we have that

$$(\gamma, \rho') \models \square = c_0 \wedge \bigwedge_{0 \leq i < n} step(c_i,\ c_{i+1}) \wedge \not\exists c_{n+1}\ step(c_n,\ c_{n+1}) \rightarrow \bigvee_{0 \leq i \leq n} \varphi[c_i/\square]$$

Since, $\gamma = \gamma_0$ and $\rho'(c_0) = \gamma_0$, it follows that $\rho' \models \square = c_0$. Further, by Lemma 8.4, since $\rho'(c_0), ..., \rho'(c_n)$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path, it must be the case that

$$\rho' \models \bigwedge_{0 \leq i < n} step(c_i,\ c_{i+1}) \wedge \not\exists c_{n+1}\ step(c_n,\ c_{n+1})$$

$$\overline{coreach}(\varphi) \quad \equiv \quad \forall n \forall a \forall b \; (\exists c \; (\beta(a,b,0,\alpha(c)) \wedge \square = c)$$
$$\wedge \forall i \; (0 \le i \wedge i < n \to \exists c \exists c' \; (\beta(a,b,i,\alpha(c)) \wedge \beta(a,b,i+1,\alpha(c')) \wedge step(c, \; c')))$$
$$\wedge \exists c \; (\beta(a,b,n,\alpha(c)) \wedge \nexists c' \; step(c, \; c'))$$
$$\to \exists i \; (0 \le i \wedge i \le n \wedge \exists c \; (\beta(a,b,i,\alpha(c)) \wedge \varphi[c/\square])))$$

Figure 6: FOL definition of $coreach(\varphi)$

Thus, as $\square$ does not occur in any $\varphi[c_i/\square]$, we conclude that $\rho' \models \bigvee_{0 \le i \le n} \varphi[c_i/\square]$, that is, $\rho' \models \varphi[c_i/\square]$ for some $0 \le i \le n$. By Lemma 8.2, $\rho' \models \varphi[c_i/\square]$ iff $(\gamma_i, \rho') \models \varphi$. Since ρ agrees with ρ' on $FreeVars(\varphi)$, we conclude that $(\gamma_i, \rho) \models \varphi$.

Conversely, assume that if τ is a finite and complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path starting with γ. Then $(\gamma', \rho) \models \varphi$ for some $\gamma' \in \tau$. Let ρ' agree with ρ on $FreeVars(\varphi)$. Then we prove that

$$(\gamma, \rho') \models \square = c_0 \wedge \bigwedge_{0 \le i < n} step(c_i, \; c_{i+1}) \wedge \nexists c_{n+1} \; step(c_n, \; c_{n+1}) \to \bigvee_{0 \le i \le n} \varphi[c_i/\square]$$

Specifically, assume $(\gamma, \rho') \models \square = c_0 \wedge \bigwedge_{0 \le i < n} step(c_i, \; c_{i+1}) \wedge \nexists c_{n+1} \; step(c_n, \; c_{n+1})$. As $\square$ does not occur in any cc_ic_{i+1}, by Lemma 8.4, it follows that $\rho'(c_0), ..., \rho'(c_n)$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path. Further, $(\gamma, \rho') \models \square = c$, implies that $\rho'(c_0), ..., \rho'(c_n)$ starts with γ. Thus, there exists some $0 \le i \le n$ such that $(\rho'(c_i), \rho) \models \varphi$, or equivalently, since ρ and ρ' agree on $FreeVars(\varphi)$, such that $(\rho'(c_i), \rho') \models \varphi$. By Lemma 8.2, $(\rho'(c_i), \rho') \models \varphi$ iff $\rho' \models \varphi[c_i/\square]$. Therefore, we have that $(\gamma, \rho') \models \bigvee_{0 \le i \le n} \varphi[c_i/\square]$. Finally, since ρ' is an arbitrary valuation which agrees with ρ on $FreeVars(\varphi)$, by the definition of $coreach(\varphi)$ we can conclude that $(\gamma, \rho) \models coreach(\varphi)$, and we are done. $\qquad \square$

The following lemma establishes a useful property of $coreach(\varphi)$.

Lemma 8.6.

$$\models coreach(\varphi) \to \varphi \vee (\exists c' \; step(c, \; c') \wedge \forall c' \; (step(c, \; c') \to coreach(\varphi)[c'/\square]))$$

Proof. We prove that if $(\gamma, \rho) \models coreach(\varphi)$ then

$$(\gamma, \rho) \models \varphi \vee (\exists c' \; step(c, \; c') \wedge \forall c' \; (step(c, \; c') \to coreach(\varphi)[c'/\square]))$$

By Lemma 8.5, we have that for all complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-paths τ starting with γ it is the case that $(\gamma'', \rho) \models \varphi$ for some $\gamma'' \in \tau$. We distinguish two cases

- $(\gamma, \rho) \models \varphi$. We are trivially done.
- $(\gamma, \rho) \not\models \varphi$. Then γ must have $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-successors. Indeed, assume the contrary. Then $\tau \equiv \gamma$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path. It follows that $(\gamma, \rho) \models \varphi$, which is a contradiction. Thus, there exists some γ' such that $\gamma \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'$. By Lemma 8.3, that is iff $\rho \models \exists c' \; step(c, \; c')$. Further, let γ' be a $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-successor of γ and τ' a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path starting with γ'. Then, $\gamma\tau$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path starting with γ. Thus, there exists some $\gamma'' \in \gamma\tau'$ such that $(\gamma'', \rho) \models \varphi$. Since $(\gamma, \rho) \not\models \varphi$, it follows that $\gamma'' \in \tau'$. Notice that γ' is an arbitrary configuration and τ' an arbitrary $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path, therefore by Lemma 8.5 and Lemma 8.2, we can conclude that $\rho \models \forall c' \; (step(c, \; c') \to coreach(\varphi)[c'/\square])$. $\qquad \square$

8.2. **Formula Gödelization.** Fig. 6 defines $\overline{coreach}(\varphi)$, the FOL equivalent of $coreach(\varphi)$ using Gödel's β predicate. Formally, we have the following

Lemma 8.7. $\models coreach(\varphi) \leftrightarrow \overline{coreach}(\varphi)$.

Proof. Let us choose some arbitrary but fixed values for n, a and b, and let ρ' such that ρ and ρ' agree on $Var \setminus \{n, a, b\}$ and $\rho'(n) = n$ and $\rho'(a) = a$ and $\rho'(b) = b$. According to the definition of the β predicate, there exists a unique integer sequence $j_0, ..., j_n$ such that $\beta(a, b, i, j_i)$ holds for each $0 \leq i \leq n$. Since α is injective, we distinguish two cases

- there exists some $0 \leq i \leq n$ such that there is not any γ_i with $\alpha(\gamma_i) = j_i$
- there exists a unique sequence $\gamma_0, ..., \gamma_n$ such that $\alpha(\gamma_i) = j_i$ for each $0 \leq i \leq n$.

In the former case, if $i = n$ we get that $\rho' \not\models \exists c \ (\beta(a, b, n, \alpha(c)) \wedge \ \nexists c' \ step(c, \ c')$ while if $0 \leq i < n$ we get that $\rho' \not\models \exists c \exists c' \ (\beta(a, b, i, \alpha(c)) \wedge \beta(a, b, i+1, \alpha(c')) \wedge step(c, \ c'))$ as in both cases we can not pick a value for c. Thus, (γ, ρ') does not satisfy left-hand-side of the implication in $\overline{coreach}(\varphi)$, and we conclude that (γ, ρ') satisfies the implication.

In the later case, we have that there is a unique way of instantiating the existentially quantified variables c and c' in each sub-formula in which they appear, as they are always arguments of the β predicate. Thus, $(\gamma, \rho') \models \exists c \ (\beta(a, b, 0, \alpha(c)) \wedge \square = c)$ iff $\gamma = \gamma_0$. By Lemma 8.4, we have that

$$\rho' \models \ \forall i \ (0 \leq i \wedge i < n \rightarrow \exists c \exists c' \ (\beta(a, b, i, \alpha(c)) \wedge \beta(a, b, i+1, \alpha(c')) \wedge step(c, \ c')))$$
$$\wedge \exists c \ (\beta(a, b, n, \alpha(c)) \wedge \ \nexists c' \ step(c, \ c'))$$

iff $\gamma_0...\gamma_n$ is a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path. Finally, by Lemma 8.2

$$\rho' \models \exists i \ (0 \leq i \wedge i \leq n \wedge \exists c \ (\beta(a, b, i, \alpha(c)) \wedge \varphi[c/\square]))$$

iff $(\gamma_i, \rho') \models \varphi$ for some $0 \leq i \leq n$.

We conclude that (γ, ρ') satisfies the implication in $\overline{coreach}(\varphi)$ iff

- there is no sequence $\gamma_0, ..., \gamma_n$ such that $\alpha(\gamma_i) = j_i$ for each $0 \leq i \leq n$
- the unique sequence $\gamma_0, ..., \gamma_n$ such that $\alpha(\gamma_i) = j_i$ for each $0 \leq i \leq n$ is either not starting at γ, not a complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-path or contains some γ' such that $(\gamma', \rho) \models \varphi$, as ρ and ρ' agree on $Var \setminus \{n, a, b\}$.

According to the property of β, for each sequence $j_0, ..., j_n$ there exist some values for a and b. Since n, a and b are chosen arbitrary, we conclude that $(\gamma, \rho) \models \overline{coreach}(\varphi)$ iff for all complete $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$-paths τ starting at γ, there exists some $\gamma' \in \tau$ such that $(\gamma', \rho) \models \varphi$. By Lemma 8.5, we have that the above iff $(\gamma, \rho) \models coreach(\varphi)$, and we are done. ☐

8.3. **Encoding Semantic Validity in FOL.** Now we can use the *coreach* predicate to encode the semantic validity of a rule $\varphi \Rightarrow^{\forall} \varphi'$ in FOL. Formally

Lemma 8.8. *If* $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$ *then* $\models \varphi \rightarrow coreach(\varphi')$.

Proof. Follows from the definition of semantic validity of $\varphi \Rightarrow^{\forall} \varphi'$ and Lemma 8.5. ☐

8.4. **Relative Completeness.** A matching logic formula ψ is *patternless* iff $\Box$ does not occur in ψ. Then we have the following lemma stating that we can derive one step on all paths

Lemma 8.9. $\mathcal{S}, \mathcal{A} \vdash_{\mathcal{C}} \Box = c \wedge \exists c' \; step(c, \; c') \wedge \psi \Rightarrow^{\forall} \exists c' \; (\Box = c' \wedge step(c, \; c')) \wedge \psi$ *where* ψ *is a patternless formula.*

Proof. We derive the rule by applying the STEP proof rule with the following prerequisites

$$\models \Box = c \wedge \exists c' \; step(c, \; c') \wedge \psi \to \bigvee_{\varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\varphi_l) \; \varphi_l$$

and for each $\varphi_l \Rightarrow^{\exists} \varphi_r \; \in \; \mathcal{S}$ (since $\Box$ does not occur in ψ)

$$\models \exists c'' \; (c'' = c \wedge \exists c' \; step(c, \; c') \wedge \varphi_l[c''/\Box]) \wedge \varphi_r \wedge \psi \to \exists c' \; (\Box = c' \wedge step(c, \; c')) \wedge \psi$$

For the first prerequisite, we have the following (using the definition of $step(c, \; c')$)

$$
\begin{aligned}
& \Box = c \wedge \exists c' \; step(c, \; c') \wedge \psi \\
\to \quad & \Box = c \wedge \exists c' \; step(c, \; c') \\
\leftrightarrow \quad & \Box = c \wedge \exists c' \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\mu) \; (\varphi_l[c/\Box] \wedge \varphi_r[c'/\Box]) \\
\to \quad & \Box = c \wedge \exists c' \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\mu) \; \varphi_l[c/\Box] \\
\to \quad & \Box = c \wedge \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\varphi_l) \; \varphi_l[c/\Box] \\
\to \quad & \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\varphi_l) \; \varphi_l
\end{aligned}
$$

For the second prerequisite, let $\varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}$. Then we have that

$$
\begin{aligned}
& \exists c'' \; (c'' = c \wedge \exists c' \; step(c, \; c') \wedge \varphi_l[c''/\Box]) \wedge \varphi_r \wedge \psi \\
\to \quad & \varphi_l[c/\Box] \wedge \varphi_r \wedge \psi \\
\to \quad & \exists c' \; (\Box = c' \wedge \varphi_l[c/\Box] \wedge \varphi_r[c'/\Box]) \wedge \psi \\
\to \quad & \exists c' \; (\Box = c' \wedge \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} (\varphi_l[c/\Box] \wedge \varphi_r[c'/\Box])) \wedge \psi \\
\to \quad & \exists c' \; (\Box = c' \wedge \bigvee_{\mu \equiv \varphi_l \Rightarrow^{\exists} \varphi_r \in \mathcal{S}} \exists FreeVars(\mu) \; (\varphi_l[c/\Box] \wedge \varphi_r[c'/\Box]) \wedge \psi \\
\to \quad & \exists c' \; (\Box = c' \wedge step(c, \; c')) \wedge \psi
\end{aligned}
$$

and we are done. $\qquad\square$

The following three lemmas show that we can derive a rule stating that all the configurations reaching φ in the transition system actually reach φ.

Lemma 8.10. *If*

$$\mathcal{S}, \mathcal{A} \vdash \Box = c \wedge \exists c' \; step(c, \; c') \wedge \forall c' \; (step(c, \; c') \to coreach(\varphi)[c'/\Box]) \Rightarrow^{\forall} \varphi$$

then $\mathcal{S}, \mathcal{A} \vdash coreach(\varphi) \Rightarrow^{\forall} \varphi$.

Proof. By Lemma 8.6

$$\models coreach(\varphi) \leftrightarrow \varphi \vee (\exists c' \; step(c, \; c') \wedge \forall c' \; (step(c, \; c') \to coreach(\varphi)[c'/\Box]))$$

Thus, by Consequence and Case Analysis, it suffices to derive

$$\mathcal{S}, \mathcal{A} \vdash \varphi \Rightarrow^{\forall} \varphi$$
$$\mathcal{S}, \mathcal{A} \vdash \square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \Rightarrow^{\forall} \varphi$$

The first sequent follows by Reflexivity. The second sequent is part of the hypothesis, and we are done. $\square$

Lemma 8.11.

$$\mathcal{S}, \mathcal{A} \vdash \square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \Rightarrow^{\forall} \varphi$$

Proof. Let μ be the rule we want to derive, namely

$$\square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \Rightarrow^{\forall} \varphi$$

Then $\mathcal{S}, \mathcal{A} \vdash \mu$ follows by Circularity from $\mathcal{S}, \mathcal{A} \vdash_{\{\mu\}} \mu$. Hence, by Transitivity, it suffices to derive the two sequents below

$$\mathcal{S}, \mathcal{A} \vdash_{\{\mu\}} \square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \Rightarrow^{\forall} \varphi'$$
$$\mathcal{S}, \mathcal{A} \cup \{\mu\} \vdash \varphi' \Rightarrow^{\forall} \varphi$$

where $\varphi' \equiv \exists c'\ (\square = c' \wedge step(c,\ c')) \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square])$. The first sequent follows by Lemma 8.9 with $\psi \equiv \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square])$. For the second sequent, by Abstraction with $\{c'\}$ and Consequence with

$$\models \square = c' \wedge step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \rightarrow coreach(\varphi)$$

it suffices to derive $\mathcal{S}, \mathcal{A} \cup \{\mu\} \vdash coreach(\varphi) \Rightarrow^{\forall} \varphi$. Then, by Lemma 8.10, we are left to derive

$$\mathcal{S}, \mathcal{A} \cup \{\mu\} \vdash \square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)[c'/\square]) \Rightarrow^{\forall} \varphi$$

that is, $\mathcal{S}, \mathcal{A} \cup \{\mu\} \vdash \mu$, which trivially follows by Axiom and we are done. $\square$

Lemma 8.12. $\mathcal{S}, \mathcal{A} \vdash coreach(\varphi) \Rightarrow^{\forall} \varphi$.

Proof. By Lemma 8.10, it suffices to derive

$$\mathcal{S}, \mathcal{A} \vdash \square = c \wedge \exists c'\ step(c,\ c') \wedge \forall c'\ (step(c,\ c') \rightarrow coreach(\varphi)) \Rightarrow^{\forall} \varphi$$

which follows by Lemma 8.11. $\square$

Finally, we can establish the main result. Assume $\mathcal{S} \models \varphi \Rightarrow^{\forall} \varphi'$. By Lemma 8.8, we have that $\models \varphi \rightarrow coreach(\varphi')$. Further, by Lemma 8.12, we have that $\mathcal{S} \vdash coreach(\varphi') \Rightarrow^{\forall} \varphi'$. Then, by Consequence, it follows that $\mathcal{S} \vdash \varphi \Rightarrow^{\forall} \varphi'$.

9. Related Work

Using Hoare logic [20] to prove concurrent programs correct dates back to Owicki and Gries [21]. In the rely-guarantee method proposed by Jones [22] each thread relies on some properties being satisfied by the other threads, and in its turn, offers some guarantees on which the other threads can rely. O'Hearn [23] advances a Separation Hypothesis in the context of separation logic [24] to achieve compositionality: the state can be partitioned into separate portions for each process and relevant resources, respectively, satisfying certain invariants. More recent research focuses on improvements over both of the above methods and even combinations of them (e.g., [25–28]).

The satisfaction of all-path-reachability rules can also be understood intuitively in the context of temporal logics. Matching logic formulae can be thought of as state formulae, and reachability rules as temporal formulae. Assuming CTL^* on finite traces, the semantics rule $\varphi \Rightarrow^\exists \varphi'$ can be expressed as $\varphi \to E \bigcirc \varphi'$, while an all-path reachability rule $\varphi \Rightarrow^\forall \varphi'$ can be expressed as $\varphi \to A \diamond \varphi'$. However, unlike in CTL^*, the φ and φ' formulae of reachability rules $\varphi \Rightarrow^\exists \varphi'$ or $\varphi \Rightarrow^\forall \varphi'$ share their free variables. Thus, existing proof systems for temporal logics (e.g., the CTL^* one by Pnueli and Kesten [29]) are not directly comparable with our approach. In general, CTL^* formulae only have atomic predicates, while in our approach we consider only a specific temporal structure ($\varphi \to A \diamond \varphi'$), but the matching logic formulae φ and φ' can express more complex state properties.

This approach has been extended to handle safety properties of distributed system, specified as rewrite theories [30]. Specifically, it can verify that φ is an invariant, or a co-invariant, of a non-terminating distributed system by checking that φ holds after each application of the STEP proof rule. The soundness of the approach is proved by transforming the underlying transition system from non-terminating to terminating by adding a transition from each state to a newly-introduced end state (where φ does not hold).

Bae et al. [31], Rocha and Meseguer [32], and Rocha et al. [33] use narrowing to perform symbolic reachability analysis in a transition system associated to a unconditional rewrite theory for the purposes of verification. There are two main differences between their work and ours. First, they express state predicates in equational theories. Matching logic is more general, being first-order logic over a model of configurations $\mathcal{T}$. Consequently, the STEP proof rule takes these issues into account when considering the successors of a state. Second, they use rewrite systems for symbolic model checking. Our work is complementary, in the sense that we use the operational semantics for program verification, and check properties more similar to those in Hoare logic.

Dynamic logic [34] adds modal operators to FOL to embed program fragments within specifications. For example, $\psi \to [\mathbf{s}]\psi'$ means "after executing $\mathbf{s}$ in a state satisfying ψ, a state may be reached which satisfies ψ'". KeY [35] offers automatic verification for Java based on dynamic logic. Dynamic logic has been extended to support reasoning about cyber-physical systems, where the transitions may have a continuous or probabilistic component [36]. Matching logic also combines programs and specifications for static properties, but dynamic properties are expressed in reachability logic which has a language-independent proof system that works with any operational semantics, while dynamic logic still requires language-specific proof rules (e.g., invariant rules for loops).

Language-independent proof systems. A first proof system is introduced in [6], while [5] presents a mechanical translation from Hoare logic proof derivations for IMP into derivations in the proof system. The CIRCULARITY proof rule is introduced in [4]. Finally, [7] supports operational semantics given with conditional rules, like small-step and big-step. All these previous results can only be applied to deterministic programs. The coinductive nature of this proof system was clarified in [37], which presents a coinduction principle that captures the essence of the CIRCULARITY rule, and which is sufficient for program verification when used within a system such as ZFC or Coq that already supports basic logical reasoning and working with collections of configurations.

10. Conclusion and Future Work

This paper introduces a sound and (relatively) complete language-independent proof system which derives program properties holding along all execution paths (capturing partial correctness for non-deterministic programs), directly from an operational semantics. The proof system separates reasoning about deterministic language features (via the operational semantics) from reasoning about non-determinism (via the proof system). Thus, all we need in order to verify programs in a language is an operational semantics for the respective language.

We believe that existing techniques such as rely-guarantee and concurrent separation logic could be used in conjunction with our proof system to achieve semantically grounded and compositional verification.

Our approach handles operational semantics given with unconditional rules, like in the $\mathbb{K}$ framework, PLT-Redex, and CHAM, but it cannot handle operational semantics given with conditional rules, like big-step and small-step (rules with premises). Extending the presented results to work with conditional rules would boil down to extending the STEP proof rule, which derives the fact that φ reaches φ' in one step along all execution paths. Such an extended STEP would have as prerequisites whether the left-hand side of a semantics rule matches (like the existing STEP) and additionally whether its premises hold. The second part would require an encoding of reachability in first-order logic, which is non-trivial and most likely would result in a first-order logic over a richer model than $\mathcal{T}$. The difficulty arises from the fact that STEP must ensure all successors of φ are in φ'. Thus, this extension is left as future work.

References

[1] Nipkow, T.: Winskel is (almost) right: Towards a mechanized semantics textbook. Formal Aspects of Computing **10** (1998) 171–186

[2] Jacobs, B.: Weakest pre-condition reasoning for Java programs with JML annotations. J. Logic and Algebraic Programming **58**(1-2) (2004) 61–88

[3] Appel, A.W.: Verified software toolchain. In: ESOP. Volume 6602 of LNCS. (2011) 1–17

[4] Roșu, G., Ștefănescu, A.: Checking reachability using matching logic. In: OOPSLA, ACM (2012) 555–574

[5] Roșu, G., Ștefănescu, A.: From Hoare logic to matching logic reachability. In: FM. Volume 7436 of LNCS. (2012) 387–402

[6] Roșu, G., Ștefănescu, A.: Towards a unified theory of operational and axiomatic semantics. In: ICALP. Volume 7392 of LNCS. (2012) 351–363

[7] Roşu, G., Ştefănescu, A., Ciobâcă, S., Moore, B.M.: One-path reachability logic. In: LICS, IEEE (2013) 358–367

[8] Roșu, G., Ellison, C., Schulte, W.: Matching logic: An alternative to Hoare/Floyd logic. In: AMAST. Volume 6486 of LNCS. (2010) 142–162

[9] Roșu, G., Șerbănuță, T.F.: An overview of the K semantic framework. J. Logic and Algebraic Programming **79**(6) (2010) 397–434

[10] Ştefănescu, A., Ciobâcă, S., Mereuţă, R., Moore, B.M., Şerbănuţă, T.F., Roşu, G.: All-path reachability logic. In: RTA. Volume 8560 of LNCS. (July 2014) 425–440

[11] Roşu, G.: Matching logic. Logical Methods in Computer Science **13**(4) (December 2017) 1–61

[12] Roşu, G.: Matching logic — extended abstract. In: RTA. Volume 36 of Leibniz International Proceedings in Informatics (LIPIcs)., Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2015) 5–21

[13] Ellison, C., Roșu, G.: An executable formal semantics of C with applications. In: POPL, ACM (2012) 533–544

[14] Hathhorn, C., Ellison, C., Rosu, G.: Defining the undefinedness of C. In: PLDI, ACM (2015) 336–345

[15] Felleisen, M., Findler, R.B., Flatt, M.: Semantics Engineering with PLT Redex. MIT (2009)

[16] Berry, G., Boudol, G.: The chemical abstract machine. Theoretical Computer Science **96**(1) (1992) 217–248

[17] Matthews, J., Findler, R.B.: An operational semantics for Scheme. JFP **18**(1) (2008) 47–86

[18] Filaretti, D., Maffeis, S.: An executable formal semantics of PHP. In: ECOOP. Volume 8586 of LNCS. (2014) 567–592

[19] de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: TACAS. Volume 4963 of LNCS. (2008) 337–340

[20] Hoare, C.A.R.: An axiomatic basis for computer programming. Communications of the ACM **12**(10) (1969) 576–580

[21] Owicki, S.S., Gries, D.: Verifying properties of parallel programs: An axiomatic approach. Communications of the ACM **19**(5) (1976) 279–285

[22] Jones, C.B.: Specification and design of (parallel) programs. In Mason, R.E.A., ed.: Information Processing 1983: World Congress Proceedings, Elsevier (1984) 321–332

[23] O'Hearn, P.W.: Resources, concurrency, and local reasoning. Theoretical Computer Science **375**(1-3) (2007) 271–307

[24] Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: LICS, IEEE (2002) 55–74

[25] Feng, X.: Local rely-guarantee reasoning. In: POPL, ACM (2009) 315–327

[26] Vafeiadis, V., Parkinson, M.J.: A marriage of rely/guarantee and separation logic. In: CONCUR. Volume 4703 of LNCS. (2007) 256–271

[27] Reddy, U.S., Reynolds, J.C.: Syntactic control of interference for separation logic. In: POPL, ACM (2012) 323–336

[28] Hayman, J.: Granularity and concurrent separation logic. In: CONCUR. Volume 6901 of LNCS. (2011) 219–234

[29] Pnueli, A., Kesten, Y.: A deductive proof system for CTL^*. In: CONCUR. Volume 2421 of LNCS. (2002) 24–40

[30] Skeirik, S., Ştefănescu, A., Meseguer, J.: A constructor-based reachability logic for rewrite theories. In: LOPSTR. Volume 10855 of LNCS., Springer (2017) 201–217

[31] Bae, K., Escobar, S., Meseguer, J.: Abstract logical model checking of infinite-state systems using narrowing. In: RTA. (2013) 81–96

[32] Rocha, C., Meseguer, J.: Proving safety properties of rewrite theories. In: CALCO. Volume 6859 of LNCS. (2011) 314–328

[33] Rocha, C., Meseguer, J., Muñoz, C.A.: Rewriting modulo SMT and open system analysis. In: WRLA. Volume 8663 of LNCS., Springer (2014) 247–262

[34] Harel, D., Kozen, D., Tiuryn, J.: Dynamic logic. In: Handbook of Philosophical Logic. (1984) 497–604

[35] Beckert, B., Hähnle, R., Schmitt, P.H.: Verification of Object-oriented Software: The KeY Approach. Springer-Verlag (2007)

[36] Madeira, A., Neves, R., Martins, M.A., Barbosa, L.S.: A dynamic logic for every season. In: SBMF. Volume 8941 of LNCS., Springer (2014) 130–145

[37] Moore, B., Peña, L., Roşu, G.: Program verification by coinduction. In: 27th European Symposium on Programming (ESOP). (April 2018)