

Program Equivalence by Circular Reasoning

Dorel Lucanu, Vlad Rusu

► **To cite this version:**

Dorel Lucanu, Vlad Rusu. Program Equivalence by Circular Reasoning. Integrated Formal Methods, Jun 2013, Turku, Finland. Springer, LNCS 7940, pp.362-377, 2013, Lecture Notes in Computer Science. <hal-00820871>

HAL Id: hal-00820871

<https://hal.inria.fr/hal-00820871>

Submitted on 6 May 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Program Equivalence by Circular Reasoning

Dorel Luca¹ and Vlad Rusu²

¹ Al. I. Cuza University of Iași, Romania dlucanu@info.uaic.ro

² Inria Lille Nord-Europe, France Vlad.Rusu@inria.fr

Abstract. We propose a logic and a deductive system for stating and automatically proving the equivalence of programs in deterministic languages having a rewriting-based operational semantics. The deductive system is circular in nature and is proved sound and weakly complete; together, these results say that, when it terminates, our system correctly solves the program-equivalence problem as we state it. We show that our approach is suitable for proving the equivalence of both terminating and non-terminating programs, and also the equivalence of both concrete and symbolic programs. The latter are programs in which some statements or expressions are symbolic variables. By proving the equivalence between symbolic programs, one proves in one shot the equivalence of (possibly, infinitely) many concrete programs obtained by replacing the variables by concrete statements or expressions. A prototype of the proof system for a particular language was implemented and can be tested on-line.

1 Introduction

In this paper we propose a formal notion of program equivalence, together with a logic for expressing this notion and a deductive system for automatically proving it. Programs can belong to any deterministic language whose semantics is specified by a set of rewrite rules. The equivalence we consider is a form of weak bisimulation, allowing several instructions of one program to be matched by several instructions of the other one. The proof system is circular: its conclusions can be re-used as hypotheses in a controlled way. It is not guaranteed to terminate, but when it does terminate, our proof system correctly solves the program-equivalence problem as stated, thanks to its soundness and weak completeness properties. These are informally presented below and are formalised and proved in the paper.

The proposed framework is also suitable for proving the equivalence of *symbolic programs*. These are programs in which some expressions and/or statements are *symbolic variables*, which denote sets of concrete programs obtained by substituting the symbolic variables by concrete expressions and/or statements. Thus, by proving the equivalence between symbolic programs, one proves in just one shot the equivalence of (possibly, infinitely) many concrete programs, which has applications in the verification of certain classes of compilers/translators. Here is an example of equivalent symbolic programs.

Example 1. Assume that we want to translate between a language that has `for`-loops into a language that only has `while`-loops. This amounts to translating the symbolic program in the left-hand side to the one in the right-hand side.

$$\text{for } I \text{ from } A \text{ to } B \text{ do} \{ S \} \quad I = A ; \text{while } I \leq B \text{ do } \{ S ; I = I + 1 \}$$

Their symbolic variables I, A, B, S can be matched by, respectively, any identifier (I), arithmetical expression (A, B), and program statement (S). If we prove the equivalence between these two symbolic programs (as we shall do in this paper as an illustrative example) then we also prove that every concrete instance of the `for`-loop is equivalent to its translation to a concrete `while`-loop (or vice-versa). Nonterminating programs can be proved equivalent as well, e.g. by replacing the test $I \leq B$ with $\text{not}(I = B)$ and by assuming nonterminating `for` loops when $A > B$, some instances of the above two symbolic programs are nonterminating.

In the rest of the paper we often refer to symbolic programs just as “programs”.

A typical use of our program-equivalence framework consists in:

1. defining the operational semantics of a programming language, say, \mathcal{L} ;
2. defining a language \mathcal{L}^{sym} , which extends the syntax of and semantics of \mathcal{L} , such that the programs in \mathcal{L}^{sym} are exactly the symbolic programs of \mathcal{L} ;
3. applying our deductive system to check the equivalence of programs in \mathcal{L}^{sym} .

Running the deductive system amounts essentially to executing the semantics of \mathcal{L}^{sym} on pairs of \mathcal{L}^{sym} -programs. This may lead to any of the following outcomes:

- termination with success, in which case the programs given as input to the deductive system are equivalent, due to the deductive system’s *soundness*;
- termination with failure, in which case the programs given as input to the deductive system are not equivalent, due to the system’s *weak completeness*;
- non-termination, in which case nothing can be concluded about equivalence.

Non-termination is inherent in any sound automatic system for proving program equivalence, because the equivalence problem is undecidable. We show, however, that our system terminates when the programs given to it as inputs terminate, and also when they do not terminate but behave in a certain regular way (by infinitely repeating pairs of so-called *observationally equivalent configurations*).

Contributions A logic and a proof system suitable for stating and proving the equivalence of concrete and of symbolic programs, as well as that of terminating and non-terminating ones. Programs can be written in any deterministic language that has a formal operational semantics based on term rewriting. We prove the soundness and weak completeness of our proof system, which ensure that the system correctly solves the program equivalence problem as we state it. A prototype implementation of the proposed deductive system is also presented.

Related Work An exhaustive bibliography on the program-equivalence problem is outside the scope of this paper, as this problem is even older than the program-verification problem. Among the recent works perhaps the closest to ours is [1]. They also deal with the equivalence of parameterised programs (symbolic, in our terminology) and define equivalence in terms of bisimulation.

Their approach is, however, very different from ours. One major difference lies in the models of programs: [1] use CFGs (control flow graphs) of programs, while we use the operational semantics of languages. CFGs are more restricted, e.g., they are not well adapted to recursive or object-oriented programs, whereas operational semantics do not have these limitations. Of course, our advantage will only become apparent when we actually apply our approach to such programs.

Other closely related recent works are [2,3,4]. The first one targets programs that include recursive procedures, the second one exploits similarities between single-threaded programs in order to prove their equivalence, and the third one extends the latter to multi-threaded programs. They use operational semantics (of a specific language) and proof systems, and formally prove their proof system's soundness. In [2] they make a useful classification of equivalence relations used in program-equivalence research, and use these relations in their work.

However, all the relations classified in [2] are of an input/output nature: for given (sequences of) inputs, programs generate equal (sequences of) outputs and/or do not terminate. Such relations are well adapted for concrete programs with inputs and outputs, but not to symbolic programs with symbolic statements, for which a clear input-output relation may not exist. Indeed, symbolic statements may denote arbitrary concrete statements - including ones that do not perform input/output - actually, when symbolic programs are concerned, one cannot even rely on the existence of inputs and outputs. One may rely, however, on the observations of the effects of symbolic statements on the program's environment (e.g., values of variables). Our notion of weak bisimulation (up to a certain observation relation) allows this, both for finitely and for infinitely many repeated observations. We also show that some of the relations from [2] can be encoded in our relation by adding information to the program environment.

Many works on program equivalence arise from the verification of compilation in a broad sense. At one end there is full compiler verification [5], and at the other end, the so-called translation validation, i.e., the individual verification of each compilation [6] (we only cite two of the most relevant recent works). As also observed by [1], symbolic program verification can also be used for certain compilers, in which one proves the equivalence of each basic instruction pattern from the source language with its translation in the target language. The application of this observation to the verification of a compiler (from another project we are involved in) is ongoing and will be presented in another paper.

Several other works have targeted specific classes of languages: functional [7], microcode [8], CLP [9]. In order to be less language-specific some works advocate the use of intermediate languages, such as [10], which works on the Boogie intermediate language. And finally, only a few approaches, among which [5,8], deal with real-life language and industrial-size programs in those languages. This is in contrast to the equivalence checking of hardware circuits, which has entered the mainstream industrial practice (see, e.g., [11] for a survey on this topic).

Our proof system is inspired by that of *circular coinduction* [12], which allows one to prove equalities of data structures such as infinite streams and regular expressions. A notable difference between the present approach and [12] is that

our specifications are essentially rewrite theories (meant to define the semantics of programming languages), whereas those of [12] are behavioural equational theories, a special class of equational specifications with visible and hidden sorts.

The rest of the paper is organised as follows. Section 2 presents our running example: IMP, a simple imperative language and its definition in \mathbb{K} [13]. \mathbb{K} is a formal framework for defining operational semantics of programming languages.

Our approach is, however, independent of the \mathbb{K} framework and the IMP language; hence, we present a general, abstract notion of language definition in Section 3, and show how the \mathbb{K} definition of IMP is an instance of that notion.

Section 4 contains our proposed definition for program equivalence, and Section 5 gives the syntax and semantics of a logic capturing the chosen equivalence.

Section 6 introduces two operations on formulas of the logic (derivatives and conjunction) which are used in our circular proof system for formula validity.

The proof system itself is presented in Section 7, together with its soundness and weak completeness results. The results say that, when it terminates, the proof system correctly answers to the question of whether its input (which is a set of formulas in our program-equivalence logic) denotes equivalent programs.

The conclusion and future work are presented in Section 8. Finally, formal proofs of the results in the paper are given in the technical report that can be found at <http://hal.archives-ouvertes.fr/hal-00744374/> .

2 A Simple Imperative Language and its Semantics in \mathbb{K}

The language we are using as running example is IMP, a simple imperative language intensively used in research papers. A full \mathbb{K} definition of it can be found in [13]. The syntax of IMP is described in Figure 1 and is mostly self-explained. The attribute (given as an annotation) *strict* from the syntax means the arguments of the annotated expression/statement are evaluated before the expression/statement itself is evaluated/executed. If the attribute has as arguments a list of natural numbers, then only the arguments in positions specified by the list are evaluated before the expression/statement. The *strict* attribute is actually syntactic sugar for a set of \mathbb{K} rules, briefly presented later in the section.

The *configuration* of an IMP program consists of code to be executed and an environment mapping identifiers to integers. In \mathbb{K} , this is written as a nested structure of *cells*: here, a top cell **cfg**, having a cell **k** and a cell **env** (see Figure 2).

The cell **k** includes the code to be executed, represented as a list of computation tasks $C_1 \curvearrowright C_2 \curvearrowright \dots$, meaning that first C_1 will be executed, then C_2 , etc. Computation tasks are typically the evaluation of statements and expressions. The cell **env** is an environment that binds the program variables to values; such a binding is written as a multiset of bindings of the form, e.g., $\mathbf{a} \mapsto 3$.

The semantics of IMP is given by a set of rules (see Figure 3) that say how the configuration evolves when the first computation task (statement or instruction) from the **k** cell is executed. The dots in a cell mean that the rest of the cell remains unchanged. Except for the conjunction, negation, and **if** statement, the semantics of each operator and statement is described by exactly one rule.

$Int ::= \text{domain of integer numbers (including operations)}$
 $Bool ::= \text{domain of boolean constants (including operations)}$
 $Id ::= \text{domain of identifiers}$
 $AExp ::= Int \mid Id$ $BExp ::= Bool$
 $\mid AExp / AExp$ [strict] $\mid AExp \leq AExp$ [strict]
 $\mid AExp * AExp$ [strict] $\mid \text{not } BExp$ [strict]
 $\mid AExp + AExp$ [strict] $\mid BExp \text{ and } BExp$ [strict(1)]
 $\mid (AExp)$ $\mid (BExp)$

$Stmt ::= \text{skip} \mid Stmt ; Stmt$ $\mid \{ Stmt \}$
 $\mid Id = AExp$ $\mid \text{while } BExp \text{ do } Stmt$
 $\mid \text{if } BExp \text{ then } Stmt$ $\mid \text{for } Id \text{ from } AExp \text{ to } AExp$
 $\text{else } Stmt$ [strict(1)] $\text{do } Stmt$ [strict(2,3)]

$Code ::= Id \mid Int \mid Bool \mid AExp \mid BExp \mid Stmt \mid Code \curvearrowright Code$

Fig. 1. \mathbb{K} Syntax of IMP

$Cfg ::= \langle \langle Code \rangle_k \langle Map \rangle_{env} \rangle_{cfg}$

Fig. 2. \mathbb{K} Configuration of IMP

In Figure 3, the operations $lookup : Map \times Id \rightarrow Int$ and $update : Map \times Id \times Int \rightarrow Map$ are part of the domain of maps and have the usual meanings: $lookup$ returns the value of an identifier in a map, and $update$ modifies the map by adding (or, if it exists, by updating) the binding of an identifier to a value.

In addition to the rules in Figure 3 there are rules induced by the strictness of some statements. For example, the `if` statement is strict only in the first argument, meaning that this argument is evaluated before the `if` statement. This amounts to the following rules (automatically generated by the \mathbb{K} tool):

$$\langle \langle \text{if } BE \text{ then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{cfg} \Rightarrow \langle \langle BE \curvearrowright \text{if } \square \text{ then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{cfg}$$

$$\langle \langle B \curvearrowright \text{if } \square \text{ then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{cfg} \Rightarrow \langle \langle \text{if } B \text{ then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{cfg}$$

where BE ranges over $BExp \setminus \{false, true\}$, B ranges over $\{false, true\}$, and \square is a special variable destined to receive the value of BE once it is computed.

3 A Generic Notion of Language Definition

Our program-equivalence approach is independent of the formal framework used for defining languages as well as from the languages being defined. We thus propose a general notion of language definition and illustrate it later in the section on the \mathbb{K} definition of IMP. We assume the reader is familiar with the basics of algebraic specification and rewriting. A language \mathcal{L} is defined by:

1. A many-sorted algebraic signature Σ , which includes at least a sort Cfg for configurations and a subsignature Σ^{Bool} for Booleans with their usual

$$\begin{aligned}
\langle\langle I_1 + I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle I_1 +_{\text{Int}} I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle I_1 * I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle I_1 *_{\text{Int}} I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle I_1 / I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \wedge I_2 \neq 0 &\Rightarrow \langle\langle I_1 /_{\text{Int}} I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle I_1 \leq I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle I_1 \leq_{\text{Int}} I_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{true and } B \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle B \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{false and } B \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \text{false} \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{not true} \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \text{false} \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{not false} \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \text{true} \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{skip} \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle S_1; S_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle S_1 \curvearrowright S_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \{ S \} \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle S \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{if true then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle S_1 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{if false then } S_1 \text{ else } S_2 \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle S_2 \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{while } B \text{ do } S \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \\
&\langle\langle \text{if } B \text{ then } \{ S; \text{while } B \text{ do } S \} \text{ else skip} \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle \text{for } X \text{ from } I_1 \text{ to } I_2 \text{ do } S \dots \rangle_k \dots \rangle_{\text{cfg}} &\Rightarrow \\
&\langle\langle X = I_1; \text{if } X \leq I_2 \text{ then } \{ S; \text{for } X \text{ from } I_1 + 1 \text{ to } I_2 \text{ do } S \} \text{ else skip} \dots \rangle_k \dots \rangle_{\text{cfg}} \\
\langle\langle X \dots \rangle_k \langle \text{Env} \rangle_{\text{env}} \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \text{lookup}(\text{Env}, X) \dots \rangle_k \langle \text{Env} \rangle_{\text{env}} \dots \rangle_{\text{cfg}} \\
\langle\langle X = I \dots \rangle_k \langle \text{Env} \rangle_{\text{env}} \dots \rangle_{\text{cfg}} &\Rightarrow \langle\langle \dots \rangle_k \langle \text{update}(\text{Env}, X, I) \rangle_{\text{env}} \dots \rangle_{\text{cfg}}
\end{aligned}$$

Fig. 3. \mathbb{K} Semantics of IMP

constants and operations. Σ may also include other subsignatures for other data sorts, depending on the language \mathcal{L} (e.g., integers, identifiers, lists, maps, ...). Let Σ^{Data} denote the subsignature of Σ consisting of all data sorts and their operations. We assume that the sort Cfg and the syntax of \mathcal{L} are not data, i.e., they are defined in $\Sigma \setminus \Sigma^{\text{Data}}$, and that terms of sort Cfg have exactly one subterm denoting statements (which are programs in the syntax of \mathcal{L}) remaining to be executed. Let T_Σ denote the Σ -algebra of ground terms and $T_{\Sigma, s}$ denote the set of ground terms of sort s . Given a sort-wise infinite set of variables Var , let $T_\Sigma(\text{Var})$ denote the free Σ -algebra of terms with variables, $T_{\Sigma, s}(\text{Var})$ denote the set of terms of sort s with variables, and $\text{var}(t)$ denote the set of variables occurring in the term t .

2. A Σ -algebra \mathcal{T} . Let \mathcal{T}_s denote the elements of \mathcal{T} that have the sort s ; the elements of \mathcal{T}_{Cfg} are called *configurations*. \mathcal{T} interprets the data sorts (those included in the subsignature Σ^{Data}) according to some Σ^{Data} -algebra \mathcal{D} . \mathcal{T} interprets the non-data sorts (statements) as ground terms over the signature

$$(\Sigma \setminus \Sigma^{\text{Data}}) \cup \bigcup_{d \in \text{sorts}(\Sigma^{\text{Data}})} \mathcal{D}_d \tag{1}$$

where \mathcal{D}_d denotes the carrier set of the sort d in the algebra \mathcal{D} , and the elements of \mathcal{D}_d are added to the signature $\Sigma \setminus \Sigma^{\text{Data}}$ as constants of sort d . That is, a language is parametric in the way its data are implemented; it

just assumes there is such an implementation Σ^{Data} . This is important for technical reasons (implementing unification by matching, discussed below). Any *valuation* $\rho : Var \rightarrow \mathcal{T}$ is extended to a (homonymous) Σ -algebra morphism $\rho : T_\Sigma(Var) \rightarrow \mathcal{T}$. The interpretation of a ground term t in \mathcal{T} is denoted by \mathcal{T}_t . If $b \in T_{\Sigma, Bool}(Var)$ then we write $\rho \models b$ iff $\rho(b) = \mathcal{D}_{true}$. For simplicity, we often write in the sequel *true*, *false* instead of \mathcal{D}_{true} , \mathcal{D}_{false} .

3. A set \mathcal{S} of rewrite rules, whose definition is given later in the section.

We explain these concepts on the IMP example. Each nonterminal from the syntax (*Int*, *Bool*, *AExp*, ...) is a sort in Σ . Each production from the syntax defines an operation in Σ ; for instance, the production $AExp ::= AExp + AExp$ defines the operation $_ + _ : AExp \times AExp \rightarrow AExp$. These operations define the constructors of the result sort. For the configuration sort *Cfg*, the only constructor is $\langle \langle _ \rangle_k \langle _ \rangle_{env} \rangle_{cfg} : Code \times Map_{Id, Int} \rightarrow Cfg$. The expression $\langle \langle X = I \curvearrowright C \rangle_k \langle Env \rangle_{env} \rangle_{cfg}$ is a term of $T_{Cfg}(Var)$, where X is a variable of sort *Id*, I is a variable of sort *Int*, C is a variable of sort *Code* (the rest of the computation), and Env is a variable of sort $Map_{Id, Int}$ (the rest of the environment). The data algebra \mathcal{D} interprets *Int* as the set of integers, the operations like $+_{Int}$ (cf. Figure 3) as the corresponding usual operation on integers, *Bool* as the set of Boolean values $\{false, true\}$, the operation like \wedge as the usual Boolean operations, the sort $Map_{Id, Int}$ as the multiset of maps $X \mapsto I$, where X ranges over identifiers *Id* and I over the integers. The fact that maps are modified only by the *update* operation ensures that each identifier is bound to at most one integer value. The other sorts, *AExp*, *BExp*, *Stmt*, and *Code*, are interpreted in the algebra \mathcal{T} as ground terms over a modification of the form (1) of the signature Σ , in which data subterms are replaced by their interpretations in \mathcal{D} . For instance, the term `if 1 >Int 0 then skip else skip` is interpreted in \mathcal{T} as `if true then skip else skip`, since \mathcal{D} interprets $1 >_{Int} 0$ as \mathcal{D}_{true} (= *true*).

Definition 1 (pattern [14]). A pattern is an expression of the form $\pi \wedge b$, where $\pi \in T_{\Sigma, Cfg}(Var)$ are basic patterns, $b \in T_{\Sigma, Bool}(Var)$, and $var(b) \subseteq var(\pi)$. If $\gamma \in \mathcal{T}_{Cfg}$ and $\rho : Var \rightarrow \mathcal{T}$ we write $(\gamma, \rho) \models \pi \wedge b$ for $\gamma = \rho(\pi)$ and $\rho \models b$.

A basic pattern π defines a set of (concrete) configurations, and the condition b gives additional constraints these configurations must satisfy. In [14] patterns are encoded as FOL formulas, hence the conjunction notation $\pi \wedge b$. In this paper we keep the notation but separate basic patterns from constraining formulas.

We identify basic patterns π with patterns $\pi \wedge true$. Examples of patterns are $\langle \langle I_1 + I_2 \curvearrowright C \rangle_k \langle Env \rangle_{env} \rangle_{cfg}$ and $\langle \langle I_1 / I_2 \curvearrowright C \rangle_k \langle Env \rangle_{env} \rangle_{cfg} \wedge I_2 \neq 0$.

Definition 2 (semantical rule and transition system). A rule is a pair of patterns of the form $l \wedge b \Rightarrow r$ (note that r is the pattern $r \wedge true$). Any set \mathcal{S} of rules defines a labelled transition system $(\mathcal{T}_{Cfg}, \Rightarrow_{\mathcal{S}}^{\mathcal{T}})$ such that $\gamma \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'$ iff there are $(l \wedge b \Rightarrow r) \in \mathcal{S}$ and $\rho : Var \rightarrow \mathcal{T}$ such that $(\gamma, \rho) \models l \wedge b$ and $(\gamma', \rho) \models r$.

A configuration γ is *final* if its program subterm is empty. A configuration γ is a *deadlock* if it is not final and there is no configuration γ' such that $\gamma \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'$.

Deadlocks are erroneous program terminations, e.g., division-by-zero attempts. A language is *deterministic* if its transition system $(\mathcal{T}, \Rightarrow_{\mathcal{S}}^{\mathcal{T}})$ is deterministic.

Assumption 1 *We assume that the transition system $(\mathcal{T}, \Rightarrow_{\mathcal{S}}^{\mathcal{T}})$ is deterministic.*

We shall be using unification in our program-equivalence deductive system. We call *symbolic unifier* of two terms t_1, t_2 any substitution $\sigma : \text{var}(t_1) \uplus \text{var}(t_2) \rightarrow T_{\Sigma}(Z)$ for some set Z of variables such that $t_1\sigma = t_2\sigma$. We call a *concrete unifier* of terms t_1, t_2 any valuation $\rho : \text{var}(t_1) \uplus \text{var}(t_2) \rightarrow \mathcal{T}$ such that $t_1\rho = t_2\rho$.

Assumption 2 *For all rules $(l \wedge b \Rightarrow r) \in \mathcal{S}$ and all patterns $\pi \in T_{\Sigma, \text{Cfg}}(\text{Var})$ with $\text{var}(l) \cap \text{var}(\pi) = \emptyset$, there is a finite, possibly empty set $U(\pi, l)$ of symbolic unifiers of π and l , which satisfy the property that for all concrete unifiers ρ of π and l , there exist substitutions $\sigma \in U(\pi, l)$ and valuations η such that $\sigma\eta = \rho$.*

In related work [15] we prove that the above assumption can always be satisfied, by implementing unification with the rules of \mathcal{L} by the *matching* with the rules of a language \mathcal{L}^{sym} , which extends the definition of \mathcal{L} such that the symbolic execution of programs in \mathcal{L} is the usual execution of programs in \mathcal{L}^{sym} . We illustrate how this is done via an example; other examples follow in the paper.

Example 2. Consider the pattern $\langle\langle \text{if } B \text{ then } S_1 \text{ else } S_2 \rangle_k, \langle M \rangle_{\text{env}}\rangle_{\text{cfg}}$ of sort *Cfg*, where B is a variable of sort *Bool* and S_1, S_2 are variables of sort *Stmt*, and the rule $\langle\langle \text{if true then } S'_1 \text{ else } S'_2 \rangle_k \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow \langle\langle S'_1 \curvearrowright S \rangle_k \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}}$. Here we have filled in the "... " from Figure 3 with actual variables, and the rule's variables were chosen so that they are distinct from those in the formula. Let π denote the basic pattern and l the left-hand side of the rule. The set $U(\pi, l)$ is a singleton given by the substitution $\sigma = (B \mapsto \text{true}, S'_1 \mapsto S_1, S'_2 \mapsto S_2, M' \mapsto M)$. On the other hand, l does not match π because the constant leaf *true* of l does not match the variable B in π . However, the rule can be equivalently rewritten as

$$\langle\langle \text{if } B' \text{ then } S'_1 \text{ else } S'_2 \rangle_k \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \wedge B' = \text{true} \Rightarrow \langle\langle S'_1 \curvearrowright S \rangle_k \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}}$$

and now, there is match between the configuration l' from the left-hand side of the new rule and π , i.e., $(B' \mapsto B, S'_1 \mapsto S_1, S'_2 \mapsto S_2, M' \mapsto M)$. This match, combined with the condition $B' = \text{true}$, amount to the above symbolic unifier σ .

4 Defining Program Equivalence

We define in this section our notion of program equivalence. We base our definition on the transition system $(\mathcal{T}_{\text{Cfg}}, \Rightarrow_{\mathcal{S}}^{\mathcal{T}})$, whose states \mathcal{T}_{Cfg} are configurations, and $\Rightarrow_{\mathcal{S}}^{\mathcal{T}}$ is the transition relation defined in the previous section (Definition 2). Our goal is to have a definition of equivalence that is equally suitable for terminating programs and non-terminating ones and for symbolic and concrete ones.

A natural approach (already chosen by [1]) is to use *strong bisimulation*: a symmetrical relation $R \subseteq \mathcal{T}_{\text{Cfg}} \times \mathcal{T}_{\text{Cfg}}$ is a strong bisimulation if for all $(\gamma_1, \gamma_2) \in R$, when $\gamma_1 \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'_1$, there is a transition $\gamma_2 \Rightarrow_{\mathcal{S}}^{\mathcal{T}} \gamma'_2$ such that $(\gamma'_1, \gamma'_2) \in R$. However, for our purpose such relations are too strong; e.g., the assignment

$i = 2$ is not equivalent to the sequence $i = 1; i = 2$ because, starting from $i = 0$, the former reaches $i = 2$ in one semantical step, whereas the latter cannot.

Hence, we need to alter strong bisimulation for our purposes. We do it, first, by removing the constraint that each step of one program is matched by exactly one step of the other one, and second, by requiring that our relation be bounded from above by a certain relation $O \subseteq \mathcal{T}_{Cfg} \times \mathcal{T}_{Cfg}$ called the *observation relation*.

Definition 3 (*O*-weak bisimulation). An *O*-weak bisimulation is a relation $R \subseteq O$ satisfying: for all $(\gamma_1, \gamma_2) \in R$,

- if $\gamma_1 \Rightarrow_S^T \gamma'_1$ then $\gamma'_1 \Rightarrow_S^* \gamma''_1$ and $\gamma_2 \Rightarrow_S^* \gamma''_2$, for some $(\gamma''_1, \gamma''_2) \in R$
- if $\gamma_2 \Rightarrow_S^T \gamma'_2$ then $\gamma_1 \Rightarrow_S^* \gamma''_1$ and $\gamma'_2 \Rightarrow_S^* \gamma''_2$ for some $(\gamma''_1, \gamma''_2) \in R$.

In the sequel we assume O to be an arbitrary, fixed parameter to our definitions. We omit it and only write "weak bisimulation" instead of "*O*-weak bisimulation". We now have our definition of program (actually, of configuration) equivalence:

Definition 4 (Configuration Equivalence). Configurations γ_1, γ_2 are equivalent, written $\gamma_1 \sim \gamma_2$, if there is a weak bisimulation R such that $(\gamma_1, \gamma_2) \in R$.

Example 3. The following configurations: $\gamma_1 \triangleq \langle \langle x = 2 \rangle_k \langle x \mapsto 0 \rangle_{env} \rangle_{cfg}$ and $\gamma'_1 \triangleq \langle \langle x = 1; x = x+1 \rangle_k \langle x \mapsto 0 \rangle_{env} \rangle_{cfg}$ are equivalent when O is defined by requiring that x has the same value in γ_1, γ_2 . The "witness" weak bisimulation R for the equivalence $\gamma_1 \sim \gamma'_1$ is defined by $\{(\gamma_1, \gamma'_1), (\gamma_2, \gamma_2)\}$, where $\gamma_2 \triangleq \langle \langle \cdot \rangle_k \langle x \mapsto 2 \rangle_{env} \rangle_{cfg}$.

The relation O gives us quite a lot of expressiveness for capturing various kinds of program equivalences. For example, *partial* equivalence [2] is: two programs are equivalent if, whenever presented with the same input, if they both terminate they produce the same output. This can be encoded by including cells in the configuration for the input and output, and by including in O the pairs of configurations satisfying: if their programs are both empty and their inputs are equal then their outputs are equal. Also, *full* equivalence from [2] is: two programs are equivalent if, whenever presented with the same input, they either both terminate and produce the same output, or they both do not terminate. This is captured by adding to the above relation all pairs of configurations from which there is an infinite execution starting from both configurations of the pair.

5 A Logic for Program Equivalence

We present in this section a logic for program equivalence. We first present the logic's syntax, then its semantics, and finally the notion of validity for formulas.

Definition 5 (Formulas). A formula is an expression of the form $\pi_1 \sim \pi_2$ if C where $\pi_1, \pi_2 \in T_{\Sigma, Cfg}(Var)$ are basic patterns and $C \in T_{\Sigma, Bool}(Var)$.

Example 4. Assume that the signature Σ for the language IMP contains a predicate $isModified : Id \times Stmt \rightarrow Bool$, expressing the fact that the value of the

given identifier is modified by the semantics of the given statement. A formula expressing the equivalence of the programs in Example 1 is

$$\begin{aligned} & \langle \langle \text{for } I \text{ from } A \text{ to } B \text{ do } \{ S \} \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \\ & \langle \langle I = A ; \text{while } I \leq B \text{ do } \{ S ; I = I + 1 \} \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \\ & \text{if not isModified}(I, S) \end{aligned}$$

where M a variable of sort Map . The condition says that the loop counter I is not modified in the loop body S . It is essential for the formula's validity.

We now define two semantics for formulas $f \triangleq \pi_1 \sim \pi_2$ if C . The first one, denoted by $\langle f \rangle$, is the set of pairs of configurations γ_1, γ_2 that satisfy, respectively, the patterns $\pi_1 \wedge C$ and $\pi_2 \wedge C$ by means of one valuation (the same valuation for both γ_1, γ_2). The second one, denoted by $\llbracket f \rrbracket$, excludes from $\langle f \rangle$ the pairs of configurations from which at least one component eventually leads to a deadlock.

Definition 6 (Semantics).

$$\begin{aligned} \langle f \rangle & \triangleq \{ (\gamma_1, \gamma_2) \mid \exists \rho : \text{Var} \rightarrow \mathcal{T}. (\gamma_i, \rho) \models \pi_i \wedge C, i = 1, 2 \}, \text{ and} \\ \llbracket f \rrbracket & \triangleq \{ (\gamma_1, \gamma_2) \in \langle f \rangle \mid \forall i \in \{1, 2\} \forall \gamma \in \mathcal{T}_{\text{Cfg}}. \gamma_i \Rightarrow^*_{\mathcal{S}} \gamma \text{ implies } \gamma \text{ is no deadlock} \}. \end{aligned}$$

We now define what it means for a formula f to be *valid*. Intuitively, we want to capture the idea that all configurations pairs $(\gamma_1, \gamma_2) \in \llbracket f \rrbracket$ satisfy $\gamma_1 \sim \gamma_2$ according to Definition 4. We use the $\llbracket \cdot \rrbracket$ semantics (not the $\langle \cdot \rangle$ one) because we are not interested in deadlocks. This is not really a restriction since deadlocks can be turned into final configurations by adding rules and, e.g., setting the content of some cell, say, **error**, to some value encoding the deadlock situation.

Definition 7 (Validity). A formula f is valid, written $\mathcal{S} \models f$, if $\llbracket f \rrbracket \neq \emptyset$ whenever $\langle f \rangle \neq \emptyset$, and for all $\gamma_1, \gamma_2 \in \llbracket f \rrbracket$, $\gamma_1 \sim \gamma_2$.

Note that f is (vacuously) valid if $\langle f \rangle = \emptyset$, and that f is not valid when $\langle f \rangle \neq \emptyset$ and $\llbracket f \rrbracket = \emptyset$ because in this case all the concrete configurations in $\langle f \rangle$ lead to deadlocks.

6 Auxiliary Operations: Derivatives and Conjunction

Our proof system consists in symbolically executing formulas according to the semantics of the language \mathcal{L} . This is achieved using the notion of *derivative*.

Definition 8 (Derivatives). Given a formula $g \triangleq \pi_1 \sim \pi_2$ if C , its derivatives are the formulas in the set $\Delta(g) = \Delta^l(g) \cup \Delta^r(g)$, where $\Delta^l(g), \Delta^r(g)$ are the smallest sets defined by: for each $(l \wedge C' \Rightarrow r) \in \mathcal{S}$, $\sigma^l \in U(\pi_1, l)$, $\sigma^r \in U(\pi_2, r)$:

- $(r\sigma^l \sim \pi_2$ if $(C \wedge C')\sigma^l \wedge \bigwedge \sigma^l) \in \Delta^l(g)$,
- $(\pi_1 \sim r\sigma^r$ if $(C \wedge C')\sigma^r \wedge \bigwedge \sigma^r) \in \Delta^r(g)$

where $\bigwedge \sigma \triangleq \bigwedge_{x \in \text{dom}(\sigma)} (x = \sigma(x))$, and $\text{dom}(\sigma)$ denotes the subset of the global set Var of variables where the substitution σ is not the identity. We naturally extend derivatives to sets F of formulas by $\Delta(F) = \bigcup_{f \in F} \Delta(f)$.

Remark 1. In Definition 8 we assume $var(l) \cap var(g) = \emptyset$, which can always be obtained by renaming the variables in the rewrite rule.

Example 5. Let B be a variable of sort *Bool* and S_1, S_2 be variables of sort *Stmt*. We consider the formula f below and compute its left-derivatives:

$$\langle\langle \text{if } B \text{ then } S_1 \text{ else } S_2 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle\langle \text{if } B' \text{ then } S_2 \text{ else } S_1 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \\ \text{if } B' = \neg B$$

The rules with a nonempty set of unifiers with the patterns in the formula are

$$\langle\langle \langle \text{if } \text{true} \text{ then } S'_1 \text{ else } S'_2 \rangle \curvearrowright S \rangle_k, \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow \langle\langle S'_1 \curvearrowright S \rangle_k, \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \\ \langle\langle \langle \text{if } \text{false} \text{ then } S'_1 \text{ else } S'_2 \rangle \curvearrowright S \rangle_k, \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \Rightarrow \langle\langle S'_2 \curvearrowright S \rangle_k, \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}}$$

The formula f has two left-derivatives, i.e., $\Delta^l(f)$ are the formulas in the set

$$\langle\langle S_1 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle\langle \text{if } B' \text{ then } S_2 \text{ else } S_1 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \text{ if } B' = \neg B \wedge B = \text{true} \\ \langle\langle S_2 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle\langle \text{if } B' \text{ then } S_2 \text{ else } S_1 \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \text{ if } B' = \neg B \wedge B = \text{false}$$

where $B = \text{true}$ and $B = \text{false}$ are induced by the symbolic unifiers: $B \mapsto \text{true}$, $S'_1 \mapsto S_1$, $S'_2 \mapsto S_2$, $M' \mapsto M$ and, respectively, $B \mapsto \text{false}$, $S'_1 \mapsto S_1$, $S'_2 \mapsto S_2$, $M' \mapsto M$. The superfluous equalities $S'_1 = S_1$, $S'_2 = S_2$, $M' = M$ were removed from conditions since S'_1 , S'_2 , and M' do not occur in the rest of the formula.

Another auxiliary operation used in our proof system is *conjunction* of formulas. We need it in order to compute the subsets of configuration pairs, denoted by formulas, which are included in the observation relation O (cf. Section 4).

Definition 9. For formulas $f : \pi_1 \sim \pi_2$ if C and $g : \pi'_1 \sim \pi'_2$ if C' , let $f \wedge g = \{\pi_1 \sigma_1 \sim \pi_2 \sigma_2 \text{ if } (C \wedge C')(\sigma_1 \cup \sigma_2) \wedge \bigwedge \sigma_1 \wedge \bigwedge \sigma_2 \mid \sigma_1 \in U(\pi_1, \pi'_1), \sigma_2 \in U(\pi_2, \pi'_2)\}$.

Example 6. Let f be the formula in Example 5 and let g denote the formula $\langle\langle P_1 \rangle_k, \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle\langle P_2 \rangle_k, \langle M'' \rangle_{\text{env}} \rangle_{\text{cfg}} \text{ if } M' = M''$. We denote by π_1, π'_1 and π_2, π'_2 their left and right-hand sides, respectively. Then, $U(\pi_1, \pi'_1)$ can be computed by matching, and consists of the unique substitution $\sigma_1 = (P_1 \mapsto \text{if } B \text{ then } S_1 \text{ else } S_2, M' \mapsto M)$. Similarly, $U(\pi_2, \pi'_2)$ consists of the substitution $\sigma_2 = (P_2 \mapsto \text{if } B' \text{ then } S_2 \text{ else } S_1, M'' \mapsto M)$. Thus, if we remove the conditions $M' = M''$, $\bigwedge \sigma_1$, and $\bigwedge \sigma_2$ (which are superfluous here since they constrain variables not occurring in the rest of the result), $f \wedge g$ is syntactically equal to f . This is consistent with the fact that \wedge is, semantically speaking, intersection, because we have $\llbracket f \rrbracket \subseteq \llbracket g \rrbracket$ and thus $\llbracket f \wedge g \rrbracket = \llbracket f \rrbracket \cap \llbracket g \rrbracket = \llbracket f \rrbracket$.

7 A Circular Proof System

In this section we define a three-rule proof system for proving program equivalence. It is inspired from *circular coinduction* [12], a coinductive proof technique for infinite data structures and coalgebras of expressions [16].

Remember that we have fixed an observation relation O . We assume a set of formulas Ω such that $\llbracket \Omega \rrbracket = O$. We also assume that for all $h \in \Omega$ and for all formula f , the conjunction $f \wedge h$ can be computed according to Definition 9:

Assumption 3 For all $(\pi_1 \sim \pi_2 \text{ if } C) \in \Omega$ and all $\pi \in T_{\Sigma, Cfg}(Var)$ with $(var(\pi_1) \cup var(\pi_2)) \cap var(\pi) = \emptyset$, there are two finite, possibly empty sets $U(\pi, \pi_1)$ and $U(\pi, \pi_2)$ of symbolic unifiers of π, π_1 and of π, π_2 , respectively.

Let also \vdash be an entailment relation satisfying $\mathcal{S}, F \vdash g$ implies $(\mathcal{S} \models g$ or $\llbracket g \rrbracket \subseteq \llbracket F \rrbracket)$. The set Ω and the relation \vdash are parameters of our proof system:

Definition 10 (Circular Proof System).

$$\begin{array}{l}
\text{[Axiom]} \quad \frac{}{\mathcal{S}, F \vdash^\circ \emptyset} \\
\text{[Reduce]} \quad \frac{\mathcal{S}, F \vdash g \quad \mathcal{S}, F \vdash^\circ G}{\mathcal{S}, F \vdash^\circ G \cup \{g\}} \\
\text{[Derive]} \quad \frac{\mathcal{S}, F \cup F' \vdash^\circ G \cup \Delta(g) \quad \mathcal{S}, g \wedge \Omega \vdash F'}{\mathcal{S}, F \vdash^\circ G \cup \{g\}} \text{ if } \Delta(g) \neq \emptyset
\end{array}$$

where $g \wedge \Omega$ denotes the set $\{g \wedge h \mid h \in \Omega\}$.

[Axiom] says that when an empty set of goals is reached, the proof is finished.

The [Reduce] rule says that if a given goal g from the current set of goals $G \cup \{g\}$ is discharged by the entailment \vdash then it is eliminated from the goals.

The last rule, [Derive], is the most complex. It says that any given goal g from the current set of goals, with a nonempty set $\Delta(g)$ of derivatives, can be replaced in the goals to be proved with the set $\Delta(g)$; and, simultaneously, any set of formulas F' that can be \vdash -entailed from $\mathcal{S}, g \wedge \Omega$ can be added as hypotheses. Note that the application of the [Derive] rule is nondeterministic in the choice of hypotheses F' , which depend on the parameters \vdash and Ω of the proof system.

Theorem 1 (soundness of \vdash°). Let Γ be a set of formulas such that $\llbracket \Gamma \rrbracket \subseteq \llbracket \Omega \rrbracket$ and for all $g \in \Gamma$, $\llbracket g \rrbracket \neq \emptyset$. If $\mathcal{S} \vdash^\circ \Gamma$ then $\mathcal{S} \models \Gamma$.

Note that we require $\llbracket \Gamma \rrbracket \subseteq \llbracket \Omega \rrbracket$ because otherwise the goals Γ have no chance of being valid. The assumption for all $g \in \Gamma$, $\llbracket g \rrbracket \neq \emptyset$ (that implies $\llbracket g \rrbracket \neq \emptyset$) is made for ensuring that g is not *vacuously* valid. Note also that initially, the set of hypotheses, denoted by F in the proof system, is empty: $\mathcal{S} \vdash^\circ \Gamma$ is $\mathcal{S}, \emptyset \vdash^\circ \Gamma$.

We now show that the circular proof system, when it terminates, always provides an answer (positive or negative) to the question $\mathcal{S} \models \Gamma$. Thus, in addition to soundness we have a *weak completeness* result. The result is "weak" because it assumes termination of the proof system. It ensures that we have a decision procedure for the equivalence of concrete, terminating programs.

In order to achieve weak completeness we need the following adaptations of Definition 8: we only keep the formulas with a *satisfiable condition*, i.e., we eliminate "empty" formulas f with $\llbracket f \rrbracket = \emptyset$. We also need additional assumptions. The first one says that non-derivable goals g that denote observationally equivalent configuration pairs are valid, and are discharged by the entailment \vdash . The second one says that deadlocks are not observationally equivalent to anything.

Assumption 4 For all formulas g such that $\llbracket g \rrbracket \subseteq \llbracket \Omega \rrbracket$ and $\Delta(g) = \emptyset$, $\mathcal{S} \vdash g$; and for all configurations γ_1, γ_2 , if γ_1 or γ_2 are deadlocks then $(\gamma_1, \gamma_2) \notin \llbracket \Omega \rrbracket$.

Theorem 2 (Weak Completeness of \vdash°). *Assume $S \models \Gamma$ and the proof system \vdash° terminates on Γ . Then, $S \vdash^\circ \Gamma$.*

Given a set of goals Γ , the proof system \vdash° may *terminate successfully* on it, which means it generates a tree that has at least one "empty" leaf (generated by [Axiom]). The proof system may also *terminate unsuccessfully* when it generates a finite tree and cannot expand it (i.e., it is blocked) and moreover that tree does not have any empty leaf. The proof system terminates on Γ if it terminates either successfully or unsuccessfully. Weak completeness thus says that if a set of goals is valid and the proof system terminates on it, then it terminates successfully.

Together, the soundness and weak completeness say that, if the proof system applied to a given set of goals terminates, then termination is successful if and only if the set of goals is valid. That is, when it terminates, the proof system correctly solves the program-equivalence problem. Of course, termination cannot be guaranteed, because the equivalence problem is undecidable. It does terminate on goals in which both programs terminate (because eventually the set of derivatives becomes empty) and also for goals in which one or both of the programs does not terminate, provided they behave in a certain regular way.

Example 7. We show the application of our proof system for proving the equivalence of our *for* and *while* programs formalised as the validity of the formula f (in which we assume for simplicity that the symbolic statement S is terminating; non-terminating statements can be handled as well but complicate the example):

$$\begin{aligned} & \langle \langle \text{for } I \text{ from } A \text{ to } B \text{ do}\{S\} \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \\ & \langle \langle I = A; \text{while } I \leq B \text{ do}\{S; I = I + 1\} \rangle_k, \langle M \rangle_{\text{env}} \rangle_{\text{cfg}} \\ & \text{if not isModified}(I, S) \end{aligned} \quad (2)$$

when the observation relation is denoted by the set $\Omega = \{ \langle \langle P_1 \rangle_k \langle M' \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle \langle P_2 \rangle_k \langle M'' \rangle_{\text{env}} \rangle_{\text{cfg}} \text{ if } M' = M'' \}$. The observation relation says that two configurations are observationally equivalent whenever they have equal environments.

The first applied rule is [Derive], which adds to the initially empty set of hypotheses the formula f , simultaneously replacing it in the goals with $\Delta(f)$. (f can be added to the hypotheses because $\langle f \rangle \subseteq \langle \Omega \rangle$, which implies $\Omega \wedge f \vdash f$).

After a certain number of applications of the [Derive] rule, the set of goals becomes (after some simplifications, which consist in removing goals with unsatisfiable conditions and logically simplifying the conditions of the remaining goals; note that A and B became (symbolic) values due to the **strict** attribute):

$$\begin{aligned} & \langle \langle \rangle_k, \langle \text{update}(M, I, A) \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \langle \langle \rangle_k, \langle \text{update}(M, I, A) \rangle_{\text{env}} \rangle_{\text{cfg}} \text{ if } A >_{\text{Int}} B \\ & \langle \langle \text{for } I \text{ from } A +_{\text{Int}} 1 \text{ to } B \text{ do}\{S\} \rangle_k, \langle \text{followup}(S, \text{update}(M, I, A)) \rangle_{\text{env}} \rangle_{\text{cfg}} \sim \\ & \langle \langle I = A +_{\text{Int}} 1; \text{while } I \leq B \text{ do}\{S; I = I + 1\} \rangle_k, \langle \text{followup}(S, \text{update}(M, I, A)) \rangle_{\text{env}} \rangle_{\text{cfg}} \\ & \text{if not isModified}(I, S) \wedge A \leq_{\text{Int}} B \end{aligned}$$

where $\text{followup}(S, M)$ denotes the effect of executing statement S on map M . Recall that S is terminating, so $\text{followup}(S, M)$ is defined. The fact that I is

not modified by S is expressed by the equation $\text{followup}(S, \text{update}(M, I, V)) = \text{update}(M, I, V)$, assumed to hold in the domain of maps for all concrete instances of S that do not modify I . Moreover, for each terminating concrete instance P of S , $\text{followup}(P, M) = M'$ iff $\langle\langle P \rangle_k, \langle M \rangle_{\text{env}}\rangle_{\text{cfg}} \Rightarrow_{\mathcal{S}}^* \langle\langle \rangle_k, \langle M' \rangle_{\text{env}}\rangle_{\text{cfg}}$.

The first goal is discharged by the [Reduce] rule (based on the fact that the \vdash relation "knows" that goals with same left and right-hand side are valid). The second goal f' is actually an *instance* of the first one: i.e., $\langle\langle f' \rangle\rangle \subseteq \langle\langle f \rangle\rangle$ since any concrete instance of f' is also a concrete instance of f . Thus, $\mathcal{S}, f \vdash f'$, and since f was added to the set hypotheses by the first application of [Derive], f' is eliminated by the [Reduce] rule. The set of goals to be proved is now empty; the proof system has terminated successfully, meaning that the formula f is valid.

8 Conclusion and Future Work

We have presented a definition for program equivalence, a logic that encodes this definition in its formulas, and a proof system for the logic, which is proved sound and weakly complete. A prototype implementation for the proof system in the \mathbb{K} framework was also presented and illustrated on a simple but paradigmatic example of equivalent programs in a language IMP defined in the \mathbb{K} framework.

The proposed approach is general: it does not depend on \mathbb{K} and IMP but only requires a formal semantics of the language of interest as a term-rewriting system. The chosen equivalence relation is a weak bisimulation, which is parametric in a certain observation relation. We show the approach is applicable for concrete and symbolic programs, as well as for terminating and non-terminating ones.

A prototype that implements the approach for the IMP language has been developed in the \mathbb{K} Framework. The implementation can be tested using the on-line interface³ of the \mathbb{K} tool.

Future Work We are currently applying our deductive system for proving the correctness of a compiler between two languages (as part of another project we are involved in). The source language is a stack-based language with control structures (loops, conditionals, dynamical function definitions). The target is also stack-based but only has (possibly, conditional) jumps. The correctness of the compiler amounts to proving the equivalence of several pairs of symbolic programs; in each pair, one component denotes a source-language control structure, and the other component is the translation of that control structure in the target language using jumps. We are also planning to combine our program-equivalence verification with matching logic [14] verification. Matching logic is an automatic, language-independent formal verification framework for languages with a rewrite-based semantics. The idea is to prove matching logic properties on programs in the source language, and guarantee, via the compiler's correctness that the compiled programs in the target language satisfy those properties as well.

³ <http://fmse.info.uaic.ro/tools/K/?tree=examples/prog-equiv/peq.k>

References

1. Sudipta Kundu, Zachary Tatlock, and Sorin Lerner. Proving optimizations correct using parameterized program equivalence. In *Programming Languages Design and Implementation*, pages 327–337, 2009.
2. Benny Godlin and Ofer Strichman. Inference rules for proving the equivalence of recursive procedures. *Acta Inf.*, 45(6):403–439, 2008.
3. Benny Godlin and Ofer Strichman. Regression verification: proving the equivalence of similar programs. *Software Testing, Verification and Reliability*, 2012. 10.1002/stvr.1472.
4. Sagar Chaki, Arie Gurfinkel, and Ofer Strichman. Regression verification for multi-threaded programs. In Viktor Kuncak and Andrey Rybalchenko, editors, *VMCAI*, volume 7148 of *Lecture Notes in Computer Science*, pages 119–135. Springer, 2012.
5. Xavier Leroy. Formal verification of a realistic compiler. *Comm. ACM*, 52(7):107–115, 2009.
6. George C. Necula. Translation validation for an optimizing compiler. In Monica S. Lam, editor, *PLDI*, pages 83–94. ACM, 2000.
7. Andrew M. Pitts. Operational semantics and program equivalence. In *Applied Semantics, International Summer School, APPSEM 2000, Caminha, Portugal, September 9-15, 2000, Advanced Lectures*, pages 378–412, London, UK, UK, 2002. Springer-Verlag.
8. Tamarah Arons, Elad Elster, Limor Fix, Sela Mador-Haim, Michael Mishaeli, Jonathan Shalev, Eli Singerman, Andreas Tiemeyer, Moshe Y. Vardi, and Lenore D. Zuck. Formal verification of backward compatibility of microcode. In *Computer-Aided Verification*, pages 185–198, 2005.
9. Sorin Craciunescu. Proving the equivalence of CLP programs. In *International Conference of Logic Programming, LNCS 2401*, pages 287–301, 2002.
10. Shuvendu K. Lahiri, Chris Hawblitzel, Ming Kawaguchi, and Henrique Rebêlo. Symdiff: A language-agnostic semantic diff tool for imperative programs. In *Computer Aided Verification, LNCS 7358*, pages 712–717, 2012.
11. Fabio Somenzi and Andreas Kuehlmann. *Electronic Design Automation For Integrated Circuits Handbook*, volume 2, chapter 4: Equivalence Checking. Taylor & Francis, 2006.
12. Grigore Roşu and Dorel Lucanu. Circular coinduction – a proof theoretical foundation. In *CALCO 2009*, volume 5728 of *LNCS*, pages 127–144. Springer, 2009.
13. G. Roşu and T.-F. Şerbănuţă. An Overview of the K Semantic Framework. *Journal of Logic and Algebraic Programming*, 79(6):397–434, 2010.
14. Grigore Roşu and Andrei Stefanescu. Checking reachability using matching logic. In *Proceedings of the 27th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA’12)*. ACM, 2012. To appear.
15. Andrei Arusoai, Dorel Lucanu, and Vlad Rusu. A Generic Approach to Symbolic Execution. Research Report RR-8189, INRIA. <http://hal.inria.fr/hal-00766220/>.
16. M. Bonsangue, G. Caltais, E. Goriac, D. Lucanu, Jan J. M. M. Rutten, and A. Silva. A decision procedure for bisimilarity of generalized regular expressions. In *Proc. 13th Brazilian Symposium on Formal Methods, LNCS 6527*, pages 226–241, 2011.